

# **StarPU Handbook**

---

for StarPU 1.1.0

---

This manual documents the usage of StarPU version 1.1.0. Its contents was last updated on 09 December 2013.

Copyright © 2009–2013 Université de Bordeaux 1

Copyright © 2010-2013 Centre National de la Recherche Scientifique

Copyright © 2011, 2012 Institut National de Recherche en Informatique et Automatique

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.



# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Motivation . . . . .	3
<b>I</b>	<b>Using StarPU</b>	<b>5</b>
<b>2</b>	<b>Building and Installing StarPU</b>	<b>7</b>
2.1	Installing a Binary Package . . . . .	7
2.2	Installing from Source . . . . .	7
2.2.1	Optional Dependencies . . . . .	7
2.2.2	Getting Sources . . . . .	7
2.2.3	Configuring StarPU . . . . .	8
2.2.4	Building StarPU . . . . .	8
2.2.5	Installing StarPU . . . . .	8
2.3	Setting up Your Own Code . . . . .	8
2.3.1	Setting Flags for Compiling, Linking and Running Applications . . . . .	8
2.3.2	Running a Basic StarPU Application . . . . .	9
2.3.3	Kernel Threads Started by StarPU . . . . .	9
2.3.4	Enabling OpenCL . . . . .	9
2.4	Benchmarking StarPU . . . . .	9
2.4.1	Task Size Overhead . . . . .	10
2.4.2	Data Transfer Latency . . . . .	10
2.4.3	Matrix-Matrix Multiplication . . . . .	10
2.4.4	Cholesky Factorization . . . . .	10
2.4.5	LU Factorization . . . . .	10
<b>3</b>	<b>Basic Examples</b>	<b>11</b>
3.1	Hello World Using The C Extension . . . . .	11
3.2	Hello World Using StarPU's API . . . . .	12
3.2.1	Required Headers . . . . .	12
3.2.2	Defining A Codelet . . . . .	12
3.2.3	Submitting A Task . . . . .	12
<b>4</b>	<b>Advanced Examples</b>	<b>13</b>
4.1	Using Multiple Implementations Of A Codelet . . . . .	13
4.2	Enabling Implementation According To Capabilities . . . . .	13
4.3	Task And Worker Profiling . . . . .	14
4.4	Partitioning Data . . . . .	15
4.5	Performance Model Example . . . . .	16
4.6	Theoretical Lower Bound On Execution Time Example . . . . .	17
4.7	Insert Task Utility . . . . .	18
4.8	Data Reduction . . . . .	19
4.9	Temporary Buffers . . . . .	20
4.9.1	Temporary Data . . . . .	20
4.9.2	Scratch Data . . . . .	21
4.10	Parallel Tasks . . . . .	21
4.10.1	Fork-mode Parallel Tasks . . . . .	21

4.10.2 SPMD-mode Parallel Tasks . . . . .	22
4.10.3 Parallel Tasks Performance . . . . .	22
4.10.4 Combined Workers . . . . .	22
4.10.5 Concurrent Parallel Tasks . . . . .	22
4.11 Debugging . . . . .	23
4.12 The Multiformat Interface . . . . .	23
4.13 Using The Driver API . . . . .	24
4.14 Defining A New Scheduling Policy . . . . .	24
4.15 On-GPU Rendering . . . . .	25
4.16 Defining A New Data Interface . . . . .	25
4.17 Setting The Data Handles For A Task . . . . .	26
4.18 More Examples . . . . .	27
<b>5 How To Optimize Performance With StarPU</b>	<b>29</b>
5.1 Data Management . . . . .	29
5.2 Task Granularity . . . . .	30
5.3 Task Submission . . . . .	30
5.4 Task Priorities . . . . .	30
5.5 Task Scheduling Policy . . . . .	30
5.6 Performance Model Calibration . . . . .	31
5.7 Task Distribution Vs Data Transfer . . . . .	33
5.8 Data Prefetch . . . . .	33
5.9 Power-based Scheduling . . . . .	33
5.10 Static Scheduling . . . . .	33
5.11 Profiling . . . . .	34
5.12 Detection Stuck Conditions . . . . .	34
5.13 CUDA-specific Optimizations . . . . .	34
5.14 Performance Debugging . . . . .	34
5.15 Simulated Performance . . . . .	35
5.15.1 Calibration . . . . .	35
5.15.2 Simulation . . . . .	35
5.15.3 Simulation On Another Machine . . . . .	36
<b>6 Performance Feedback</b>	<b>37</b>
6.1 Using The Temanejo Task Debugger . . . . .	37
6.2 On-line Performance Feedback . . . . .	37
6.2.1 Enabling On-line Performance Monitoring . . . . .	37
6.2.2 Per-task Feedback . . . . .	37
6.2.3 Per-codelet Feedback . . . . .	38
6.2.4 Per-worker Feedback . . . . .	38
6.2.5 Bus-related Feedback . . . . .	38
<b>7 Tips and Tricks To Know About</b>	<b>39</b>
7.1 How To Initialize A Computation Library Once For Each Worker? . . . . .	39
7.2 How to limit memory per node . . . . .	39
7.3 Thread Binding on NetBSD . . . . .	40
7.4 Using StarPU With MKL 11 (Intel Composer XE 2013) . . . . .	40
<b>8 MPI Support</b>	<b>41</b>
8.1 Simple Example . . . . .	41
8.2 Point To Point Communication . . . . .	42
8.3 Exchanging User Defined Data Interface . . . . .	42
8.4 MPI Insert Task Utility . . . . .	43
8.5 MPI Data migration . . . . .	44
8.6 MPI Collective Operations . . . . .	45
<b>9 FFT Support</b>	<b>47</b>
9.1 Compilation . . . . .	47

<b>10 C Extensions</b>	<b>49</b>
10.1 Defining Tasks . . . . .	49
<b>11 SOCL OpenCL Extensions</b>	<b>51</b>
<b>12 Scheduling Contexts</b>	<b>53</b>
12.1 General Ideas . . . . .	53
12.2 Creating A Context . . . . .	53
12.3 Modifying A Context . . . . .	54
12.4 Submitting Tasks To A Context . . . . .	54
12.5 Deleting A Context . . . . .	54
12.6 Emptying A Context . . . . .	54
12.7 Contexts Sharing Workers . . . . .	55
<b>13 Scheduling Context Hypervisor</b>	<b>57</b>
13.1 What Is The Hypervisor . . . . .	57
13.2 Start the Hypervisor . . . . .	57
13.3 Interrogate The Runtime . . . . .	57
13.4 Trigger the Hypervisor . . . . .	57
13.5 Resizing Strategies . . . . .	58
13.6 a new hypervisor policy . . . . .	59
<b>II Inside StarPU</b>	<b>61</b>
<b>14 Execution Configuration Through Environment Variables</b>	<b>63</b>
14.1 Configuring Workers . . . . .	63
14.2 Configuring The Scheduling Engine . . . . .	64
14.3 Extensions . . . . .	65
14.4 Miscellaneous And Debug . . . . .	65
14.5 Configuring The Hypervisor . . . . .	66
<b>15 Compilation Configuration</b>	<b>69</b>
15.1 Common Configuration . . . . .	69
15.2 Configuring Workers . . . . .	69
15.3 Extension Configuration . . . . .	70
15.4 Advanced Configuration . . . . .	71
<b>16 Module Index</b>	<b>73</b>
16.1 Modules . . . . .	73
<b>17 Module Documentation a.k.a StarPU's API</b>	<b>75</b>
17.1 Versioning . . . . .	75
17.1.1 Detailed Description . . . . .	75
17.1.2 Macro Definition Documentation . . . . .	75
STARPU_MAJOR_VERSION . . . . .	75
STARPU_MINOR_VERSION . . . . .	75
STARPU_RELEASE_VERSION . . . . .	75
17.1.3 Function Documentation . . . . .	75
starpu_get_version . . . . .	75
17.2 Initialization and Termination . . . . .	75
17.2.1 Detailed Description . . . . .	76
17.2.2 Data Structure Documentation . . . . .	76
struct starpu_driver . . . . .	76
union starpu_driver.id . . . . .	76
struct starpu_vector_interface . . . . .	76
struct starpu_conf . . . . .	76
17.2.3 Function Documentation . . . . .	78
starpu_init . . . . .	78
starpu_conf_init . . . . .	79

starpu_shutdown . . . . .	79
starpu_asynchronous_copy_disabled . . . . .	79
starpu_asynchronous_cuda_copy_disabled . . . . .	79
starpu_asynchronous_opencl_copy_disabled . . . . .	79
starpu_topology_print . . . . .	79
<b>17.3 Standard Memory Library . . . . .</b>	<b>79</b>
<b>17.3.1 Detailed Description . . . . .</b>	<b>79</b>
<b>17.3.2 Macro Definition Documentation . . . . .</b>	<b>79</b>
starpu_data_malloc_pinned_if_possible . . . . .	79
starpu_data_free_pinned_if_possible . . . . .	79
STARPU_MALLOC_PINNED . . . . .	80
STARPU_MALLOC_COUNT . . . . .	80
<b>17.3.3 Function Documentation . . . . .</b>	<b>80</b>
starpu_malloc_flags . . . . .	80
starpu_malloc_set_align . . . . .	80
starpu_malloc . . . . .	80
starpu_free . . . . .	80
starpu_free_flags . . . . .	80
starpu_memory_get_available . . . . .	80
<b>17.4 Toolbox . . . . .</b>	<b>80</b>
<b>17.4.1 Detailed Description . . . . .</b>	<b>81</b>
<b>17.4.2 Macro Definition Documentation . . . . .</b>	<b>81</b>
STARPU_GNUC_PREREQ . . . . .	81
STARPU_UNLIKELY . . . . .	81
STARPU_LIKELY . . . . .	81
STARPU_ATTRIBUTE_UNUSED . . . . .	81
STARPU_ATTRIBUTE_INTERNAL . . . . .	81
STARPU_ATTRIBUTE_MALLOC . . . . .	81
STARPU_ATTRIBUTE_WARN_UNUSED_RESULT . . . . .	81
STARPU_ATTRIBUTE_PURE . . . . .	81
STARPU_ATTRIBUTE_ALIGNED . . . . .	81
STARPU_WARN_UNUSED_RESULT . . . . .	82
STARPU_POISON_PTR . . . . .	82
STARPU_MIN . . . . .	82
STARPU_MAX . . . . .	82
STARPU_ASSERT . . . . .	82
STARPU_ASSERT_MSG . . . . .	82
STARPU_ABORT . . . . .	82
STARPU_ABORT_MSG . . . . .	82
STARPU_CHECK_RETURN_VALUE . . . . .	82
STARPU_CHECK_RETURN_VALUE_IS . . . . .	82
STARPU_RMB . . . . .	82
STARPU_WMB . . . . .	82
<b>17.4.3 Function Documentation . . . . .</b>	<b>82</b>
starpu_get_env_number . . . . .	82
<b>17.5 Threads . . . . .</b>	<b>82</b>
<b>17.5.1 Detailed Description . . . . .</b>	<b>84</b>
<b>17.5.2 Macro Definition Documentation . . . . .</b>	<b>84</b>
STARPU_PTHREAD_CREATE_ON . . . . .	84
STARPU_PTHREAD_CREATE . . . . .	84
STARPU_PTHREAD_MUTEX_INIT . . . . .	84
STARPU_PTHREAD_MUTEX_DESTROY . . . . .	84
STARPU_PTHREAD_MUTEX_LOCK . . . . .	84
STARPU_PTHREAD_MUTEX_UNLOCK . . . . .	84
STARPU_PTHREAD_KEY_CREATE . . . . .	84
STARPU_PTHREAD_KEY_DELETE . . . . .	84
STARPU_PTHREAD_SETSPECIFIC . . . . .	84
STARPU_PTHREAD_GETSPECIFIC . . . . .	84

STARPU_PTHREAD_RWLOCK_INIT . . . . .	84
STARPU_PTHREAD_RWLOCK_RDLOCK . . . . .	84
STARPU_PTHREAD_RWLOCK_WRLOCK . . . . .	84
STARPU_PTHREAD_RWLOCK_UNLOCK . . . . .	85
STARPU_PTHREAD_RWLOCK_DESTROY . . . . .	85
STARPU_PTHREAD_COND_INIT . . . . .	85
STARPU_PTHREAD_COND_DESTROY . . . . .	85
STARPU_PTHREAD_COND_SIGNAL . . . . .	85
STARPU_PTHREAD_COND_BROADCAST . . . . .	85
STARPU_PTHREAD_COND_WAIT . . . . .	85
STARPU_PTHREAD_BARRIER_INIT . . . . .	85
STARPU_PTHREAD_BARRIER_DESTROY . . . . .	85
STARPU_PTHREAD_BARRIER_WAIT . . . . .	85
STARPU_PTHREAD_MUTEX_INITIALIZER . . . . .	85
STARPU_PTHREAD_COND_INITIALIZER . . . . .	85
<b>17.5.3 Function Documentation . . . . .</b>	<b>85</b>
starpu_pthread_create . . . . .	85
starpu_pthread_join . . . . .	85
starpu_pthread_attr_init . . . . .	86
starpu_pthread_attr_destroy . . . . .	86
starpu_pthread_attr_setdetachstate . . . . .	86
starpu_pthread_mutex_init . . . . .	86
starpu_pthread_mutex_destroy . . . . .	86
starpu_pthread_mutex_lock . . . . .	86
starpu_pthread_mutex_unlock . . . . .	86
starpu_pthread_mutex_trylock . . . . .	86
starpu_pthread_key_create . . . . .	86
starpu_pthread_key_delete . . . . .	86
starpu_pthread_setspecific . . . . .	86
starpu_pthread_getspecific . . . . .	87
starpu_pthread_cond_init . . . . .	87
starpu_pthread_cond_signal . . . . .	87
starpu_pthread_cond_broadcast . . . . .	87
starpu_pthread_cond_wait . . . . .	87
starpu_pthread_cond_timedwait . . . . .	87
starpu_pthread_cond_destroy . . . . .	87
starpu_pthread_rwlock_init . . . . .	87
starpu_pthread_rwlock_destroy . . . . .	87
starpu_pthread_rwlock_rdlock . . . . .	87
starpu_pthread_rwlock_wrlock . . . . .	87
starpu_pthread_rwlock_unlock . . . . .	87
<b>17.6 Workers' Properties . . . . .</b>	<b>88</b>
<b>17.6.1 Detailed Description . . . . .</b>	<b>88</b>
<b>17.6.2 Data Structure Documentation . . . . .</b>	<b>88</b>
struct starpu_worker_collection . . . . .	88
struct starpu_sched_ctx_iterator . . . . .	89
<b>17.6.3 Macro Definition Documentation . . . . .</b>	<b>89</b>
STARPU_NMAXWORKERS . . . . .	89
<b>17.6.4 Enumeration Type Documentation . . . . .</b>	<b>89</b>
starpu_node_kind . . . . .	89
starpu_worker_archtype . . . . .	90
starpu_worker_collection_type . . . . .	90
<b>17.6.5 Function Documentation . . . . .</b>	<b>90</b>
starpu_worker_get_count . . . . .	90
starpu_worker_get_count_by_type . . . . .	90
starpu_cpu_worker_get_count . . . . .	90
starpu_cuda_worker_get_count . . . . .	90
starpu_opencl_worker_get_count . . . . .	90

starpu_worker_get_id . . . . .	90
starpu_worker_get_ids_by_type . . . . .	90
starpu_worker_get_by_type . . . . .	91
starpu_worker_get_by_devid . . . . .	91
starpu_worker_get_devid . . . . .	91
starpu_worker_get_type . . . . .	91
starpu_worker_get_name . . . . .	91
starpu_worker_get_memory_node . . . . .	91
starpu_node_get_kind . . . . .	91
<b>17.7 Data Management . . . . .</b>	<b>91</b>
<b>17.7.1 Detailed Description . . . . .</b>	<b>92</b>
<b>17.7.2 Macro Definition Documentation . . . . .</b>	<b>93</b>
<b>STARPU_DATA_ACQUIRE_CB . . . . .</b>	<b>93</b>
<b>17.7.3 Typedef Documentation . . . . .</b>	<b>93</b>
<b>starpu_data_handle_t . . . . .</b>	<b>93</b>
<b>17.7.4 Enumeration Type Documentation . . . . .</b>	<b>93</b>
<b>starpu_data_access_mode . . . . .</b>	<b>93</b>
<b>17.7.5 Function Documentation . . . . .</b>	<b>93</b>
<b>starpu_data_register . . . . .</b>	<b>93</b>
<b>starpu_data_register_same . . . . .</b>	<b>93</b>
<b>starpu_data_unregister . . . . .</b>	<b>94</b>
<b>starpu_data_unregister_no_coherency . . . . .</b>	<b>94</b>
<b>starpu_data_unregister_submit . . . . .</b>	<b>94</b>
<b>starpu_data_invalidate . . . . .</b>	<b>94</b>
<b>starpu_data_invalidate_submit . . . . .</b>	<b>94</b>
<b>starpu_data_set_wt_mask . . . . .</b>	<b>94</b>
<b>starpu_data_prefetch_on_node . . . . .</b>	<b>94</b>
<b>starpu_data_lookup . . . . .</b>	<b>94</b>
<b>starpu_data_request_allocation . . . . .</b>	<b>94</b>
<b>starpu_data_query_status . . . . .</b>	<b>94</b>
<b>starpu_data_advise_as_important . . . . .</b>	<b>94</b>
<b>starpu_data_set_reduction_methods . . . . .</b>	<b>95</b>
<b>starpu_data_acquire . . . . .</b>	<b>95</b>
<b>starpu_data_acquire_cb . . . . .</b>	<b>95</b>
<b>starpu_data_acquire_on_node . . . . .</b>	<b>95</b>
<b>starpu_data_acquire_on_node_cb . . . . .</b>	<b>95</b>
<b>starpu_data_release . . . . .</b>	<b>95</b>
<b>starpu_data_release_on_node . . . . .</b>	<b>95</b>
<b>17.8 Data Interfaces . . . . .</b>	<b>95</b>
<b>17.8.1 Detailed Description . . . . .</b>	<b>99</b>
<b>17.8.2 Data Structure Documentation . . . . .</b>	<b>99</b>
<b>struct starpu_data_interface_ops . . . . .</b>	<b>99</b>
<b>struct starpu_data_copy_methods . . . . .</b>	<b>100</b>
<b>struct starpu_variable_interface . . . . .</b>	<b>102</b>
<b>struct starpu_vector_interface . . . . .</b>	<b>103</b>
<b>struct starpu_matrix_interface . . . . .</b>	<b>103</b>
<b>struct starpu_block_interface . . . . .</b>	<b>103</b>
<b>struct starpu_bcsr_interface . . . . .</b>	<b>103</b>
<b>struct starpu_csr_interface . . . . .</b>	<b>104</b>
<b>struct starpu_coo_interface . . . . .</b>	<b>104</b>
<b>17.8.3 Macro Definition Documentation . . . . .</b>	<b>104</b>
<b>STARPU_VARIABLE_GET_PTR . . . . .</b>	<b>104</b>
<b>STARPU_VARIABLE_GET_ELEMSIZE . . . . .</b>	<b>104</b>
<b>STARPU_VARIABLE_GET_DEV_HANDLE . . . . .</b>	<b>104</b>
<b>STARPU_VARIABLE_GET_OFFSET . . . . .</b>	<b>104</b>
<b>STARPU_VECTOR_GET_PTR . . . . .</b>	<b>104</b>
<b>STARPU_VECTOR_GET_DEV_HANDLE . . . . .</b>	<b>105</b>
<b>STARPU_VECTOR_GET_OFFSET . . . . .</b>	<b>105</b>

STARPU_VECTOR_GET_NX . . . . .	105
STARPU_VECTOR_GET_ELEMSIZE . . . . .	105
STARPU_MATRIX_GET_PTR . . . . .	105
STARPU_MATRIX_GET_DEV_HANDLE . . . . .	105
STARPU_MATRIX_GET_OFFSET . . . . .	105
STARPU_MATRIX_GET_NX . . . . .	105
STARPU_MATRIX_GET_NY . . . . .	105
STARPU_MATRIX_GET_LD . . . . .	105
STARPU_MATRIX_GET_ELEMSIZE . . . . .	105
STARPU_BLOCK_GET_PTR . . . . .	105
STARPU_BLOCK_GET_DEV_HANDLE . . . . .	105
STARPU_BLOCK_GET_OFFSET . . . . .	105
STARPU_BLOCK_GET_NX . . . . .	106
STARPU_BLOCK_GET_NY . . . . .	106
STARPU_BLOCK_GET_NZ . . . . .	106
STARPU_BLOCK_GET_LDY . . . . .	106
STARPU_BLOCK_GET_LDZ . . . . .	106
STARPU_BLOCK_GET_ELEMSIZE . . . . .	106
STARPU_BCSR_GET_NNZ . . . . .	106
STARPU_BCSR_GET_NZVAL . . . . .	106
STARPU_BCSR_GET_NZVAL_DEV_HANDLE . . . . .	106
STARPU_BCSR_GET_COLIND . . . . .	106
STARPU_BCSR_GET_COLIND_DEV_HANDLE . . . . .	106
STARPU_BCSR_GET_ROWPTR . . . . .	106
STARPU_BCSR_GET_ROWPTR_DEV_HANDLE . . . . .	106
STARPU_BCSR_GET_OFFSET . . . . .	106
STARPU_CSR_GET_NNZ . . . . .	107
STARPU_CSR_GET_NROW . . . . .	107
STARPU_CSR_GET_NZVAL . . . . .	107
STARPU_CSR_GET_NZVAL_DEV_HANDLE . . . . .	107
STARPU_CSR_GET_COLIND . . . . .	107
STARPU_CSR_GET_COLIND_DEV_HANDLE . . . . .	107
STARPU_CSR_GET_ROWPTR . . . . .	107
STARPU_CSR_GET_ROWPTR_DEV_HANDLE . . . . .	107
STARPU_CSR_GET_OFFSET . . . . .	107
STARPU_CSR_GET_FIRSTENTRY . . . . .	107
STARPU_CSR_GET_ELEMSIZE . . . . .	107
STARPU_COO_GET_COLUMNS . . . . .	107
STARPU_COO_GET_COLUMNS_DEV_HANDLE . . . . .	107
STARPU_COO_GET_ROWS . . . . .	107
STARPU_COO_GET_ROWS_DEV_HANDLE . . . . .	108
STARPU_COO_GET_VALUES . . . . .	108
STARPU_COO_GET_VALUES_DEV_HANDLE . . . . .	108
STARPU_COO_GET_OFFSET . . . . .	108
STARPU_COO_GET_NX . . . . .	108
STARPU_COO_GET_NY . . . . .	108
STARPU_COO_GET_NVALUES . . . . .	108
STARPU_COO_GET_ELEMSIZE . . . . .	108
17.8.4 Enumeration Type Documentation . . . . .	108
starpu.data.interface.id . . . . .	108
17.8.5 Function Documentation . . . . .	109
starpu_void_data_register . . . . .	109
starpu_variable_data_register . . . . .	109
starpu_vector_data_register . . . . .	109
starpu_matrix_data_register . . . . .	109
starpu_block_data_register . . . . .	109
starpu_bcsr_data_register . . . . .	110
starpu_csr_data_register . . . . .	110

starpu_coo_data_register . . . . .	110
starpu_data_get_interface_on_node . . . . .	110
starpu_data_handle_to_pointer . . . . .	110
starpu_data_get_local_ptr . . . . .	110
starpu_data_get_interface_id . . . . .	110
starpu_data_get_size . . . . .	110
starpu_data_pack . . . . .	110
starpu_data_unpack . . . . .	110
starpu_variable_get_elemsize . . . . .	111
starpu_variable_get_local_ptr . . . . .	111
starpu_vector_get_nx . . . . .	111
starpu_vector_get_elemsize . . . . .	111
starpu_vector_get_local_ptr . . . . .	111
starpu_matrix_get_nx . . . . .	111
starpu_matrix_get_ny . . . . .	111
starpu_matrix_get_local_id . . . . .	111
starpu_matrix_get_local_ptr . . . . .	111
starpu_matrix_get_elemsize . . . . .	111
starpu_block_get_nx . . . . .	111
starpu_block_get_ny . . . . .	111
starpu_block_get_nz . . . . .	111
starpu_block_get_local_ldy . . . . .	111
starpu_block_get_local_ldz . . . . .	112
starpu_block_get_local_ptr . . . . .	112
starpu_block_get_elemsize . . . . .	112
starpu_bcsr_get_nnz . . . . .	112
starpu_bcsr_get_nrow . . . . .	112
starpu_bcsr_get_firstentry . . . . .	112
starpu_bcsr_get_local_nzval . . . . .	112
starpu_bcsr_get_local_colind . . . . .	112
starpu_bcsr_get_local_rowptr . . . . .	112
starpu_bcsr_get_r . . . . .	112
starpu_bcsr_get_c . . . . .	112
starpu_bcsr_get_elemsize . . . . .	112
starpu_csr_get_nnz . . . . .	112
starpu_csr_get_nrow . . . . .	112
starpu_csr_get_firstentry . . . . .	112
starpu_csr_get_local_nzval . . . . .	113
starpu_csr_get_local_colind . . . . .	113
starpu_csr_get_local_rowptr . . . . .	113
starpu_csr_get_elemsize . . . . .	113
starpu_malloc_on_node . . . . .	113
starpu_free_on_node . . . . .	113
starpu_interface_copy . . . . .	113
starpu_hash_crc32c_be_n . . . . .	113
starpu_hash_crc32c_be . . . . .	113
starpu_hash_crc32c_string . . . . .	113
starpu_data_interface_get_next_id . . . . .	113
17.9 Data Partition . . . . .	113
17.9.1 Detailed Description . . . . .	115
17.9.2 Data Structure Documentation . . . . .	115
struct starpu_data_filter . . . . .	115
17.9.3 Function Documentation . . . . .	115
starpu_data_partition . . . . .	115
starpu_data_unpartition . . . . .	116
starpu_data_get_nb_children . . . . .	116
starpu_data_get_child . . . . .	116
starpu_data_get_sub_data . . . . .	116

starpu_data_vget_sub_data . . . . .	116
starpu_data_map_filters . . . . .	116
starpu_data_vmap_filters . . . . .	116
starpu_vector_filter_block . . . . .	116
starpu_vector_filter_block_shadow . . . . .	116
starpu_vector_filter_list . . . . .	117
starpu_vector_filter_divide_in_2 . . . . .	117
starpu_matrix_filter_block . . . . .	117
starpu_matrix_filter_block_shadow . . . . .	117
starpu_matrix_filter_vertical_block . . . . .	117
starpu_matrix_filter_vertical_block_shadow . . . . .	117
starpu_block_filter_block . . . . .	117
starpu_block_filter_block_shadow . . . . .	117
starpu_block_filter_vertical_block . . . . .	118
starpu_block_filter_vertical_block_shadow . . . . .	118
starpu_block_filter_depth_block . . . . .	118
starpu_block_filter_depth_block_shadow . . . . .	118
starpu_bcsr_filter_canonical_block . . . . .	118
starpu_csr_filter_vertical_block . . . . .	118
<b>17.10 Multiformat Data Interface . . . . .</b>	<b>118</b>
<b>17.10.1 Detailed Description . . . . .</b>	<b>119</b>
<b>17.10.2 Data Structure Documentation . . . . .</b>	<b>119</b>
<b>struct starpu_multiformat_data_interface_ops . . . . .</b>	<b>119</b>
<b>struct starpu_multiformat_interface . . . . .</b>	<b>120</b>
<b>17.10.3 Macro Definition Documentation . . . . .</b>	<b>120</b>
<b>STARPU_MULTIFORMAT_GET_CPU_PTR . . . . .</b>	<b>120</b>
<b>STARPU_MULTIFORMAT_GET_CUDA_PTR . . . . .</b>	<b>120</b>
<b>STARPU_MULTIFORMAT_GET_OPENCL_PTR . . . . .</b>	<b>120</b>
<b>STARPU_MULTIFORMAT_GET_NX . . . . .</b>	<b>120</b>
<b>17.10.4 Function Documentation . . . . .</b>	<b>120</b>
<b>starpu_multiformat_data_register . . . . .</b>	<b>120</b>
<b>17.11 Codelet And Tasks . . . . .</b>	<b>121</b>
<b>17.11.1 Detailed Description . . . . .</b>	<b>122</b>
<b>17.11.2 Data Structure Documentation . . . . .</b>	<b>122</b>
<b>struct starpu_codelet . . . . .</b>	<b>122</b>
<b>struct starpu_data_descr . . . . .</b>	<b>124</b>
<b>struct starpu_task . . . . .</b>	<b>124</b>
<b>17.11.3 Macro Definition Documentation . . . . .</b>	<b>128</b>
<b>STARPU_CPU . . . . .</b>	<b>128</b>
<b>STARPU_CUDA . . . . .</b>	<b>128</b>
<b>STARPU_OPENCL . . . . .</b>	<b>128</b>
<b>STARPU_MULTIPLE_CPU_IMPLEMENTATIONS . . . . .</b>	<b>128</b>
<b>STARPU_MULTIPLE_CUDA_IMPLEMENTATIONS . . . . .</b>	<b>128</b>
<b>STARPU_MULTIPLE_OPENCL_IMPLEMENTATIONS . . . . .</b>	<b>128</b>
<b>STARPU_NMAXBUFS . . . . .</b>	<b>128</b>
<b>STARPU_TASK_INITIALIZER . . . . .</b>	<b>128</b>
<b>STARPU_TASK_GET_HANDLE . . . . .</b>	<b>128</b>
<b>STARPU_TASK_SET_HANDLE . . . . .</b>	<b>128</b>
<b>STARPU_CODELET_GET_MODE . . . . .</b>	<b>129</b>
<b>STARPU_CODELET_SET_MODE . . . . .</b>	<b>129</b>
<b>STARPU_TASK_INVALID . . . . .</b>	<b>129</b>
<b>17.11.4 Typedef Documentation . . . . .</b>	<b>129</b>
<b>starpu_cpu_func_t . . . . .</b>	<b>129</b>
<b>starpu_cuda_func_t . . . . .</b>	<b>129</b>
<b>starpu_opencl_func_t . . . . .</b>	<b>129</b>
<b>17.11.5 Enumeration Type Documentation . . . . .</b>	<b>129</b>
<b>starpu_codelet_type . . . . .</b>	<b>129</b>
<b>starpu_task_status . . . . .</b>	<b>129</b>

17.11.6 Function Documentation . . . . .	130
starpu_codelet_init . . . . .	130
starpu_task_init . . . . .	130
starpu_task_create . . . . .	130
starpu_task_dup . . . . .	130
starpu_task_clean . . . . .	130
starpu_task_destroy . . . . .	130
starpu_task_wait . . . . .	130
starpu_task_submit . . . . .	130
starpu_task_submit_to_ctx . . . . .	131
starpu_task_wait_for_all . . . . .	131
starpu_task_wait_for_all_in_ctx . . . . .	131
starpu_task_nready . . . . .	131
starpu_task_nsubmitted . . . . .	131
starpu_task_get_current . . . . .	131
starpu_codelet_display_stats . . . . .	131
starpu_task_wait_for_no_ready . . . . .	131
starpu_task_set_implementation . . . . .	131
starpu_task_get_implementation . . . . .	131
starpu_create_sync_task . . . . .	131
17.12 Insert_Task . . . . .	131
17.12.1 Detailed Description . . . . .	132
17.12.2 Macro Definition Documentation . . . . .	132
STARPU_VALUE . . . . .	132
STARPU_CALLBACK . . . . .	132
STARPU_CALLBACK_WITH_ARG . . . . .	132
STARPU_CALLBACK_ARG . . . . .	132
STARPU_PRIORITY . . . . .	132
STARPU_DATA_ARRAY . . . . .	132
STARPU_TAG . . . . .	132
STARPU_FLOPS . . . . .	132
STARPU_SCHED_CTX . . . . .	132
17.12.3 Function Documentation . . . . .	133
starpu_insert_task . . . . .	133
starpu_codelet_pack_args . . . . .	133
starpu_codelet_unpack_args . . . . .	133
starpu_task_build . . . . .	133
17.13 Explicit Dependencies . . . . .	133
17.13.1 Detailed Description . . . . .	134
17.13.2 Typedef Documentation . . . . .	134
starpu_tag_t . . . . .	134
17.13.3 Function Documentation . . . . .	134
starpu_task_declare_deps_array . . . . .	134
starpu_tag_declare_deps . . . . .	134
starpu_tag_declare_deps_array . . . . .	134
starpu_tag_wait . . . . .	134
starpu_tag_wait_array . . . . .	134
starpu_tag_restart . . . . .	135
starpu_tag_remove . . . . .	135
starpu_tag_notify_from_apps . . . . .	135
17.14 Implicit Data Dependencies . . . . .	135
17.14.1 Detailed Description . . . . .	135
17.14.2 Function Documentation . . . . .	135
starpu_data_set_default_sequential_consistency_flag . . . . .	135
starpu_data_get_default_sequential_consistency_flag . . . . .	135
starpu_data_set_sequential_consistency_flag . . . . .	136
17.15 Performance Model . . . . .	136
17.15.1 Detailed Description . . . . .	136

17.15.2 Data Structure Documentation . . . . .	136
struct starpu_perfmodel . . . . .	136
struct starpu_perfmodel_regression_model . . . . .	137
struct starpu_perfmodel_per_arch . . . . .	138
struct starpu_perfmodel_history_list . . . . .	138
struct starpu_perfmodel_history_entry . . . . .	139
17.15.3 Enumeration Type Documentation . . . . .	139
starpu_perfmodel_archtype . . . . .	139
starpu_perfmodel_type . . . . .	139
17.15.4 Function Documentation . . . . .	140
starpu_perfmodel_load_symbol . . . . .	140
starpu_perfmodel_unload_model . . . . .	140
starpu_perfmodel_debugfilepath . . . . .	140
starpu_perfmodel_get_arch_name . . . . .	140
starpu_worker_get_perf_archtype . . . . .	140
starpu_perfmodel_list . . . . .	140
starpu_perfmodel_directory . . . . .	140
starpu_perfmodel_print . . . . .	140
starpu_perfmodel_print_all . . . . .	140
starpu_bus_print_bandwidth . . . . .	140
starpu_bus_print_affinity . . . . .	140
starpu_perfmodel_update_history . . . . .	141
starpu_transfer_bandwidth . . . . .	141
starpu_transfer_latency . . . . .	141
starpu_transfer_predict . . . . .	141
17.16 Profiling . . . . .	141
17.16.1 Detailed Description . . . . .	141
17.16.2 Data Structure Documentation . . . . .	141
struct starpu_profiling_task_info . . . . .	141
struct starpu_profiling_worker_info . . . . .	142
struct starpu_profiling_bus_info . . . . .	142
17.16.3 Macro Definition Documentation . . . . .	143
STARPU_PROFILING_DISABLE . . . . .	143
STARPU_PROFILING_ENABLE . . . . .	143
17.16.4 Function Documentation . . . . .	143
starpu_profiling_status_set . . . . .	143
starpu_profiling_status_get . . . . .	143
starpu_profiling_init . . . . .	143
starpu_profiling_set_id . . . . .	143
starpu_profiling_worker_get_info . . . . .	143
starpu_bus_get_profiling_info . . . . .	143
starpu_bus_get_count . . . . .	143
starpu_bus_get_id . . . . .	143
starpu_bus_get_src . . . . .	143
starpu_bus_get_dst . . . . .	144
starpu_timing_timespec_delay_us . . . . .	144
starpu_timing_timespec_to_us . . . . .	144
starpu_profiling_bus_helper_display_summary . . . . .	144
starpu_profiling_worker_helper_display_summary . . . . .	144
starpu_data_display_memory_stats . . . . .	144
17.17 Theoretical Lower Bound on Execution Time . . . . .	144
17.17.1 Detailed Description . . . . .	144
17.17.2 Function Documentation . . . . .	144
starpu_bound_start . . . . .	144
starpu_bound_stop . . . . .	144
starpu_bound_print_dot . . . . .	144
starpu_bound_compute . . . . .	145
starpu_bound_print_lp . . . . .	145

starpu_bound_print_mps . . . . .	145
starpu_bound_print . . . . .	145
<b>17.18CUDA Extensions . . . . .</b>	<b>145</b>
<b>17.18.1 Detailed Description . . . . .</b>	<b>145</b>
<b>17.18.2 Macro Definition Documentation . . . . .</b>	<b>145</b>
STARPU_USE_CUDA . . . . .	145
STARPU_MAXCUDADEVS . . . . .	145
STARPU_CUDA_REPORT_ERROR . . . . .	145
STARPU_CUBLAS_REPORT_ERROR . . . . .	145
<b>17.18.3 Function Documentation . . . . .</b>	<b>146</b>
starpu_cuda_get_local_stream . . . . .	146
starpu_cuda_get_device_properties . . . . .	146
starpu_cuda_report_error . . . . .	146
starpu_cuda_copy_async_sync . . . . .	146
starpu_cuda_set_device . . . . .	146
starpu_cublas_init . . . . .	146
starpu_cublas_shutdown . . . . .	146
starpu_cublas_report_error . . . . .	146
<b>17.19OpenCL Extensions . . . . .</b>	<b>146</b>
<b>17.19.1 Detailed Description . . . . .</b>	<b>148</b>
<b>17.19.2 Data Structure Documentation . . . . .</b>	<b>148</b>
struct starpu_opencl_program . . . . .	148
<b>17.19.3 Macro Definition Documentation . . . . .</b>	<b>149</b>
STARPU_USE_OPENCL . . . . .	149
STARPU_MAXOPENCLDEVS . . . . .	149
STARPU_OPENCL_DATADIR . . . . .	149
STARPU_OPENCL_DISPLAY_ERROR . . . . .	149
STARPU_OPENCL_REPORT_ERROR . . . . .	149
STARPU_OPENCL_REPORT_ERROR_WITH_MSG . . . . .	149
<b>17.19.4 Function Documentation . . . . .</b>	<b>149</b>
starpu_opencl_get_context . . . . .	149
starpu_opencl_get_device . . . . .	149
starpu_opencl_get_queue . . . . .	149
starpu_opencl_get_current_context . . . . .	149
starpu_opencl_get_current_queue . . . . .	149
starpu_opencl_set_kernel_args . . . . .	150
starpu_opencl_load_opencl_from_file . . . . .	150
starpu_opencl_load_opencl_from_string . . . . .	150
starpu_opencl_unload_opencl . . . . .	150
starpu_opencl_load_program_source . . . . .	150
starpu_opencl_compile_opencl_from_file . . . . .	150
starpu_opencl_compile_opencl_from_string . . . . .	150
starpu_opencl_load_binary_opencl . . . . .	150
starpu_opencl_load_kernel . . . . .	151
starpu_opencl_release_kernel . . . . .	151
starpu_opencl_collect_stats . . . . .	151
starpu_opencl_error_string . . . . .	151
starpu_opencl_display_error . . . . .	151
starpu_opencl_report_error . . . . .	151
starpu_opencl_allocate_memory . . . . .	151
starpu_opencl_copy_ram_to_opencl . . . . .	151
starpu_opencl_copy_opencl_to_ram . . . . .	151
starpu_opencl_copy_opencl_to_opencl . . . . .	152
starpu_opencl_copy_async_sync . . . . .	152
<b>17.20Miscellaneous Helpers . . . . .</b>	<b>152</b>
<b>17.20.1 Detailed Description . . . . .</b>	<b>152</b>
<b>17.20.2 Function Documentation . . . . .</b>	<b>152</b>
starpu_data_cpy . . . . .	152

starpu_execute_on_each_worker . . . . .	152
17.21 FxT Support . . . . .	152
17.21.1 Detailed Description . . . . .	153
17.21.2 Data Structure Documentation . . . . .	153
struct starpu_fxt_codelet_event . . . . .	153
struct starpu_fxt_options . . . . .	153
17.21.3 Function Documentation . . . . .	154
starpu_fxt_options_init . . . . .	154
starpu_fxt_generate_trace . . . . .	154
starpu_fxt_start_profiling . . . . .	154
starpu_fxt_stop_profiling . . . . .	154
17.22 FFT Support . . . . .	154
17.22.1 Detailed Description . . . . .	154
17.22.2 Function Documentation . . . . .	154
starpufft_malloc . . . . .	154
starpufft_free . . . . .	154
starpufft_plan_dft_1d . . . . .	155
starpufft_plan_dft_2d . . . . .	155
starpufft_start . . . . .	155
starpufft_start_handle . . . . .	155
starpufft_execute . . . . .	155
starpufft_execute_handle . . . . .	155
starpufft_cleanup . . . . .	155
starpufft_destroy_plan . . . . .	155
17.23 MPI Support . . . . .	155
17.23.1 Detailed Description . . . . .	156
17.23.2 Macro Definition Documentation . . . . .	156
STARPU_USE_MPI . . . . .	156
STARPU_EXECUTE_ON_NODE . . . . .	156
STARPU_EXECUTE_ON_DATA . . . . .	157
17.23.3 Function Documentation . . . . .	157
starpu_mpi_init . . . . .	157
starpu_mpi_initialize . . . . .	157
starpu_mpi_initialize_extended . . . . .	157
starpu_mpi_shutdown . . . . .	157
starpu_mpi_comm_amounts_retrieve . . . . .	157
starpu_mpi_send . . . . .	157
starpu_mpi_recv . . . . .	157
starpu_mpi_isend . . . . .	157
starpu_mpi_irecv . . . . .	157
starpu_mpi_isend_detached . . . . .	158
starpu_mpi_irecv_detached . . . . .	158
starpu_mpi_wait . . . . .	158
starpu_mpi_test . . . . .	158
starpu_mpi_barrier . . . . .	158
starpu_mpi_isend_detached_unlock_tag . . . . .	158
starpu_mpi_irecv_detached_unlock_tag . . . . .	158
starpu_mpi_isend_array_detached_unlock_tag . . . . .	158
starpu_mpi_irecv_array_detached_unlock_tag . . . . .	158
starpu_mpi_cache_flush . . . . .	158
starpu_mpi_cache_flush_all_data . . . . .	159
starpu_data_set_tag . . . . .	159
starpu_data_get_tag . . . . .	159
starpu_data_set_rank . . . . .	159
starpu_data_get_rank . . . . .	159
starpu_mpi_insert_task . . . . .	159
starpu_mpi_get_data_on_node . . . . .	159
starpu_mpi_get_data_on_node_detached . . . . .	159

starpu_mpi_redux_data . . . . .	160
starpu_mpi_scatter_detached . . . . .	160
starpu_mpi_gather_detached . . . . .	160
<b>17.24 Task Bundles . . . . .</b>	<b>160</b>
<b>17.24.1 Detailed Description . . . . .</b>	<b>160</b>
<b>17.24.2 Typedef Documentation . . . . .</b>	<b>160</b>
starpu_task_bundle_t . . . . .	160
<b>17.24.3 Function Documentation . . . . .</b>	<b>160</b>
starpu_task_bundle_create . . . . .	160
starpu_task_bundle_insert . . . . .	161
starpu_task_bundle_remove . . . . .	161
starpu_task_bundle_close . . . . .	161
starpu_task_bundle_expected_length . . . . .	161
starpu_task_bundle_expected_power . . . . .	161
starpu_task_bundle_expected_data_transfer_time . . . . .	161
<b>17.25 Task Lists . . . . .</b>	<b>161</b>
<b>17.25.1 Detailed Description . . . . .</b>	<b>162</b>
<b>17.25.2 Data Structure Documentation . . . . .</b>	<b>162</b>
struct starpu_task_list . . . . .	162
<b>17.25.3 Function Documentation . . . . .</b>	<b>163</b>
starpu_task_list_init . . . . .	163
starpu_task_list_push_front . . . . .	163
starpu_task_list_push_back . . . . .	163
starpu_task_list_front . . . . .	163
starpu_task_list_back . . . . .	163
starpu_task_list_empty . . . . .	163
starpu_task_list_erase . . . . .	163
starpu_task_list_pop_front . . . . .	163
starpu_task_list_pop_back . . . . .	163
starpu_task_list_begin . . . . .	163
starpu_task_list_next . . . . .	163
<b>17.26 Parallel Tasks . . . . .</b>	<b>164</b>
<b>17.26.1 Detailed Description . . . . .</b>	<b>164</b>
<b>17.26.2 Function Documentation . . . . .</b>	<b>164</b>
starpu_combined_worker_get_size . . . . .	164
starpu_combined_worker_get_rank . . . . .	164
starpu_combined_worker_get_count . . . . .	164
starpu_combined_worker_get_id . . . . .	164
starpu_combined_worker_assign_workerid . . . . .	164
starpu_combined_worker_get_description . . . . .	164
starpu_combined_worker_can_execute_task . . . . .	164
starpu_parallel_task_barrier_init . . . . .	164
<b>17.27 Running Drivers . . . . .</b>	<b>165</b>
<b>17.27.1 Detailed Description . . . . .</b>	<b>165</b>
<b>17.27.2 Function Documentation . . . . .</b>	<b>165</b>
starpu_driver_run . . . . .	165
starpu_driver_init . . . . .	165
starpu_driver_run_once . . . . .	165
starpu_driver_deinit . . . . .	165
starpu_drivers_request_termination . . . . .	165
<b>17.28 Expert Mode . . . . .</b>	<b>165</b>
<b>17.28.1 Detailed Description . . . . .</b>	<b>165</b>
<b>17.28.2 Function Documentation . . . . .</b>	<b>165</b>
starpu_wake_all_blocked_workers . . . . .	165
starpu_progression_hook_register . . . . .	165
starpu_progression_hook_deregister . . . . .	165
<b>17.29 StarPU-Top Interface . . . . .</b>	<b>166</b>
<b>17.29.1 Detailed Description . . . . .</b>	<b>166</b>

17.29.2 Data Structure Documentation . . . . .	166
struct starpu_top_data . . . . .	166
struct starpu_top_param . . . . .	167
17.29.3 Enumeration Type Documentation . . . . .	168
starpu_top_data_type . . . . .	168
starpu_top_param_type . . . . .	168
starpu_top_message_type . . . . .	168
17.29.4 Function Documentation . . . . .	168
starpu_top_add_data_boolean . . . . .	168
starpu_top_add_data_integer . . . . .	168
starpu_top_add_data_float . . . . .	169
starpu_top_register_parameter_boolean . . . . .	169
starpu_top_register_parameter_float . . . . .	169
starpu_top_register_parameter_integer . . . . .	169
starpu_top_register_parameter_enum . . . . .	169
starpu_top_init_and_wait . . . . .	169
starpu_top_update_parameter . . . . .	169
starpu_top_update_data_boolean . . . . .	169
starpu_top_update_data_integer . . . . .	169
starpu_top_update_data_float . . . . .	169
starpu_top_task_prevision . . . . .	169
starpu_top_debug_log . . . . .	170
starpu_top_debug_lock . . . . .	170
17.30 Scheduling Contexts . . . . .	170
17.30.1 Detailed Description . . . . .	171
17.30.2 Data Structure Documentation . . . . .	171
struct starpu_sched_ctx_performance_counters . . . . .	171
17.30.3 Macro Definition Documentation . . . . .	172
STARPU_SCHED_CTX_POLICY_NAME . . . . .	172
STARPU_SCHED_CTX_POLICY_STRUCT . . . . .	172
STARPU_SCHED_CTX_POLICY_MIN_PRIO . . . . .	172
STARPU_SCHED_CTX_POLICY_MAX_PRIO . . . . .	172
STARPU_MIN_PRIO . . . . .	172
STARPU_MAX_PRIO . . . . .	172
STARPU_DEFAULT_PRIO . . . . .	172
17.30.4 Function Documentation . . . . .	172
starpu_sched_ctx_create . . . . .	172
starpu_sched_ctx_create_inside_interval . . . . .	172
starpu_sched_ctx_register_close_callback . . . . .	172
starpu_sched_ctx_add_workers . . . . .	173
starpu_sched_ctx_remove_workers . . . . .	173
starpu_sched_ctx_delete . . . . .	173
starpu_sched_ctx_set_inheritor . . . . .	173
starpu_sched_ctx_set_context . . . . .	173
starpu_sched_ctx_get_context . . . . .	173
starpu_sched_ctx_stop_task_submission . . . . .	173
starpu_sched_ctx_finished_submit . . . . .	173
starpu_sched_ctx_get_workers_list . . . . .	173
starpu_sched_ctx_get_nworkers . . . . .	173
starpu_sched_ctx_get_nshared_workers . . . . .	173
starpu_sched_ctx_contains_worker . . . . .	173
starpu_sched_ctx_worker_get_id . . . . .	173
starpu_sched_ctx_overlapping_ctxs_on_worker . . . . .	174
starpu_sched_ctx_set_min_priority . . . . .	174
starpu_sched_ctx_set_max_priority . . . . .	174
starpu_sched_ctx_get_min_priority . . . . .	174
starpu_sched_ctx_get_max_priority . . . . .	174
starpu_sched_ctx_create_worker_collection . . . . .	174

starpu_sched_ctx_delete_worker_collection . . . . .	174
starpu_sched_ctx_get_worker_collection . . . . .	174
starpu_sched_ctx_set_perf_counters . . . . .	174
starpu_sched_ctx_call_pushed_task_cb . . . . .	174
starpu_sched_ctx_notify_hypervisor_exists . . . . .	174
starpu_sched_ctx_check_if_hypervisor_exists . . . . .	174
starpu_sched_ctx_set_policy_data . . . . .	174
starpu_sched_ctx_get_policy_data . . . . .	175
<b>17.31 Scheduling Policy . . . . .</b>	<b>175</b>
<b>17.31.1 Detailed Description . . . . .</b>	<b>175</b>
<b>17.31.2 Data Structure Documentation . . . . .</b>	<b>175</b>
<b>struct starpu_sched_policy . . . . .</b>	<b>175</b>
<b>17.31.3 Function Documentation . . . . .</b>	<b>177</b>
<b>starpu_sched_get_predefined_policies . . . . .</b>	<b>177</b>
<b>starpu_worker_get_sched_condition . . . . .</b>	<b>177</b>
<b>starpu_sched_set_min_priority . . . . .</b>	<b>177</b>
<b>starpu_sched_set_max_priority . . . . .</b>	<b>177</b>
<b>starpu_sched_get_min_priority . . . . .</b>	<b>177</b>
<b>starpu_sched_get_max_priority . . . . .</b>	<b>177</b>
<b>starpu_push_local_task . . . . .</b>	<b>177</b>
<b>starpu_push_task_end . . . . .</b>	<b>177</b>
<b>starpu_worker_can_execute_task . . . . .</b>	<b>177</b>
<b>starpu_timing_now . . . . .</b>	<b>177</b>
<b>starpu_task_footprint . . . . .</b>	<b>177</b>
<b>starpu_task_expected_length . . . . .</b>	<b>178</b>
<b>starpu_worker_get_relative_speedup . . . . .</b>	<b>178</b>
<b>starpu_task_expected_data_transfer_time . . . . .</b>	<b>178</b>
<b>starpu_data_expected_transfer_time . . . . .</b>	<b>178</b>
<b>starpu_task_expected_power . . . . .</b>	<b>178</b>
<b>starpu_task_expected_conversion_time . . . . .</b>	<b>178</b>
<b>starpu_get_prefetch_flag . . . . .</b>	<b>178</b>
<b>starpu_prefetch_task_input_on_node . . . . .</b>	<b>178</b>
<b>starpu_sched_ctx_worker_shares_tasks_lists . . . . .</b>	<b>178</b>
<b>17.32 Scheduling Context Hypervisor - Regular usage . . . . .</b>	<b>178</b>
<b>17.32.1 Detailed Description . . . . .</b>	<b>179</b>
<b>17.32.2 Macro Definition Documentation . . . . .</b>	<b>179</b>
<b>SC_HYPERVERSOR_MAX_IDLE . . . . .</b>	<b>179</b>
<b>SC_HYPERVERSOR_PRIORITY . . . . .</b>	<b>179</b>
<b>SC_HYPERVERSOR_MIN_WORKERS . . . . .</b>	<b>179</b>
<b>SC_HYPERVERSOR_MAX_WORKERS . . . . .</b>	<b>179</b>
<b>SC_HYPERVERSOR_GRANULARITY . . . . .</b>	<b>179</b>
<b>SC_HYPERVERSOR_FIXED_WORKERS . . . . .</b>	<b>179</b>
<b>SC_HYPERVERSOR_MIN_TASKS . . . . .</b>	<b>180</b>
<b>SC_HYPERVERSOR_NEW_WORKERS_MAX_IDLE . . . . .</b>	<b>180</b>
<b>SC_HYPERVERSOR_TIME_TO_APPLY . . . . .</b>	<b>180</b>
<b>SC_HYPERVERSOR_ISPEED_W_SAMPLE . . . . .</b>	<b>180</b>
<b>SC_HYPERVERSOR_ISPEED_CTX_SAMPLE . . . . .</b>	<b>180</b>
<b>SC_HYPERVERSOR_NULL . . . . .</b>	<b>180</b>
<b>17.32.3 Function Documentation . . . . .</b>	<b>180</b>
<b>sc.hypervisor_init . . . . .</b>	<b>180</b>
<b>sc.hypervisor_shutdown . . . . .</b>	<b>180</b>
<b>sc.hypervisor_register_ctx . . . . .</b>	<b>180</b>
<b>sc.hypervisor_unregister_ctx . . . . .</b>	<b>180</b>
<b>sc.hypervisor_stop_resize . . . . .</b>	<b>181</b>
<b>sc.hypervisor_start_resize . . . . .</b>	<b>181</b>
<b>sc.hypervisor_get_policy . . . . .</b>	<b>181</b>
<b>sc.hypervisor_add_workers_to_sched_ctx . . . . .</b>	<b>181</b>
<b>sc.hypervisor_remove_workers_from_sched_ctx . . . . .</b>	<b>181</b>

sc.hypervisor_move_workers . . . . .	181
sc.hypervisor_size_ctxs . . . . .	181
sc.hypervisor_set_type_of_task . . . . .	181
sc.hypervisor_update_diff_total_flops . . . . .	181
sc.hypervisor_update_diff_elapsed_flops . . . . .	181
sc.hypervisor_ctl . . . . .	181
<b>17.33 Scheduling Context Hypervisor - Building a new resizing policy</b> . . . . .	181
<b>17.33.1 Detailed Description</b> . . . . .	182
<b>17.33.2 Data Structure Documentation</b> . . . . .	182
struct sc_hypervisor_policy . . . . .	182
struct sc_hypervisor_policy_config . . . . .	183
struct sc_hypervisor_wrapper . . . . .	183
struct sc_hypervisor_resize_ack . . . . .	185
struct sc_hypervisor_policy_task_pool . . . . .	185
<b>17.33.3 Function Documentation</b> . . . . .	186
sc_hypervisor_post_resize_request . . . . .	186
sc_hypervisor_get_size_req . . . . .	186
sc_hypervisor_save_size_req . . . . .	186
sc_hypervisor_free_size_req . . . . .	186
sc_hypervisor_can_resize . . . . .	186
sc_hypervisor_get_config . . . . .	186
sc_hypervisor_set_config . . . . .	186
sc_hypervisor_get_sched_ctxs . . . . .	186
sc_hypervisor_get_nsched_ctxs . . . . .	186
sc_hypervisor_get_wrapper . . . . .	186
sc_hypervisor_get_elapsed_flops_per_sched_ctx . . . . .	186
<b>18 File Index</b> . . . . .	187
<b>18.1 File List</b> . . . . .	187
<b>19 File Documentation</b> . . . . .	189
<b>19.1 starpu.h File Reference</b> . . . . .	189
<b>19.2 starpu_bound.h File Reference</b> . . . . .	190
<b>19.3 starpu_config.h File Reference</b> . . . . .	190
<b>19.4 starpu_cublas.h File Reference</b> . . . . .	191
<b>19.5 starpu_cuda.h File Reference</b> . . . . .	191
<b>19.6 starpu_data.h File Reference</b> . . . . .	192
<b>19.7 starpu_data_filters.h File Reference</b> . . . . .	193
<b>19.8 starpu_data_interfaces.h File Reference</b> . . . . .	195
<b>19.9 starpu_DEPRECATED_api.h File Reference</b> . . . . .	199
<b>19.10 starpu_driver.h File Reference</b> . . . . .	199
<b>19.11 starpu_expert.h File Reference</b> . . . . .	199
<b>19.12 starpu_fxt.h File Reference</b> . . . . .	199
<b>19.13 starpu_hash.h File Reference</b> . . . . .	200
<b>19.14 starpu_opencl.h File Reference</b> . . . . .	200
<b>19.15 starpu_perfmodel.h File Reference</b> . . . . .	201
<b>19.16 starpu_profiling.h File Reference</b> . . . . .	202
<b>19.17 starpu_rand.h File Reference</b> . . . . .	203
<b>19.18 starpu_sched_ctx.h File Reference</b> . . . . .	203
<b>19.19 starpu_sched_ctx_hypervisor.h File Reference</b> . . . . .	204
<b>19.20 starpu_scheduler.h File Reference</b> . . . . .	204
<b>19.21 starpu_stdlib.h File Reference</b> . . . . .	205
<b>19.22 starpu_task.h File Reference</b> . . . . .	206
<b>19.23 starpu_task_bundle.h File Reference</b> . . . . .	207
<b>19.24 starpu_task_list.h File Reference</b> . . . . .	207
<b>19.25 starpu_task_util.h File Reference</b> . . . . .	208
<b>19.26 starpu_thread.h File Reference</b> . . . . .	208
<b>19.27 starpu_thread_util.h File Reference</b> . . . . .	209
<b>19.28 starpu_top.h File Reference</b> . . . . .	210

19.29starpu_util.h File Reference . . . . .	211
19.30starpu_worker.h File Reference . . . . .	212
19.31starpu_mpi.h File Reference . . . . .	213
19.32sc_hypervisor.h File Reference . . . . .	214
19.33sc_hypervisor_config.h File Reference . . . . .	215
19.34sc_hypervisor_lp.h File Reference . . . . .	215
19.35sc_hypervisor_monitoring.h File Reference . . . . .	216
19.36sc_hypervisor_policy.h File Reference . . . . .	216
19.36.1 Data Structure Documentation . . . . .	217
struct types_of_workers . . . . .	217
<b>20 Deprecated List</b>	<b>219</b>
<b>Index</b>	<b>219</b>
<b>III Appendix</b>	<b>221</b>
<b>21 Full Source Code for the 'Scaling a Vector' Example</b>	<b>223</b>
21.1 Main Application . . . . .	223
21.2 CPU Kernel . . . . .	224
21.3 CUDA Kernel . . . . .	225
21.4 OpenCL Kernel . . . . .	225
21.4.1 Invoking the Kernel . . . . .	225
21.4.2 Source of the Kernel . . . . .	226
<b>22 GNU Free Documentation License</b>	<b>227</b>
22.1 ADDENDUM: How to use this License for your documents . . . . .	231



# **Chapter 1**

## **Introduction**

### **1.1 Motivation**



# **Part I**

## **Using StarPU**



# Chapter 2

## Building and Installing StarPU

### 2.1 Installing a Binary Package

One of the StarPU developers being a Debian Developer, the packages are well integrated and very up-to-date. To see which packages are available, simply type:

```
$ apt-cache search starpu
```

To install what you need, type:

```
$ sudo apt-get install libstarpu-1.1 libstarpu-dev
```

### 2.2 Installing from Source

StarPU can be built and installed by the standard means of the GNU autotools. The following chapter is intended to briefly remind how these tools can be used to install StarPU.

#### 2.2.1 Optional Dependencies

The `hwloc topology discovery library` is not mandatory to use StarPU but strongly recommended. It allows for topology aware scheduling, which improves performance. `hwloc` is available in major free operating system distributions, and for most operating systems.

If `hwloc` is not available on your system, the option `--without-hwloc` should be explicitly given when calling the `configure` script. If `hwloc` is installed with a `pkg-config` file, no option is required, it will be detected automatically, otherwise `--with-hwloc` should be used to specify the location of `hwloc`.

#### 2.2.2 Getting Sources

StarPU's sources can be obtained from the [download page of the StarPU website](#).

All releases and the development tree of StarPU are freely available on INRIA's gforge under the LGPL license. Some releases are available under the BSD license.

The latest release can be downloaded from the [INRIA's gforge](#) or directly from the [StarPU download page](#).

The latest nightly snapshot can be downloaded from the [StarPU gforge website](#).

```
$ wget http://starpu.gforge.inria.fr/testing/starpu-nightly-latest.tar.gz
```

And finally, current development version is also accessible via svn. It should be used only if you need the very latest changes (i.e. less than a day!). Note that the client side of the software Subversion can be obtained from <http://subversion.tigris.org>. If you are running on Windows, you will probably prefer to use [TortoiseSVN](#).

```
$ svn checkout svn://scm.gforge.inria.fr/svn/starpu/trunk StarPU
```

### 2.2.3 Configuring StarPU

Running `autogen.sh` is not necessary when using the tarball releases of StarPU. If you are using the source code from the `svn` repository, you first need to generate the configure scripts and the Makefiles. This requires the availability of `autoconf`, `automake >= 2.60`.

```
$ ./autogen.sh
```

You then need to configure StarPU. Details about options that are useful to give to `./configure` are given in [Compilation Configuration](#).

```
$ ./configure
```

If `configure` does not detect some software or produces errors, please make sure to post the content of `config.log` when reporting the issue.

By default, the files produced during the compilation are placed in the source directory. As the compilation generates a lot of files, it is advised to put them all in a separate directory. It is then easier to cleanup, and this allows to compile several configurations out of the same source tree. For that, simply enter the directory where you want the compilation to produce its files, and invoke the `configure` script located in the StarPU source directory.

```
$ mkdir build
$ cd build
$ ../configure
```

### 2.2.4 Building StarPU

```
$ make
```

Once everything is built, you may want to test the result. An extensive set of regression tests is provided with StarPU. Running the tests is done by calling `make check`. These tests are run every night and the result from the main profile is publicly [available](#).

```
$ make check
```

### 2.2.5 Installing StarPU

In order to install StarPU at the location that was specified during configuration:

```
$ make install
```

Libtool interface versioning information are included in libraries names (`libstarpumpi-1.1.so` and `libstarpufft-1.1.so`).

## 2.3 Setting up Your Own Code

### 2.3.1 Setting Flags for Compiling, Linking and Running Applications

StarPU provides a `pkg-config` executable to obtain relevant compiler and linker flags. Compiling and linking an application against StarPU may require to use specific flags or libraries (for instance `CUDA` or `libspe2`). To this end, it is possible to use the tool `pkg-config`.

If StarPU was not installed at some standard location, the path of StarPU's library must be specified in the environment variable `PKG_CONFIG_PATH` so that `pkg-config` can find it. For example if StarPU was installed in `$prefix_dir`:

```
$ PKG_CONFIG_PATH=$PKG_CONFIG_PATH:$prefix_dir/lib/pkgconfig
```

The flags required to compile or link against StarPU are then accessible with the following commands:

```
$ pkg-config --cflags starpu-1.1    # options for the compiler
$ pkg-config --libs starpu-1.1      # options for the linker
```

Note that it is still possible to use the API provided in the version 0.9 of StarPU by calling `pkg-config` with the `libstarpu` package. Similar packages are provided for `libstarpumpi` and `libstarpufft`.

Make sure that `pkg-config --libs starpu-1.1` actually produces some output before going further: `PKG_CONFIG_PATH` has to point to the place where `starpu-1.1.pc` was installed during `make install`.

Also pass the option `--static` if the application is to be linked statically.

It is also necessary to set the environment variable `LD_LIBRARY_PATH` to locate dynamic libraries at runtime.

```
$ LD_LIBRARY_PATH=$prefix_dir/lib:$LD_LIBRARY_PATH
```

When using a Makefile, the following lines can be added to set the options for the compiler and the linker:

```
CFLAGS      +=      $(shell pkg-config --cflags starpu-1.1)
LDFLAGS     +=      $(shell pkg-config --libs starpu-1.1)
```

### 2.3.2 Running a Basic StarPU Application

Basic examples using StarPU are built in the directory `examples/basic_examples/` (and installed in `$prefix_dir/lib/starpu/examples/`). You can for example run the example `vector_scal`.

```
$ ./examples/basic_examples/vector_scal
BEFORE: First element was 1.000000
AFTER: First element is 3.140000
```

When StarPU is used for the first time, the directory `$STARPU_HOME/.starpu/` is created, performance models will be stored in that directory (`STARPU_HOME`).

Please note that buses are benchmarked when StarPU is launched for the first time. This may take a few minutes, or less if `hwloc` is installed. This step is done only once per user and per machine.

### 2.3.3 Kernel Threads Started by StarPU

StarPU automatically binds one thread per CPU core. It does not use SMT/hyperthreading because kernels are usually already optimized for using a full core, and using hyperthreading would make kernel calibration rather random.

Since driving GPUs is a CPU-consuming task, StarPU dedicates one core per GPU.

While StarPU tasks are executing, the application is not supposed to do computations in the threads it starts itself, tasks should be used instead.

TODO: add a StarPU function to bind an application thread (e.g. the main thread) to a dedicated core (and thus disable the corresponding StarPU CPU worker).

### 2.3.4 Enabling OpenCL

When both CUDA and OpenCL drivers are enabled, StarPU will launch an OpenCL worker for NVIDIA GPUs only if CUDA is not already running on them. This design choice was necessary as OpenCL and CUDA can not run at the same time on the same NVIDIA GPU, as there is currently no interoperability between them.

To enable OpenCL, you need either to disable CUDA when configuring StarPU:

```
$ ./configure --disable-cuda
```

or when running applications:

```
$ STARPU_NCUDA=0 ./application
```

OpenCL will automatically be started on any device not yet used by CUDA. So on a machine running 4 GPUS, it is therefore possible to enable CUDA on 2 devices, and OpenCL on the 2 other devices by doing so:

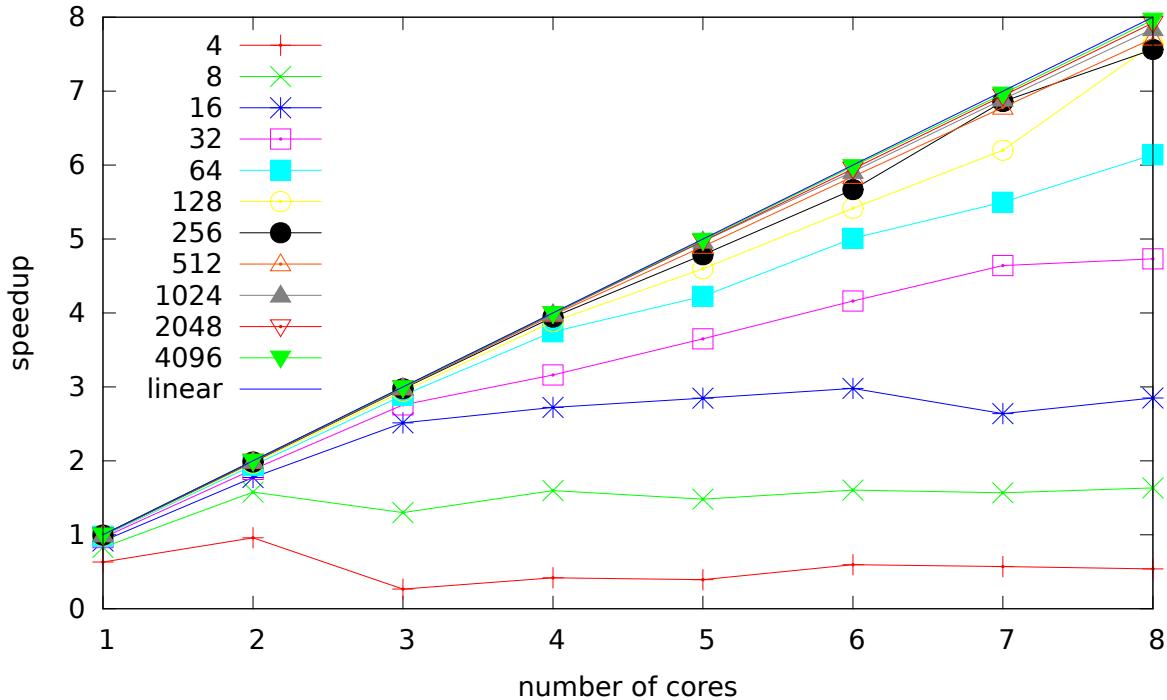
```
$ STARPU_NCUDA=2 ./application
```

## 2.4 Benchmarking StarPU

Some interesting benchmarks are installed among examples in `$prefix_dir/lib/starpu/examples/`. Make sure to try various schedulers, for instance `STARPU_SCHED=dmda`.

### 2.4.1 Task Size Overhead

This benchmark gives a glimpse into how long a task should be (in  $\mu\text{s}$ ) for StarPU overhead to be low enough to keep efficiency. Run `tasks_size_overhead.sh`, it will generate a plot of the speedup of tasks of various sizes, depending on the number of CPUs being used.



### 2.4.2 Data Transfer Latency

`local_pingpong` performs a ping-pong between the first two CUDA nodes, and prints the measured latency.

### 2.4.3 Matrix-Matrix Multiplication

`sgemm` and `dgemm` perform a blocked matrix-matrix multiplication using BLAS and cuBLAS. They output the obtained GFlops.

### 2.4.4 Cholesky Factorization

`cholesky/*` perform a Cholesky factorization (single precision). They use different dependency primitives.

### 2.4.5 LU Factorization

`lu/*` perform an LU factorization. They use different dependency primitives.

# Chapter 3

## Basic Examples

### 3.1 Hello World Using The C Extension

This section shows how to implement a simple program that submits a task to StarPU using the StarPU C extension ([C Extensions](#)). The complete example, and additional examples, is available in the directory `gcc-plugin/examples` of the StarPU distribution. A similar example showing how to directly use the StarPU's API is shown in [Hello World Using StarPU's API](#).

GCC from version 4.5 permit to use the StarPU GCC plug-in ([C Extensions](#)). This makes writing a task both simpler and less error-prone. In a nutshell, all it takes is to declare a task, declare and define its implementations (for CPU, OpenCL, and/or CUDA), and invoke the task like a regular C function. The example below defines `my_task` which has a single implementation for CPU:

```
#include <stdio.h>

/* Task declaration. */
static void my_task (int x) __attribute__ ((task));

/* Definition of the CPU implementation of 'my_task'. */
static void my_task (int x)
{
    printf ("Hello, world! With x = %d\n", x);
}

int main ()
{
    /* Initialize StarPU. */
#pragma starpu initialize

    /* Do an asynchronous call to 'my_task'. */
    my_task (42);

    /* Wait for the call to complete. */
#pragma starpu wait

    /* Terminate. */
#pragma starpu shutdown

    return 0;
}
```

The code can then be compiled and linked with GCC and the flag `-fplugin`:

```
$ gcc `pkg-config starpu-1.1 --cflags` hello-starpu.c \
-fplugin='pkg-config starpu-1.1 --variable=gccplugin' \
`pkg-config starpu-1.1 --libs`
```

The code can also be compiled without the StarPU C extension and will behave as a normal sequential code.

```
$ gcc hello-starpu.c
hello-starpu.c:33:1: warning: 'task' attribute directive ignored [-Wattributes]
$ ./a.out
Hello, world! With x = 42
```

As can be seen above, the C extensions allows programmers to use StarPU tasks by essentially annotating "regular" C code.

## 3.2 Hello World Using StarPU's API

This section shows how to achieve the same result as in the previous section using StarPU's standard C API.

### 3.2.1 Required Headers

The header `starpu.h` should be included in any code using StarPU.

```
#include <starpu.h>
```

### 3.2.2 Defining A Codelet

A codelet is a structure that represents a computational kernel. Such a codelet may contain an implementation of the same kernel on different architectures (e.g. CUDA, x86, ...). For compatibility, make sure that the whole structure is properly initialized to zero, either by using the function `starpu_codelet_init()`, or by letting the compiler implicitly do it as exemplified above.

The field `starpu_codelet::nbuffers` specifies the number of data buffers that are manipulated by the codelet: here the codelet does not access or modify any data that is controlled by our data management library.

We create a codelet which may only be executed on the CPUs. When a CPU core will execute a codelet, it will call the function `cpu_func`, which *must* have the following prototype:

```
void (*cpu_func)(void *buffers[], void *cl_arg);
```

In this example, we can ignore the first argument of this function which gives a description of the input and output buffers (e.g. the size and the location of the matrices) since there is none. We also ignore the second argument which is a pointer to optional arguments for the codelet.

```
void cpu_func(void *buffers[], void *cl_arg)
{
    printf("Hello world\n");
}

struct starpu_codelet cl =
{
    .cpu_funcs = { cpu_func, NULL },
    .nbuffers = 0
};
```

### 3.2.3 Submitting A Task

Before submitting any tasks to StarPU, `starpu_init()` must be called. The `NULL` argument specifies that we use the default configuration. Tasks cannot be submitted after the termination of StarPU by a call to `starpu_shutdown()`.

In the example above, a task structure is allocated by a call to `starpu_task_create()`. This function only allocates and fills the corresponding structure with the default settings, but it does not submit the task to StarPU.

# Chapter 4

## Advanced Examples

### 4.1 Using Multiple Implementations Of A Codelet

One may want to write multiple implementations of a codelet for a single type of device and let StarPU choose which one to run. As an example, we will show how to use SSE to scale a vector. The codelet can be written as follows:

```
#include <xmmmintrin.h>

void scal_sse_func(void *buffers[], void *cl_arg)
{
    float *vector = (float *) STARPU_VECTOR_GET_PTR(buffers[0]);
    unsigned int n = STARPU_VECTOR_GET_NX(buffers[0]);
    unsigned int n_iterations = n/4;
    if (n % 4 != 0)
        n_iterations++;

    __m128 *VECTOR = (__m128*) vector;
    __m128 factor __attribute__((aligned(16)));
    factor = _mm_set1_ps(*(float *) cl_arg);

    unsigned int i;
    for (i = 0; i < n_iterations; i++)
        VECTOR[i] = _mm_mul_ps(factor, VECTOR[i]);
}

struct starpu_codelet cl = {
    .where = STARPU_CPU,
    .cpu_funcs = { scal_cpu_func, scal_sse_func, NULL },
    .nbuffers = 1,
    .modes = { STARPU_RW }
};
```

Schedulers which are multi-implementation aware (only `dmda` and `pheft` for now) will use the performance models of all the implementations it was given, and pick the one that seems to be the fastest.

### 4.2 Enabling Implementation According To Capabilities

Some implementations may not run on some devices. For instance, some CUDA devices do not support double floating point precision, and thus the kernel execution would just fail; or the device may not have enough shared memory for the implementation being used. The field `starpu_codelet::can_execute` permits to express this. For instance:

```
static int can_execute(unsigned workerid, struct starpu_task *task, unsigned nimpl)
{
    const struct cudaDeviceProp *props;
    if (starpu_worker_get_type(workerid) ==
        STARPU_CPU_WORKER)
        return 1;
    /* Cuda device */
    props = starpu_cuda_get_device_properties(workerid);
    if ((props->major >= 2 || props->minor >= 3)
        /* At least compute capability 1.3, supports doubles */
        return 1;
    /* Old card, does not support doubles */
    return 0;
}

struct starpu_codelet cl = {
```

```

    .where = STARPU_CPU|STARPU_CUDA,
    .can_execute = can_execute,
    .cpu_funcs = { cpu_func, NULL },
    .cuda_funcs = { gpu_func, NULL }
    .nbuffers = 1,
    .modes = { STARPU_RW }
};

}

```

This can be essential e.g. when running on a machine which mixes various models of CUDA devices, to take benefit from the new models without crashing on old models.

Note: the function `starpu_codelet::can_execute` is called by the scheduler each time it tries to match a task with a worker, and should thus be very fast. The function `starpu_cuda_get_device_properties()` provides a quick access to CUDA properties of CUDA devices to achieve such efficiency.

Another example is to compile CUDA code for various compute capabilities, resulting with two CUDA functions, e.g. `scal_gpu_13` for compute capability 1.3, and `scal_gpu_20` for compute capability 2.0. Both functions can be provided to StarPU by using `starpu_codelet::cuda_funcs`, and `starpu_codelet::can_execute` can then be used to rule out the `scal_gpu_20` variant on a CUDA device which will not be able to execute it:

```

static int can_execute(unsigned workerid, struct starpu_task *task, unsigned nimpl)
{
    const struct cudaDeviceProp *props;
    if (starpu_worker_get_type(workerid) ==
        STARPU_CPU_WORKER)
        return 1;
    /* Cuda device */
    if (nimpl == 0)
        /* Trying to execute the 1.3 capability variant, we assume it is ok in all cases. */
        return 1;
    /* Trying to execute the 2.0 capability variant, check that the card can do it. */
    props = starpu_cuda_get_device_properties(workerid);
    if (props->major >= 2 || props->minor >= 0)
        /* At least compute capability 2.0, can run it */
        return 1;
    /* Old card, does not support 2.0, will not be able to execute the 2.0 variant. */
    return 0;
}

struct starpu_codelet cl = {
    .where = STARPU_CPU|STARPU_CUDA,
    .can_execute = can_execute,
    .cpu_funcs = { cpu_func, NULL },
    .cuda_funcs = { scal_gpu_13, scal_gpu_20, NULL },
    .nbuffers = 1,
    .modes = { STARPU_RW }
};

```

Note: the most generic variant should be provided first, as some schedulers are not able to try the different variants.

## 4.3 Task And Worker Profiling

A full example showing how to use the profiling API is available in the StarPU sources in the directory `examples/profiling/`.

```

struct starpu_task *task = starpu_task_create();
task->cl = &cl;
task->synchronous = 1;
/* We will destroy the task structure by hand so that we can
 * query the profiling info before the task is destroyed. */
task->destroy = 0;

/* Submit and wait for completion (since synchronous was set to 1) */
starpu_task_submit(task);

/* The task is finished, get profiling information */
struct starpu_profiling_task_info *info = task->
    profiling_info;

/* How much time did it take before the task started ? */
double delay += starpu_timing_timespec_delay_us(&info->
    submit_time, &info->start_time);

/* How long was the task execution ? */
double length += starpu_timing_timespec_delay_us(&info->
    start_time, &info->end_time);

/* We don't need the task structure anymore */
starpu_task_destroy(task);

```

```

/* Display the occupancy of all workers during the test */
int worker;
for (worker = 0; worker < starpu.worker.get_count(); worker++)
{
    struct starpu_profiling_worker_info worker_info;
    int ret = starpu.profiling_worker_get_info(worker, &worker_info);
    STARPU_ASSERT(!ret);

    double total_time = starpu_timing.timespec_to_us(&worker_info
        .total_time);
    double executing_time = starpu_timing.timespec_to_us(&
        worker_info.executing_time);
    double sleeping_time = starpu_timing.timespec_to_us(&
        worker_info.sleeping_time);
    double overhead_time = total_time - executing_time - sleeping_time;

    float executing_ratio = 100.0 * executing_time / total_time;
    float sleeping_ratio = 100.0 * sleeping_time / total_time;
    float overhead_ratio = 100.0 - executing_ratio - sleeping_ratio;

    char workername[128];
    starpu.worker.get_name(worker, workername, 128);
    fprintf(stderr, "Worker %s:\n", workername);
    fprintf(stderr, "\ttotal time: %.2lf ms\n", total_time * 1e-3);
    fprintf(stderr, "\texec time: %.2lf ms (%.2f %%)\n",
        executing_time * 1e-3, executing_ratio);
    fprintf(stderr, "\tblocked time: %.2lf ms (%.2f %%)\n",
        sleeping_time * 1e-3, sleeping_ratio);
    fprintf(stderr, "\toverhead time: %.2lf ms (%.2f %%)\n",
        overhead_time * 1e-3, overhead_ratio);
}

```

## 4.4 Partitioning Data

An existing piece of data can be partitioned in sub parts to be used by different tasks, for instance:

```

int vector[NX];
starpu_data.handle_t handle;

/* Declare data to StarPU */
starpu_vector_data_register(&handle, 0, (uintptr_t)vector,
                           NX, sizeof(vector[0]));

/* Partition the vector in PARTS sub-vectors */
starpu_data.filter f =
{
    .filter_func = starpu_vector_filter_block,
    .nchildren = PARTS
};
starpu_data.partition(handle, &f);

```

The task submission then uses the function `starpu_data_get_sub_data()` to retrieve the sub-handles to be passed as tasks parameters.

```

/* Submit a task on each sub-vector */
for (i=0; i<starpu_data.get_nb_children(handle); i++) {
    /* Get subdata number i (there is only 1 dimension) */
    starpu_data.handle_t sub_handle =
        starpu_data.get_sub_data(handle, 1, i);
    struct starputask *task = starputask_create();

    task->handles[0] = sub_handle;
    task->cl = &cl;
    task->synchronous = 1;
    task->cl_arg = &factor;
    task->cl_arg_size = sizeof(factor);

    starpu_task_submit(task);
}

```

Partitioning can be applied several times, see `examples/basic_examples/mult.c` and `examples/filters/`.

Wherever the whole piece of data is already available, the partitioning will be done in-place, i.e. without allocating new buffers but just using pointers inside the existing copy. This is particularly important to be aware of when using OpenCL, where the kernel parameters are not pointers, but handles. The kernel thus needs to be also passed the offset within the OpenCL buffer:

```

void opencl_func(void *buffers[], void *cl_arg)
{
    cl_mem vector = (cl_mem) STARPU_VECTOR_GET_DEV_HANDLE(buffers[0]);
    unsigned offset = STARPU_BLOCK_GET_OFFSET(buffers[0]);

```

```
...
    clSetKernelArg(kernel, 0, sizeof(vector), &vector);
    clSetKernelArg(kernel, 1, sizeof(offset), &offset);
}
...
```

And the kernel has to shift from the pointer passed by the OpenCL driver:

```
__kernel void opencl_kernel(__global int *vector, unsigned offset)
{
    block = (__global void *)block + offset;
}
...
```

StarPU provides various interfaces and filters for matrices, vectors, etc., but applications can also write their own data interfaces and filters, see `examples/interface` and `examples/filters/custom_mf` for an example.

## 4.5 Performance Model Example

To achieve good scheduling, StarPU scheduling policies need to be able to estimate in advance the duration of a task. This is done by giving to codelets a performance model, by defining a structure `starpu_perfmodel` and providing its address in the field `starpu_codelet::model`. The fields `starpu_perfmodel::symbol` and `starpu_perfmodel::type` are mandatory, to give a name to the model, and the type of the model, since there are several kinds of performance models. For compatibility, make sure to initialize the whole structure to zero, either by using explicit `memset()`, or by letting the compiler implicitly do it as exemplified below.

- Measured at runtime (model type `STARPU_HISTORY_BASED`). This assumes that for a given set of data input/output sizes, the performance will always be about the same. This is very true for regular kernels on GPUs for instance (<0.1% error), and just a bit less true on CPUs ( $\sim=1\%$  error). This also assumes that there are few different sets of data input/output sizes. StarPU will then keep record of the average time of previous executions on the various processing units, and use it as an estimation. History is done per task size, by using a hash of the input and output sizes as an index. It will also save it in `$STARPU_HOME/.starpu/sampling/codelets` for further executions, and can be observed by using the tool `starpu_perfmodel_display`, or drawn by using the tool `starpu_perfmodel_plot` ([Performance Model Calibration](#)). The models are indexed by machine name. To share the models between machines (e.g. for a homogeneous cluster), use `export STARPU_HOSTNAME=some_global_name`. Measurements are only done when using a task scheduler which makes use of it, such as `dmda`. Measurements can also be provided explicitly by the application, by using the function `starpu_perfmodel_update_history()`.

The following is a small code example.

If e.g. the code is recompiled with other compilation options, or several variants of the code are used, the symbol string should be changed to reflect that, in order to recalibrate a new model from zero. The symbol string can even be constructed dynamically at execution time, as long as this is done before submitting any task using it.

```
static struct starpu_perfmodel mult_perf_model = {
    .type = STARPU_HISTORY_BASED,
    .symbol = "mult_perf_model"
};

struct starpu_codelet cl = {
    .where = STARPU_CPU,
    .cpu.funcs = { cpu_mult, NULL },
    .nbuffers = 3,
    .modes = { STARPU_R, STARPU_R, STARPU_W },
    /* for the scheduling policy to be able to use performance models */
    .model = &mult_perf_model
};
```

- Measured at runtime and refined by regression (model types `STARPU_REGRESSION_BASED` and `STARPU_NL_REGRESSION_BASED`). This still assumes performance regularity, but works with various data input sizes, by applying regression over observed execution times. `STARPU_REGRESSION_BASED` uses an  $a \cdot n^b$  regression form, `STARPU_NL_REGRESSION_BASED` uses an  $a \cdot n^b + c$  (more precise than `STARPU_REGRESSION_BASED`, but costs a lot more to compute).

For instance, `tests/perfmodels/regression_based.c` uses a regression-based performance model for the function `memset()`.

Of course, the application has to issue tasks with varying size so that the regression can be computed. StarPU will not trust the regression unless there is at least 10% difference between the minimum and maximum observed input size. It can be useful to set the environment variable `STARPU_CALIBRATE` to 1 and run the application on varying input sizes with `STARPU_SCHED` set to `dmda` scheduler, so as to feed the performance model for a variety of inputs. The application can also provide the measurements explicitly by using the function `starpu_perfmodel_update_history()`. The tools `starpu_perfmodel_display` and `starpu_perfmodel_plot` can be used to observe how much the performance model is calibrated ([Performance Model Calibration](#)); when their output look good, `STARPU_CALIBRATE` can be reset to 0 to let StarPU use the resulting performance model without recording new measures, and `STARPU_SCHED` can be set to `dmda` to benefit from the performance models. If the data input sizes vary a lot, it is really important to set `STARPU_CALIBRATE` to 0, otherwise StarPU will continue adding the measures, and result with a very big performance model, which will take time a lot of time to load and save.

For non-linear regression, since computing it is quite expensive, it is only done at termination of the application. This means that the first execution of the application will use only history-based performance model to perform scheduling, without using regression.

- Provided as an estimation from the application itself (model type `STARPU_COMMON` and field `starpu_perfmodel::cost_function`), see for instance `examples/commonblas_model.h` and `examples/commonblas_model.c`.
- Provided explicitly by the application (model type `STARPU_PER_ARCH`): the fields `.per_arch[arch][nimpl].cost_function` have to be filled with pointers to functions which return the expected duration of the task in micro-seconds, one per architecture.

For `STARPU_HISTORY_BASED`, `STARPU_REGRESSION_BASED`, and `STARPU_NL_REGRESSION_BASED`, the total size of task data (both input and output) is used as an index by default. The field `starpu_perfmodel::size_base` however permits the application to override that, when for instance some of the data do not matter for task cost (e.g. mere reference table), or when using sparse structures (in which case it is the number of non-zeros which matter), or when there is some hidden parameter such as the number of iterations, or when the application actually has a very good idea of the complexity of the algorithm, and just not the speed of the processor, etc. The example in the directory `examples/pi` uses this to include the number of iterations in the base.

StarPU will automatically determine when the performance model is calibrated, or rather, it will assume the performance model is calibrated until the application submits a task for which the performance can not be predicted. For `STARPU_HISTORY_BASED`, StarPU will require 10 (`_STARPU_CALIBRATION_MINIMUM`) measurements for a given size before estimating that an average can be taken as estimation for further executions with the same size. For `STARPU_REGRESSION_BASED` and `STARPU_NL_REGRESSION_BASED`, StarPU will require 10 (`_STARPU_CALIBRATION_MINIMUM`) measurements, and that the minimum measured data size is smaller than 90% of the maximum measured data size (i.e. the measurement interval is large enough for a regression to have a meaning). Calibration can also be forced by setting the `STARPU_CALIBRATE` environment variable to 1, or even reset by setting it to 2.

How to use schedulers which can benefit from such performance model is explained in [Task Scheduling Policy](#).

The same can be done for task power consumption estimation, by setting the field `starpu_codelet::power_model` the same way as the field `starpu_codelet::model`. Note: for now, the application has to give to the power consumption performance model a name which is different from the execution time performance model.

The application can request time estimations from the StarPU performance models by filling a task structure as usual without actually submitting it. The data handles can be created by calling any of the functions `starpu_*_data_register` with a `NULL` pointer and `-1` node and the desired data sizes, and need to be unregistered as usual. The functions `starpu_task_expected_length()` and `starpu_task_expected_power()` can then be called to get an estimation of the task cost on a given arch. `starpu_task_footprint()` can also be used to get the footprint used for indexing history-based performance models. `starpu_task_destroy()` needs to be called to destroy the dummy task afterwards. See `tests/perfmodels/regression_based.c` for an example.

## 4.6 Theoretical Lower Bound On Execution Time Example

For kernels with history-based performance models (and provided that they are completely calibrated), StarPU can very easily provide a theoretical lower bound for the execution time of a whole set of tasks. See for instance

examples/lu/lu\_example.c: before submitting tasks, call the function `starpu_bound_start()`, and after complete execution, call `starpu_bound_stop()`. `starpu_bound_print_lp()` or `starpu_bound_print_mps()` can then be used to output a Linear Programming problem corresponding to the schedule of your tasks. Run it through `lp_solve` or any other linear programming solver, and that will give you a lower bound for the total execution time of your tasks. If StarPU was compiled with the library `glpk` installed, `starpu_bound_compute()` can be used to solve it immediately and get the optimized minimum, in ms. Its parameter `integer` allows to decide whether integer resolution should be computed and returned

The `deps` parameter tells StarPU whether to take tasks, implicit data, and tag dependencies into account. Tags released in a callback or similar are not taken into account, only tags associated with a task are. It must be understood that the linear programming problem size is quadratic with the number of tasks and thus the time to solve it will be very long, it could be minutes for just a few dozen tasks. You should probably use `lp_solve -timeout 1 test.pl -wmps test.mps` to convert the problem to MPS format and then use a better solver, `glpsol` might be better than `lp_solve` for instance (the `--pcost` option may be useful), but sometimes doesn't manage to converge. `cbc` might look slower, but it is parallel. For `lp_solve`, be sure to try at least all the `-B` options. For instance, we often just use `lp_solve -cc -B1 -Bb -Bg -Bp -Bf -Br -BG -Bd -Bs -BB -Bo -Bc -Bi`, and the `-gr` option can also be quite useful. The resulting schedule can be observed by using the tool `starpu_lp2paje`, which converts it into the Paje format.

Data transfer time can only be taken into account when `deps` is set. Only data transfers inferred from implicit data dependencies between tasks are taken into account. Other data transfers are assumed to be completely overlapped.

Setting `deps` to 0 will only take into account the actual computations on processing units. It however still properly takes into account the varying performances of kernels and processing units, which is quite more accurate than just comparing StarPU performances with the fastest of the kernels being used.

The `prio` parameter tells StarPU whether to simulate taking into account the priorities as the StarPU scheduler would, i.e. schedule prioritized tasks before less prioritized tasks, to check to which extend this results to a less optimal solution. This increases even more computation time.

## 4.7 Insert Task Utility

StarPU provides the wrapper function `starpu_insert_task()` to ease the creation and submission of tasks.

Here the implementation of the codelet:

```
void func_cpu(void *descr[], void *_args)
{
    int *x0 = (int *)STARPU_VARIABLE_GET_PTR(descr[0]);
    float *x1 = (float *)STARPU_VARIABLE_GET_PTR(descr[1]);
    int ifactor;
    float ffactor;

    starpu_codelet_unpack_args(_args, &ifactor, &ffactor);
    *x0 = *x0 * ifactor;
    *x1 = *x1 * ffactor;
}

struct starpu_codelet mycodelet = {
    .where = STARPU_CPU,
    .cpu_funcs = { func_cpu, NULL },
    .nbufers = 2,
    .modes = { STARPU_RW, STARPU_RW }
};
```

And the call to the function `starpu_insert_task()`:

```
starpu_insert_task(&mycodelet,
                   STARPU_VALUE, &ifactor, sizeof(ifactor),
                   STARPU_VALUE, &ffactor, sizeof(ffactor),
                   STARPU_RW, data_handles[0], STARPU_RW, data_handles[1],
                   0);
```

The call to `starpu_insert_task()` is equivalent to the following code:

```
struct starpu_task *task = starpu_task_create();
task->c1 = &mycodelet;
task->handles[0] = data_handles[0];
task->handles[1] = data_handles[1];
char *arg_buffer;
size_t arg_buffer_size;
starpu_codelet_pack_args(&arg_buffer, &arg_buffer_size,
                        STARPU_VALUE, &ifactor, sizeof(ifactor),
```

```

        STARPU_VALUE, &ffactor, sizeof(ffactor),
        0);
task->cl.arg = arg.buffer;
task->cl.arg.size = arg.buffer.size;
int ret = starpu_task_submit(task);

```

Here a similar call using `STARPU_DATA_ARRAY`.

```

starpu_insert_task(&mycodelet,
                   STARPU_DATA_ARRAY, data_handles, 2,
                   STARPU_VALUE, &ifactor, sizeof(ifactor),
                   STARPU_VALUE, &ffactor, sizeof(ffactor),
                   0);

```

If some part of the task insertion depends on the value of some computation, the macro `STARPU_DATA_ACQUIRE_CB` can be very convenient. For instance, assuming that the index variable `i` was registered as handle `A_handle[i]`:

```

/* Compute which portion we will work on, e.g. pivot */
starpu_insert_task(&which_index, STARPU_W, i.handle, 0);

/* And submit the corresponding task */
STARPU_DATA_ACQUIRE_CB(i.handle, STARPU_R,
                       starpu_insert_task(&work, STARPU_RW, A_handle[i], 0));

```

The macro `STARPU_DATA_ACQUIRE_CB` submits an asynchronous request for acquiring data `i` for the main application, and will execute the code given as third parameter when it is acquired. In other words, as soon as the value of `i` computed by the codelet `which_index` can be read, the portion of code passed as third parameter of `STARPU_DATA_ACQUIRE_CB` will be executed, and is allowed to read from `i` to use it e.g. as an index. Note that this macro is only available when compiling StarPU with the compiler gcc.

## 4.8 Data Reduction

In various cases, some piece of data is used to accumulate intermediate results. For instances, the dot product of a vector, maximum/minimum finding, the histogram of a photograph, etc. When these results are produced along the whole machine, it would not be efficient to accumulate them in only one place, incurring data transmission each and access concurrency.

StarPU provides a mode `STARPU_REDUX`, which permits to optimize that case: it will allocate a buffer on each memory node, and accumulate intermediate results there. When the data is eventually accessed in the normal mode `STARPU_R`, StarPU will collect the intermediate results in just one buffer.

For this to work, the user has to use the function `starpu.data.set_reduction_methods()` to declare how to initialize these buffers, and how to assemble partial results.

For instance, `cg` uses that to optimize its dot product: it first defines the codelets for initialization and reduction:

```

struct starpu_codelet bzero_variable.cl =
{
    .cpu_funcs = { bzero.variable.cpu, NULL },
    .cuda_funcs = { bzero.variable.cuda, NULL },
    .nbuffers = 1,
}

static void accumulate_variable.cpu(void *descr[], void *cl_arg)
{
    double *v_dst = (double *)STARPU_VARIABLE_GET_PTR(descr[0]);
    double *v_src = (double *)STARPU_VARIABLE_GET_PTR(descr[1]);
    *v_dst = *v_dst + *v_src;
}

static void accumulate_variable.cuda(void *descr[], void *cl_arg)
{
    double *v_dst = (double *)STARPU_VARIABLE_GET_PTR(descr[0]);
    double *v_src = (double *)STARPU_VARIABLE_GET_PTR(descr[1]);
    cublasaxpy(1, (double)1.0, v_src, 1, v_dst, 1);
    cudaStreamSynchronize(starpu_cuda.get.local.stream());
}

struct starpu_codelet accumulate_variable.cl =
{
    .cpu_funcs = { accumulate.variable.cpu, NULL },
    .cuda_funcs = { accumulate.variable.cuda, NULL },
    .nbuffers = 1,
}

```

and attaches them as reduction methods for its handle `dtq`:

```
starpu_variable_data_register(&dtq_handle, -1, NULL, sizeof(
    type));
starpu_data_set_reduction_methods(dtq_handle,
    &accumulate_variable_cl, &bzero_variable_cl);
```

and `dtq_handle` can now be used in mode `STARPU_REDUX` for the dot products with partitioned vectors:

```
for (b = 0; b < nblocks; b++)
    starpu_insert_task(&dot_kernel_cl,
        STARPU_REDUX, dtq_handle,
        STARPU_R, starpu_data_get_sub_data(v1, 1, b),
        STARPU_R, starpu_data_get_sub_data(v2, 1, b),
        0);
```

During registration, we have here provided `NULL`, i.e. there is no initial value to be taken into account during reduction. StarPU will thus only take into account the contributions from the tasks `dot_kernel_cl`. Also, it will not allocate any memory for `dtq_handle` before tasks `dot_kernel_cl` are ready to run.

If another dot product has to be performed, one could unregister `dtq_handle`, and re-register it. But one can also call `starpu_data_invalidate_submit()` with the parameter `dtq_handle`, which will clear all data from the handle, thus resetting it back to the initial status `register(NULL)`.

The example `cg` also uses reduction for the blocked gemv kernel, leading to yet more relaxed dependencies and more parallelism.

`STARPU_REDUX` can also be passed to `starpu_mpi_insert_task()` in the MPI case. That will however not produce any MPI communication, but just pass `STARPU_REDUX` to the underlying `starpu.insert_task()`. It is up to the application to call `starpu_mpi_redux_data()`, which posts tasks that will reduce the partial results among MPI nodes into the MPI node which owns the data. For instance, some hypothetical application which collects partial results into data `res`, then uses it for other computation, before looping again with a new reduction:

```
for (i = 0; i < 100; i++) {
    starpu_mpi_insert_task(MPI_COMM_WORLD, &init_res,
        STARPU_W, res, 0);
    starpu_mpi_insert_task(MPI_COMM_WORLD, &work,
        STARPU_RW, A,
        STARPU_R, B, STARPU_REDUX, res, 0);
    starpu_mpi_redux_data(MPI_COMM_WORLD, res);
    starpu_mpi_insert_task(MPI_COMM_WORLD, &work2,
        STARPU_RW, B, STARPU_R, res, 0);
}
```

## 4.9 Temporary Buffers

There are two kinds of temporary buffers: temporary data which just pass results from a task to another, and scratch data which are needed only internally by tasks.

### 4.9.1 Temporary Data

Data can sometimes be entirely produced by a task, and entirely consumed by another task, without the need for other parts of the application to access it. In such case, registration can be done without prior allocation, by using the special memory node number `-1`, and passing a zero pointer. StarPU will actually allocate memory only when the task creating the content gets scheduled, and destroy it on unregistration.

In addition to that, it can be tedious for the application to have to unregister the data, since it will not use its content anyway. The unregistration can be done lazily by using the function `starpu_data_unregister_submit()`, which will record that no more tasks accessing the handle will be submitted, so that it can be freed as soon as the last task accessing it is over.

The following code exemplifies both points: it registers the temporary data, submits three tasks accessing it, and records the data for automatic unregistration.

```
starpu_vector_data_register(&handle, -1, 0, n, sizeof(float));
starpu_insert_task(&produce_data, STARPU_W, handle, 0);
starpu_insert_task(&compute_data, STARPU_RW, handle, 0);
starpu_insert_task(&summarize_data, STARPU_R, handle,
    STARPU_W, result.handle, 0);
starpu_data_unregister_submit(handle);
```

## 4.9.2 Scratch Data

Some kernels sometimes need temporary data to achieve the computations, i.e. a workspace. The application could allocate it at the start of the codelet function, and free it at the end, but that would be costly. It could also allocate one buffer per worker (similarly to [How To Initialize A Computation Library Once For Each Worker?](#)), but that would make them systematic and permanent. A more optimized way is to use the data access mode `STARPU_SCRATCH`, as exemplified below, which provides per-worker buffers without content consistency.

```
starpu_vector_data_register(&workspace, -1, 0, sizeof(float));
for (i = 0; i < N; i++)
    starpu_insert_task(&compute, STARPU_R, input[i],
                      STARPU_SCRATCH, workspace, STARPU_W, output[i], 0);
```

StarPU will make sure that the buffer is allocated before executing the task, and make this allocation per-worker: for CPU workers, notably, each worker has its own buffer. This means that each task submitted above will actually have its own workspace, which will actually be the same for all tasks running one after the other on the same worker. Also, if for instance GPU memory becomes scarce, StarPU will notice that it can free such buffers easily, since the content does not matter.

The example `examples/pi` uses scratches for some temporary buffer.

## 4.10 Parallel Tasks

StarPU can leverage existing parallel computation libraries by the means of parallel tasks. A parallel task is a task which gets worked on by a set of CPUs (called a parallel or combined worker) at the same time, by using an existing parallel CPU implementation of the computation to be achieved. This can also be useful to improve the load balance between slow CPUs and fast GPUs: since CPUs work collectively on a single task, the completion time of tasks on CPUs become comparable to the completion time on GPUs, thus relieving from granularity discrepancy concerns. `hwloc` support needs to be enabled to get good performance, otherwise StarPU will not know how to better group cores.

Two modes of execution exist to accomodate with existing usages.

### 4.10.1 Fork-mode Parallel Tasks

In the Fork mode, StarPU will call the codelet function on one of the CPUs of the combined worker. The codelet function can use `starpu_combined_worker_get_size()` to get the number of threads it is allowed to start to achieve the computation. The CPU binding mask for the whole set of CPUs is already enforced, so that threads created by the function will inherit the mask, and thus execute where StarPU expected, the OS being in charge of choosing how to schedule threads on the corresponding CPUs. The application can also choose to bind threads by hand, using e.g. `sched_getaffinity` to know the CPU binding mask that StarPU chose.

For instance, using OpenMP (full source is available in `examples/openmp/vector_scal.c`):

```
void scal_cpu_func(void *buffers[], void *_args)
{
    unsigned i;
    float *factor = _args;
    struct starpu_vector_interface *vector = buffers[0];
    unsigned n = STARPU_VECTOR_GET_NX(vector);
    float *val = (float *)STARPU_VECTOR_GET_PTR(vector);

#pragma omp parallel for num_threads(starpu_combined_worker_get_size())
    for (i = 0; i < n; i++)
        val[i] *= *factor;
}

static struct starpu_codelet cl =
{
    .modes = { STARPU_RW },
    .where = STARPU_CPU,
    .type = STARPU_FORKJOIN,
    .max.parallelism = INT_MAX,
    .cpu_funcs = {scal_cpu_func, NULL},
    .nbuffers = 1,
};
```

Other examples include for instance calling a BLAS parallel CPU implementation (see `examples/mult/xgemm.c`).

### 4.10.2 SPMD-mode Parallel Tasks

In the SPMD mode, StarPU will call the codelet function on each CPU of the combined worker. The codelet function can use `starpu_combined_worker_get_size()` to get the total number of CPUs involved in the combined worker, and thus the number of calls that are made in parallel to the function, and `starpu_combined_worker_get_rank()` to get the rank of the current CPU within the combined worker. For instance:

```
static void func(void *buffers[], void *args)
{
    unsigned i;
    float *factor = args;
    struct starpu_vector_interface *vector = buffers[0];
    unsigned n = STARPU_VECTOR_GET_NX(vector);
    float *val = (float *)STARPU_VECTOR_GET_PTR(vector);

    /* Compute slice to compute */
    unsigned m = starpu_combined_worker_get_size();
    unsigned j = starpu_combined_worker_get_rank();
    unsigned slice = (n+m-1)/m;

    for (i = j * slice; i < (j+1) * slice && i < n; i++)
        val[i] *= *factor;
}

static struct starpu_codelet cl =
{
    .modes = { STARPU_RW },
    .where = STARPU_CPU,
    .type = STARPU_SPMD,
    .max_parallelism = INT_MAX,
    .cpu.funcs = { func, NULL },
    .nbuffers = 1,
}
```

Of course, this trivial example will not really benefit from parallel task execution, and was only meant to be simple to understand. The benefit comes when the computation to be done is so that threads have to e.g. exchange intermediate results, or write to the data in a complex but safe way in the same buffer.

### 4.10.3 Parallel Tasks Performance

To benefit from parallel tasks, a parallel-task-aware StarPU scheduler has to be used. When exposed to codelets with a flag `STARPU_FORKJOIN` or `STARPU_SPMD`, the schedulers `pheft` (parallel-heft) and `peager` (parallel eager) will indeed also try to execute tasks with several CPUs. It will automatically try the various available combined worker sizes (making several measurements for each worker size) and thus be able to avoid choosing a large combined worker if the codelet does not actually scale so much.

### 4.10.4 Combined Workers

By default, StarPU creates combined workers according to the architecture structure as detected by `hwloc`. It means that for each object of the `hwloc` topology (NUMA node, socket, cache, ...) a combined worker will be created. If some nodes of the hierarchy have a big arity (e.g. many cores in a socket without a hierarchy of shared caches), StarPU will create combined workers of intermediate sizes. The variable `STARPU_SYNTHESIZE_ARITY_COMBINED_WORKER` permits to tune the maximum arity between levels of combined workers.

The combined workers actually produced can be seen in the output of the tool `starpu_machine_display` (the environment variable `STARPU_SCHED` has to be set to a combined worker-aware scheduler such as `pheft` or `peager`).

### 4.10.5 Concurrent Parallel Tasks

Unfortunately, many environments and libraries do not support concurrent calls.

For instance, most OpenMP implementations (including the main ones) do not support concurrent `pragma omp parallel` statements without nesting them in another `pragma omp parallel` statement, but StarPU does not yet support creating its CPU workers by using such pragma.

Other parallel libraries are also not safe when being invoked concurrently from different threads, due to the use of global variables in their sequential sections for instance.

The solution is then to use only one combined worker at a time. This can be done by setting the field `starpu_conf::single_combined_worker` to 1, or setting the environment variable `STARPU_SINGLE_COMBINED_WORKER` to 1. StarPU will then run only one parallel task at a time (but other CPU and GPU tasks are not affected and can

be run concurrently). The parallel task scheduler will however still try varying combined worker sizes to look for the most efficient ones.

## 4.11 Debugging

StarPU provides several tools to help debugging applications. Execution traces can be generated and displayed graphically, see [Generating Traces With FxT](#). Some gdb helpers are also provided to show the whole StarPU state:

```
(gdb) source tools/gdbinit
(gdb) help starpu
```

The Temanejo task debugger can also be used, see [Using The Temanejo Task Debugger](#).

## 4.12 The Multiformat Interface

It may be interesting to represent the same piece of data using two different data structures: one that would only be used on CPUs, and one that would only be used on GPUs. This can be done by using the multiformat interface. StarPU will be able to convert data from one data structure to the other when needed. Note that the scheduler dmdu is the only one optimized for this interface. The user must provide StarPU with conversion codelets:

```
#define NX 1024
struct point array_of_structs[NX];
starpu_data_handle_t handle;

/*
 * The conversion of a piece of data is itself a task, though it is created,
 * submitted and destroyed by StarPU internals and not by the user. Therefore,
 * we have to define two codelets.
 * Note that for now the conversion from the CPU format to the GPU format has to
 * be executed on the GPU, and the conversion from the GPU to the CPU has to be
 * executed on the CPU.
 */
#ifndef STARPU_USE_OPENCL
void cpu_to_opencl_func(void *buffers[], void *args);
struct starpu_codelet cpu_to_opencl_cl = {
    .where = STARPU_OPENCL,
    .opencl_funcs = { cpu_to_opencl_func, NULL },
    .nbuffers = 1,
    .modes = { STARPU_RW }
};

void opencl_to_cpu_func(void *buffers[], void *args);
struct starpu_codelet opencl_to_cpu_cl = {
    .where = STARPU_CPU,
    .cpu_funcs = { opencl_to_cpu_func, NULL },
    .nbuffers = 1,
    .modes = { STARPU_RW }
};
#endif

struct starpu_multiformat_data_interface_ops format_ops = {
#ifndef STARPU_USE_OPENCL
    .opencl_elemsize = 2 * sizeof(float),
    .cpu_to_opencl_cl = &cpu_to_opencl_cl,
    .opencl_to_cpu_cl = &opencl_to_cpu_cl,
#endif
    .cpu_elemsize = 2 * sizeof(float),
    ...
};

starpu_multiformat_data_register(handle, 0, &array_of_structs, NX, &
format_ops);
```

Kernels can be written almost as for any other interface. Note that `STARPU_MULTIFORMAT_GET_CPU_PTR` shall only be used for CPU kernels. CUDA kernels must use `STARPU_MULTIFORMAT_GET_CUDA_PTR`, and OpenCL kernels must use `STARPU_MULTIFORMAT_GET_OPENCL_PTR`. `STARPU_MULTIFORMAT_GET_NX` may be used in any kind of kernel.

```
static void
multiformat_scal_cpu_func(void *buffers[], void *args)
{
    struct point *aos;
    unsigned int n;

    aos = STARPU_MULTIFORMAT_GET_CPU_PTR(buffers[0]);
```

```

n = STARPU_MULTIFORMAT_GET_NX(buffers[0]);
...
}

extern "C" void multiformat_scal_cuda_func(void *buffers[], void *_args)
{
    unsigned int n;
    struct struct_of_arrays *soa;

    soa = (struct struct_of_arrays *) STARPU_MULTIFORMAT_GET_CUDA_PTR(
        buffers[0]);
    n = STARPU_MULTIFORMAT_GET_NX(buffers[0]);

    ...
}

```

A full example may be found in `examples/basic_examples/multiformat.c`.

## 4.13 Using The Driver API

### Running Drivers

```

int ret;
struct starpu_driver {
    .type = STARPU_CUDA_WORKER,
    .id.cuda_id = 0
};
ret = starpu_driver_init(&d);
if (ret != 0)
    error();
while (some_condition) {
    ret = starpu_driver_run_once(&d);
    if (ret != 0)
        error();
}
ret = starpu_driver_deinit(&d);
if (ret != 0)
    error();

```

To add a new kind of device to the structure `starpu.driver`, one needs to:

1. Add a member to the union `starpu_driver::id`
2. Modify the internal function `_starpu_launch_drivers()` to make sure the driver is not always launched.
3. Modify the function `starpu_driver_run()` so that it can handle another kind of architecture.
4. Write the new function `_starpu_run_foobar()` in the corresponding driver.

## 4.14 Defining A New Scheduling Policy

A full example showing how to define a new scheduling policy is available in the StarPU sources in the directory `examples/scheduler/`.

See [Scheduling Policy](#)

```

static struct starpu_sched_policy dummy_sched_policy = {
    .init_sched = init_dummy_sched,
    .deinit_sched = deinit_dummy_sched,
    .add_workers = dummy_sched.add_workers,
    .remove_workers = dummy_sched.remove_workers,
    .push_task = push_task_dummy,
    .push_prio_task = NULL,
    .pop_task = pop_task_dummy,
    .post_exec_hook = NULL,
    .pop_every_task = NULL,
    .policy_name = "dummy",
    .policy_description = "dummy scheduling strategy"
};

```

## 4.15 On-GPU Rendering

Graphical-oriented applications need to draw the result of their computations, typically on the very GPU where these happened. Technologies such as OpenGL/CUDA interoperability permit to let CUDA directly work on the OpenGL buffers, making them thus immediately ready for drawing, by mapping OpenGL buffer, textures or renderbuffer objects into CUDA. CUDA however imposes some technical constraints: peer memcpy has to be disabled, and the thread that runs OpenGL has to be the one that runs CUDA computations for that GPU.

To achieve this with StarPU, pass the option `--disable-cuda-memcpy-peer` to `./configure` (TODO: make it dynamic), OpenGL/GLUT has to be initialized first, and the interoperability mode has to be enabled by using the field `starpu_conf::cuda_opengl_interoperability`, and the driver loop has to be run by the application, by using the field `starpu_conf::not_launched_drivers` to prevent StarPU from running it in a separate thread, and by using `starpu_driver_run()` to run the loop. The examples `gl_interop` and `gl_interop_idle` show how it articulates in a simple case, where rendering is done in task callbacks. The former uses `glutMainLoopEvent` to make GLUT progress from the StarPU driver loop, while the latter uses `glutIdleFunc` to make StarPU progress from the GLUT main loop.

Then, to use an OpenGL buffer as a CUDA data, StarPU simply needs to be given the CUDA pointer at registration, for instance:

```
/* Get the CUDA worker id */
for (workerid = 0; workerid < starpu_worker_get_count(); workerid++)
    if (starpu_worker_get_type(workerid) ==
        STARPU_CUDA_WORKER)
        break;

/* Build a CUDA pointer pointing at the OpenGL buffer */
cudaGraphicsResourceGetMappedPointer((void**)&output, &num_bytes, resource);

/* And register it to StarPU */
starpu_vector_data_register(&handle,
                           starpu_worker_get_memory_node(workerid),
                           output, num_bytes / sizeof(float4), sizeof(float4));

/* The handle can now be used as usual */
starpu_insert_task(&c1, STARPU_RW, handle, 0);

/* ... */

/* This gets back data into the OpenGL buffer */
starpu_data_unregister(handle);
```

and display it e.g. in the callback function.

## 4.16 Defining A New Data Interface

Let's define a new data interface to manage complex numbers.

```
/* interface for complex numbers */
struct starpu_complex_interface
{
    double *real;
    double *imaginary;
    int nx;
};
```

Registering such a data to StarPU is easily done using the function `starpu_data_register()`. The last parameter of the function, `interface_complex_ops`, will be described below.

```
void starpu_complex_data_register(starpu_data_handle_t *handle,
                                  unsigned home_node, double *real, double *imaginary, int nx)
{
    struct starpu_complex_interface complex =
    {
        .real = real,
        .imaginary = imaginary,
        .nx = nx
    };

    if (interface_complex_ops.interfaceid ==
        STARPU_UNKNOWN_INTERFACE_ID)
    {
        interface_complex_ops.interfaceid =
        starpu_data_interface_get_next_id();
    }

    starpu_data_register(handleptr, home_node, &complex, &interface_complex_ops);
}
```

Different operations need to be defined for a data interface through the type `starpu_data_interface_ops`. We only define here the basic operations needed to run simple applications. The source code for the different functions can be found in the file `examples/interface/complex_interface.c`.

```
static struct starpu_data_interface_ops interface_complex.ops =
{
    .register_data_handle = complex_register_data_handle,
    .allocate_data_on_node = complex_allocate_data_on_node,
    .copy_methods = &complex_copy_methods,
    .get_size = complex_get_size,
    .footprint = complex_footprint,
    .interfaceid = STARPU_UNKNOWN_INTERFACE_ID,
    .interface_size = sizeof(struct starpu_complex_interface),
};
```

Functions need to be defined to access the different fields of the complex interface from a StarPU data handle.

```
double *starpu_complex_get_real(starpu_data_handle_t handle)
{
    struct starpu_complex_interface *complex_interface =
        (struct starpu_complex_interface *) starpu_data_get_interface_on_node
    (handle, 0);
    return complex_interface->real;
}

double *starpu_complex_get_imaginary(starpu_data_handle_t handle);
int starpu_complex_get_nx(starpu_data_handle_t handle);
```

Similar functions need to be defined to access the different fields of the complex interface from a `void *` pointer to be used within codelet implementations.

```
#define STARPU_COMPLEX_GET_REAL(interface) \
    ((struct starpu_complex_interface *) (interface))->real)
#define STARPU_COMPLEX_GET_IMAGINARY(interface) \
    ((struct starpu_complex_interface *) (interface))->imaginary)
#define STARPU_COMPLEX_GET_NX(interface) \
    ((struct starpu_complex_interface *) (interface))->nx)
```

Complex data interfaces can then be registered to StarPU.

```
double real = 45.0;
double imaginary = 12.0;
starpu_complex_data_register(&handle1, 0, &real, &imaginary, 1);
starpu_insert_task(&cl_display, STARPU_R, handle1, 0);
```

and used by codelets.

```
void display_complex_codelet(void *descr[], __attribute__((unused)) void *_args)
{
    int nx = STARPU_COMPLEX_GET_NX(descr[0]);
    double *real = STARPU_COMPLEX_GET_REAL(descr[0]);
    double *imaginary = STARPU_COMPLEX_GET_IMAGINARY(descr[0]);
    int i;

    for(i=0 ; i<nx ; i++)
    {
        fprintf(stderr, "Complex[%d] = %3.2f + %3.2f i\n", i, real[i], imaginary[i]);
    }
}
```

The whole code for this complex data interface is available in the directory `examples/interface/`.

## 4.17 Setting The Data Handles For A Task

The number of data a task can manage is fixed by the environment variable `STARPU_NMAXBUFS` which has a default value which can be changed through the configure option `--enable-maxbuffers`.

However, it is possible to define tasks managing more data by using the field `starpu_task::dyn_handles` when defining a task and the field `starpu_codelet::dyn_modes` when defining the corresponding codelet.

```
enum starpu_data_access_mode modes[STARPU_NMAXBUFS+1] = {
    STARPU_R, STARPU_R, ...
};

struct starpu_codelet dummy_big_cl =
{
    .cuda_funcs = {dummy_big_kernel, NULL},
    .opencl_funcs = {dummy_big_kernel, NULL},
```

```

.cpu_funcs = {dummy_big.kernel, NULL},
.nbuffers = STARPU_NMAXBUFS+1,
.dyn_modes = modes
};

task = starpu_task_create();
task->cl = &dummy_big.cl;
task->dyn_handles = malloc(task->cl->nbuffers * sizeof(
    starpu_data_handle_t));
for(i=0 ; i<task->cl->nbuffers ; i++)
{
    task->dyn_handles[i] = handle;
}
starpu_task_submit(task);

starpu_data_handle_t *handles = malloc(dummy_big.cl.nbuffers * sizeof(
    starpu_data_handle_t));
for(i=0 ; i<dummy_big.cl.nbuffers ; i++)
{
    handles[i] = handle;
}
starpu_insert_task(&dummy_big.cl,
    STARPU_VALUE, &dummy_big.cl.nbuffers, sizeof(dummy_big.cl.
nbuffers),
    STARPU_DATA_ARRAY, handles, dummy_big.cl.
nbuffers,
    0);

```

The whole code for this complex data interface is available in the directory `examples/basic-examples/dynamic_handles.c`.

## 4.18 More Examples

More examples are available in the StarPU sources in the directory `examples/`. Simple examples include:

**incrementer/** Trivial incrementation test.

**basic\_examples/** Simple documented Hello world and vector/scalar product (as shown in [Basic Examples](#)), matrix product examples (as shown in [Performance Model Example](#)), an example using the blocked matrix data interface, an example using the variable data interface, and an example using different formats on CPUs and GPUs.

**matvecmult/** OpenCL example from NVidia, adapted to StarPU.

**axpy/** AXPY CUBLAS operation adapted to StarPU.

**fortran/** Example of Fortran bindings.

More advanced examples include:

**filters/** Examples using filters, as shown in [Partitioning Data](#).

**lu/** LU matrix factorization, see for instance `xlu_implicit.c`

**cholesky/** Cholesky matrix factorization, see for instance `cholesky_implicit.c`.



# Chapter 5

## How To Optimize Performance With StarPU

TODO: improve!

Simply encapsulating application kernels into tasks already permits to seamlessly support CPU and GPUs at the same time. To achieve good performance, a few additional changes are needed.

### 5.1 Data Management

When the application allocates data, whenever possible it should use the function `starpu_malloc()`, which will ask CUDA or OpenCL to make the allocation itself and pin the corresponding allocated memory. This is needed to permit asynchronous data transfer, i.e. permit data transfer to overlap with computations. Otherwise, the trace will show that the `DriverCopyAsync` state takes a lot of time, this is because CUDA or OpenCL then reverts to synchronous transfers.

By default, StarPU leaves replicates of data wherever they were used, in case they will be re-used by other tasks, thus saving the data transfer time. When some task modifies some data, all the other replicates are invalidated, and only the processing unit which ran that task will have a valid replicate of the data. If the application knows that this data will not be re-used by further tasks, it should advise StarPU to immediately replicate it to a desired list of memory nodes (given through a bitmask). This can be understood like the write-through mode of CPU caches.

```
starpu_data_set_wt_mask(img_handle, 1<<0);
```

will for instance request to always automatically transfer a replicate into the main memory (node 0), as bit 0 of the write-through bitmask is being set.

```
starpu_data_set_wt_mask(img_handle, ~0U);
```

will request to always automatically broadcast the updated data to all memory nodes.

Setting the write-through mask to `~0U` can also be useful to make sure all memory nodes always have a copy of the data, so that it is never evicted when memory gets scarce.

Implicit data dependency computation can become expensive if a lot of tasks access the same piece of data. If no dependency is required on some piece of data (e.g. because it is only accessed in read-only mode, or because write accesses are actually commutative), use the function `starpu_data_set_sequential_consistency_flag()` to disable implicit dependencies on that data.

In the same vein, accumulation of results in the same data can become a bottleneck. The use of the mode `STARPU_REDUX` permits to optimize such accumulation (see [Data Reduction](#)).

Applications often need a data just for temporary results. In such a case, registration can be made without an initial value, for instance this produces a vector data:

```
starpu_vector_data_register(&handle, -1, 0, n, sizeof(float));
```

StarPU will then allocate the actual buffer only when it is actually needed, e.g. directly on the GPU without allocating in main memory.

In the same vein, once the temporary results are not useful any more, the data should be thrown away. If the handle is not to be reused, it can be unregistered:

```
starpu_unregister_submit(handle);
```

actual unregistration will be done after all tasks working on the handle terminate.

If the handle is to be reused, instead of unregistering it, it can simply be invalidated:

```
starpu_invalidate_submit(handle);
```

the buffers containing the current value will then be freed, and reallocated only when another task writes some value to the handle.

## 5.2 Task Granularity

Like any other runtime, StarPU has some overhead to manage tasks. Since it does smart scheduling and data management, that overhead is not always neglectable. The order of magnitude of the overhead is typically a couple of microseconds, which is actually quite smaller than the CUDA overhead itself. The amount of work that a task should do should thus be somewhat bigger, to make sure that the overhead becomes neglectable. The offline performance feedback can provide a measure of task length, which should thus be checked if bad performance are observed. To get a grasp at the scalability possibility according to task size, one can run `tests/microbenchs/tasks_size_overhead.sh` which draws curves of the speedup of independent tasks of very small sizes.

The choice of scheduler also has impact over the overhead: for instance, the scheduler `dmda` takes time to make a decision, while `eager` does not. `tasks_size_overhead.sh` can again be used to get a grasp at how much impact that has on the target machine.

## 5.3 Task Submission

To let StarPU make online optimizations, tasks should be submitted asynchronously as much as possible. Ideally, all the tasks should be submitted, and mere calls to `starpu_task_wait_for_all()` or `starpu_data_unregister()` be done to wait for termination. StarPU will then be able to rework the whole schedule, overlap computation with communication, manage accelerator local memory usage, etc.

## 5.4 Task Priorities

By default, StarPU will consider the tasks in the order they are submitted by the application. If the application programmer knows that some tasks should be performed in priority (for instance because their output is needed by many other tasks and may thus be a bottleneck if not executed early enough), the field `starpu_task::priority` should be set to transmit the priority information to StarPU.

## 5.5 Task Scheduling Policy

By default, StarPU uses the simple greedy scheduler `eager`. This is because it provides correct load balance even if the application codelets do not have performance models. If your application codelets have performance models ([Performance Model Example](#)), you should change the scheduler thanks to the environment variable `STARPU_SCHED`. For instance `export STARPU_SCHED=dmda`. Use `help` to get the list of available schedulers.

The `eager` scheduler uses a central task queue, from which workers draw tasks to work on. This however does not permit to prefetch data since the scheduling decision is taken late. If a task has a non-0 priority, it is put at the front of the queue.

The `prio` scheduler also uses a central task queue, but sorts tasks by priority (between -5 and 5).

The `random` scheduler distributes tasks randomly according to assumed worker overall performance.

The `ws` (work stealing) scheduler schedules tasks on the local worker by default. When a worker becomes idle, it steals a task from the most loaded worker.

The `dm` (deque model) scheduler uses task execution performance models into account to perform an HEFT-like scheduling strategy: it schedules tasks where their termination time will be minimal.

The `dmda` (deque model data aware) scheduler is similar to dm, it also takes into account data transfer time.

The `dmdar` (deque model data aware ready) scheduler is similar to dmda, it also sorts tasks on per-worker queues by number of already-available data buffers.

The `dmdas` (deque model data aware sorted) scheduler is similar to dmda, it also supports arbitrary priority values.

The `heft` (heterogeneous earliest finish time) scheduler is deprecated. It is now just an alias for `dmda`.

The `pheft` (parallel HEFT) scheduler is similar to heft, it also supports parallel tasks (still experimental). Should not be used when several contexts using it are being executed simultaneously.

The **peager** (parallel eager) scheduler is similar to eager, it also supports parallel tasks (still experimental). Should not be used when several contexts using it are being executed simultaneously.

## 5.6 Performance Model Calibration

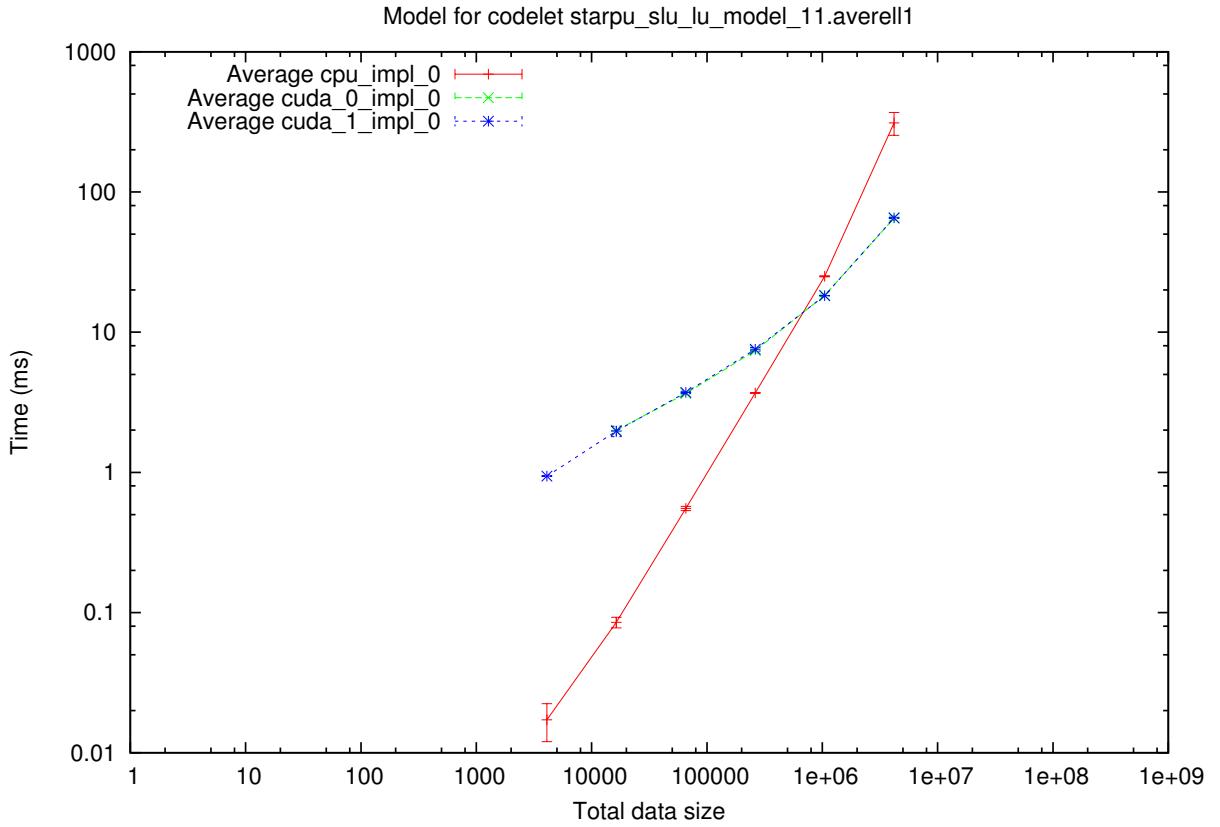
Most schedulers are based on an estimation of codelet duration on each kind of processing unit. For this to be possible, the application programmer needs to configure a performance model for the codelets of the application (see [Performance Model Example](#) for instance). History-based performance models use on-line calibration. StarPU will automatically calibrate codelets which have never been calibrated yet, and save the result in `$STARPU_HOME/.starpu/sampling/codelets`. The models are indexed by machine name. To share the models between machines (e.g. for a homogeneous cluster), use `export STARPU_HOSTNAME=some_global_name`. To force continuing calibration, use `export STARPU_CALIBRATE=1`. This may be necessary if your application has not-so-stable performance. StarPU will force calibration (and thus ignore the current result) until 10 (`_STARPU_CALIBRATION_MINIMUM`) measurements have been made on each architecture, to avoid badly scheduling tasks just because the first measurements were not so good. Details on the current performance model status can be obtained from the command `starpu_perfmodel_display`: the `-l` option lists the available performance models, and the `-s` option permits to choose the performance model to be displayed. The result looks like:

```
$ starpu_perfmodel_display -s starpu_slu_lu_model_11
performance model for cpu_impl_0
# hash      size      flops      mean      dev      n
914f3bef  1048576  0.000000e+00  2.503577e+04  1.982465e+02  8
3e921964   65536   0.000000e+00  5.527003e+02  1.848114e+01  7
e5a07e31    4096   0.000000e+00  1.717457e+01  5.190038e+00  14
...
...
```

Which shows that for the LU 11 kernel with a 1MiB matrix, the average execution time on CPUs was about 25ms, with a 0.2ms standard deviation, over 8 samples. It is a good idea to check this before doing actual performance measurements.

A graph can be drawn by using the tool `starpu_perfmodel_plot`:

```
$ starpu_perfmodel_plot -s starpu_slu_lu_model_11
4096 16384 65536 262144 1048576 4194304
$ gnuplot starpu_starpu_slu_lu_model_11.gp
$ gv starpu_starpu_slu_lu_model_11.eps
```



If a kernel source code was modified (e.g. performance improvement), the calibration information is stale and should be dropped, to re-calibrate from start. This can be done by using `export STARPU_CALIBRATE=2`.

Note: due to CUDA limitations, to be able to measure kernel duration, calibration mode needs to disable asynchronous data transfers. Calibration thus disables data transfer / computation overlapping, and should thus not be used for eventual benchmarks. Note 2: history-based performance models get calibrated only if a performance-model-based scheduler is chosen.

The history-based performance models can also be explicitly filled by the application without execution, if e.g. the application already has a series of measurements. This can be done by using `starpu_perfmodel_update_history()`, for instance:

```

static struct starpu_perfmodel perf_model = {
    .type = STARPU_HISTORY_BASED,
    .symbol = "my_perfmodel",
};

struct starpu_codelet cl = {
    .where = STARPU_CUDA,
    .cuda_funcs = { cuda_func1, cuda_func2, NULL },
    .nbuffers = 1,
    .modes = {STARPU_W},
    .model = &perf_model
};

void feed(void) {
    struct my_measure *measure;
    struct starpu_task task;
    starpu_task_init(&task);

    task.cl = &cl;

    for (measure = &measures[0]; measure < measures[last]; measure++) {
        starpu_data_handle_t handle;
        starpu_vector_data_register(&handle, -1, 0, measure->size, sizeof(float))
    };
    task.handles[0] = handle;
    starpu_perfmodel_update_history(&perf_model, &task,
                                   STARPU_CUDA_DEFAULT + measure->cudadev, 0,
                                   measure->implementation, measure->time);
    starpu_task_clean(&task);
    starpu_data_unregister(handle);
}
}

```

Measurement has to be provided in milliseconds for the completion time models, and in Joules for the energy consumption models.

## 5.7 Task Distribution Vs Data Transfer

Distributing tasks to balance the load induces data transfer penalty. StarPU thus needs to find a balance between both. The target function that the scheduler `dmda` of StarPU tries to minimize is  $\alpha * T_{execution} + \beta * T_{data\_transfer}$ , where  $T_{execution}$  is the estimated execution time of the codelet (usually accurate), and  $T_{data\_transfer}$  is the estimated data transfer time. The latter is estimated based on bus calibration before execution start, i.e. with an idle machine, thus without contention. You can force bus re-calibration by running the tool `starpu_calibrate_bus`. The beta parameter defaults to 1, but it can be worth trying to tweak it by using `export STARPU_SCHED_BETA=2` for instance, since during real application execution, contention makes transfer times bigger. This is of course imprecise, but in practice, a rough estimation already gives the good results that a precise estimation would give.

## 5.8 Data Prefetch

The scheduling policies `heft`, `dmda` and `pheft` perform data prefetch (see [STARPU\\_PREFETCH](#)): as soon as a scheduling decision is taken for a task, requests are issued to transfer its required data to the target processing unit, if needed, so that when the processing unit actually starts the task, its data will hopefully be already available and it will not have to wait for the transfer to finish.

The application may want to perform some manual prefetching, for several reasons such as excluding initial data transfers from performance measurements, or setting up an initial statically-computed data distribution on the machine before submitting tasks, which will thus guide StarPU toward an initial task distribution (since StarPU will try to avoid further transfers).

This can be achieved by giving the function `starpu_data_prefetch_on_node()` the handle and the desired target memory node.

## 5.9 Power-based Scheduling

If the application can provide some power performance model (through the field `starpu_codelet::power_model`), StarPU will take it into account when distributing tasks. The target function that the scheduler `dmda` minimizes becomes  $\alpha * T_{execution} + \beta * T_{data\_transfer} + \gamma * Consumption$ , where `Consumption` is the estimated task consumption in Joules. To tune this parameter, use `export STARPU_SCHED_GAMMA=3000` for instance, to express that each Joule (i.e kW during 1000us) is worth 3000us execution time penalty. Setting `alpha` and `beta` to zero permits to only take into account power consumption.

This is however not sufficient to correctly optimize power: the scheduler would simply tend to run all computations on the most energy-conservative processing unit. To account for the consumption of the whole machine (including idle processing units), the idle power of the machine should be given by setting `export STARPU_IDLE_POWER=200` for 200W, for instance. This value can often be obtained from the machine power supplier.

The power actually consumed by the total execution can be displayed by setting `export STARPU_PROFILING=1 STARPU_WORKER_STATS=1`.

On-line task consumption measurement is currently only supported through the `CL_PROFILING_POWER_CONSUMED` OpenCL extension, implemented in the MoviSim simulator. Applications can however provide explicit measurements by using the function `starpu_perfmodel_update_history()` (examplified in [Performance Model Example](#) with the `power_model` performance model). Fine-grain measurement is often not feasible with the feedback provided by the hardware, so the user can for instance run a given task a thousand times, measure the global consumption for that series of tasks, divide it by a thousand, repeat for varying kinds of tasks and task sizes, and eventually feed StarPU with these manual measurements through `starpu_perfmodel_update_history()`.

## 5.10 Static Scheduling

In some cases, one may want to force some scheduling, for instance force a given set of tasks to GPU0, another set to GPU1, etc. while letting some other tasks be scheduled on any other device. This can indeed be useful to guide

StarPU into some work distribution, while still letting some degree of dynamism. For instance, to force execution of a task on CUDA0:

```
task->execute_on_a_specific_worker = 1;
task->worker = starpu_worker_get_by_type(
    STARPU_CUDA_WORKER, 0);
```

Note however that using scheduling contexts while statically scheduling tasks on workers could be tricky. Be careful to schedule the tasks exactly on the workers of the corresponding contexts, otherwise the workers' corresponding scheduling structures may not be allocated or the execution of the application may deadlock. Moreover, the hypervisor should not be used when statically scheduling tasks.

## 5.11 Profiling

A quick view of how many tasks each worker has executed can be obtained by setting `export STARPU_WORKER_STATS=1`. This is a convenient way to check that execution did happen on accelerators without penalizing performance with the profiling overhead.

A quick view of how much data transfers have been issued can be obtained by setting `export STARPU_BUS_STATS=1`.

More detailed profiling information can be enabled by using `export STARPU_PROFILING=1` or by calling `starpu_profiling_status_set()` from the source code. Statistics on the execution can then be obtained by using `export STARPU_BUS_STATS=1` and `export STARPU_WORKER_STATS=1`. More details on performance feedback are provided by the next chapter.

## 5.12 Detection Stuck Conditions

It may happen that for some reason, StarPU does not make progress for a long period of time. Reason are sometimes due to contention inside StarPU, but sometimes this is due to external reasons, such as stuck MPI driver, or CUDA driver, etc.

```
export STARPU_WATCHDOG_TIMEOUT=10000
```

allows to make StarPU print an error message whenever StarPU does not terminate any task for 10ms. In addition to that,

```
export STARPU_WATCHDOG_CRASH=1
```

triggers a crash in that condition, thus allowing to catch the situation in gdb etc.

## 5.13 CUDA-specific Optimizations

Due to CUDA limitations, StarPU will have a hard time overlapping its own communications and the codelet computations if the application does not use a dedicated CUDA stream for its computations instead of the default stream, which synchronizes all operations of the GPU. StarPU provides one by the use of `starpu_cuda_get_local_stream()` which can be used by all CUDA codelet operations to avoid this issue. For instance:

```
func <<<grid,block,0,starpu_cuda_get_local_stream()>>> (foo, bar);
cudaStreamSynchronize(starpu_cuda_get_local_stream());
```

StarPU already does appropriate calls for the CUBLAS library.

Unfortunately, some CUDA libraries do not have stream variants of kernels. That will lower the potential for overlapping.

## 5.14 Performance Debugging

To get an idea of what is happening, a lot of performance feedback is available, detailed in the next chapter. The various informations should be checked for.

- What does the Gantt diagram look like? (see [Creating a Gantt Diagram](#))
  - If it's mostly green (tasks running in the initial context) or context specific color prevailing, then the machine is properly utilized, and perhaps the codelets are just slow. Check their performance, see [Performance Of Codelets](#).

- If it's mostly purple (FetchingInput), tasks keep waiting for data transfers, do you perhaps have far more communication than computation? Did you properly use CUDA streams to make sure communication can be overlapped? Did you use data-locality aware schedulers to avoid transfers as much as possible?
- If it's mostly red (Blocked), tasks keep waiting for dependencies, do you have enough parallelism? It might be a good idea to check what the DAG looks like (see [Creating a DAG With Graphviz](#)).
- If only some workers are completely red (Blocked), for some reason the scheduler didn't assign tasks to them. Perhaps the performance model is bogus, check it (see [Performance Of Codelets](#)). Do all your codelets have a performance model? When some of them don't, the schedulers switches to a greedy algorithm which thus performs badly.

You can also use the Temanejo task debugger (see [Using The Temanejo Task Debugger](#)) to visualize the task graph more easily.

## 5.15 Simulated Performance

StarPU can use Simgrid in order to simulate execution on an arbitrary platform.

### 5.15.1 Calibration

The idea is to first compile StarPU normally, and run the application, so as to automatically benchmark the bus and the codelets.

```
$ ./configure && make
$ STARPU_SCHED=dmda ./examples/matvecmult/matvecmult
[starpu] [_starpu_load_history_based_model] Warning: model matvecmult
    is not calibrated, forcing calibration for this run. Use the
    STARPU_CALIBRATE environment variable to control this.
$ ...
$ STARPU_SCHED=dmda ./examples/matvecmult/matvecmult
TEST PASSED
```

Note that we force to use the scheduler `dmda` to generate performance models for the application. The application may need to be run several times before the model is calibrated.

### 5.15.2 Simulation

Then, recompile StarPU, passing `--enable-simgrid` to `./configure`, and re-run the application:

```
$ ./configure --enable-simgrid && make
$ STARPU_SCHED=dmda ./examples/matvecmult/matvecmult
TEST FAILED !!!
```

It is normal that the test fails: since the computation are not actually done (that is the whole point of simgrid), the result is wrong, of course.

If the performance model is not calibrated enough, the following error message will be displayed

```
$ STARPU_SCHED=dmda ./examples/matvecmult/matvecmult
[starpu] [_starpu_load_history_based_model] Warning: model matvecmult
    is not calibrated, forcing calibration for this run. Use the
    STARPU_CALIBRATE environment variable to control this.
[starpu] [_starpu_simgrid_execute_job] [assert failure] Codelet
    matvecmult does not have a perfmodel, or is not calibrated enough
```

The number of devices can be chosen as usual with `STARPU_NCPU`, `STARPU_NCUDA`, and `STARPU_NOPE-NCL`. For now, only the number of cpus can be arbitrarily chosen. The number of CUDA and OpenCL devices have to be lower than the real number on the current machine.

The amount of simulated GPU memory is for now unbound by default, but it can be chosen by hand through the `STARPU_LIMIT_CUDA_MEM`, `STARPU_LIMIT_CUDA_devid_MEM`, `STARPU_LIMIT_OPENCL_MEM`, and `STARPU_LIMIT_OPENCL_devid_MEM` environment variables.

The Simgrid default stack size is small; to increase it use the parameter `--cfg=contexts/stack_size`, for example:

```
$ ./example --cfg=contexts/stack_size:8192
TEST FAILED !!!
```

Note: of course, if the application uses `gettimeofday` to make its performance measurements, the real time will be used, which will be bogus. To get the simulated time, it has to use `starpu_timing_now()` which returns the virtual timestamp in ms.

### 5.15.3 Simulation On Another Machine

The simgrid support even permits to perform simulations on another machine, your desktop, typically. To achieve this, one still needs to perform the Calibration step on the actual machine to be simulated, then copy them to your desktop machine (the `$STARPU_HOME/.starpu` directory). One can then perform the Simulation step on the desktop machine, by setting the environment variable `STARPU_HOSTNAME` to the name of the actual machine, to make StarPU use the performance models of the simulated machine even on the desktop machine.

If the desktop machine does not have CUDA or OpenCL, StarPU is still able to use simgrid to simulate execution with CUDA/OpenCL devices, but the application source code will probably disable the CUDA and OpenCL codelets in that case. Since during simgrid execution, the functions of the codelet are actually not called, one can use dummy functions such as the following to still permit CUDA or OpenCL execution:

```
static struct starpu_codelet chol_ll =
{
    .cpu_funcs = {chol.cpu_codelet.update_ull, NULL},
#ifndef STARPU_USE_CUDA
    .cuda_funcs = {chol.cublas_codelet.update_ull, NULL},
#endif
    .model = &chol.model_ll
};
```

# Chapter 6

# Performance Feedback

## 6.1 Using The Temanejo Task Debugger

StarPU can connect to Temanejo >= 1.0rc2 (see <http://www.hlrs.de/temanejo>), to permit nice visual task debugging. To do so, build Temanejo's libayudame.so, install Ayudame.h to e.g. /usr/local/include, apply the tools/patch-ayudame to it to fix C build, re-. ./configure, make sure that it found it, rebuild StarPU. Run the Temanejo GUI, give it the path to your application, any options you want to pass it, the path to libayudame.so.

Make sure to specify at least the same number of CPUs in the dialog box as your machine has, otherwise an error will happen during execution. Future versions of Temanejo should be able to tell StarPU the number of CPUs to use.

Tag numbers have to be below 4000000000000000ULL to be usable for Temanejo (so as to distinguish them from tasks).

## 6.2 On-line Performance Feedback

### 6.2.1 Enabling On-line Performance Monitoring

In order to enable online performance monitoring, the application can call `starpu_profiling_status_set()` with the parameter `STARPU_PROFILING_ENABLE`. It is possible to detect whether monitoring is already enabled or not by calling `starpu_profiling_status_get()`. Enabling monitoring also reinitialize all previously collected feedback. The environment variable `STARPU_PROFILING` can also be set to 1 to achieve the same effect. The function `starpu_profiling_init()` can also be called during the execution to reinitialize performance counters and to start the profiling if the environment variable `STARPU_PROFILING` is set to 1.

Likewise, performance monitoring is stopped by calling `starpu_profiling_status_set()` with the parameter `STARPU_PROFILING_DISABLE`. Note that this does not reset the performance counters so that the application may consult them later on.

More details about the performance monitoring API are available in [Profiling](#).

### 6.2.2 Per-task Feedback

If profiling is enabled, a pointer to a structure `starpu_profiling_task_info` is put in the field `starpu_task::profiling_info` when a task terminates. This structure is automatically destroyed when the task structure is destroyed, either automatically or by calling `starpu_task_destroy()`.

The structure `starpu_profiling_task_info` indicates the date when the task was submitted (`starpu_profiling_task_info::submit_time`), started (`starpu_profiling_task_info::start_time`), and terminated (`starpu_profiling_task_info::end_time`), relative to the initialization of StarPU with `starpu_init()`. It also specifies the identifier of the worker that has executed the task (`starpu_profiling_task_info::workerid`). These date are stored as `timespec` structures which the user may convert into micro-seconds using the helper function `starpu_timing_timespec_to_us()`.

It is worth noting that the application may directly access this structure from the callback executed at the end of the task. The structure `starpu_task` associated to the callback currently being executed is indeed accessible with the function `starpu_task_get_current()`.

### 6.2.3 Per-codelet Feedback

The field `starpu_codelet::per_worker_stats` is an array of counters. The i-th entry of the array is incremented every time a task implementing the codelet is executed on the i-th worker. This array is not reinitialized when profiling is enabled or disabled.

### 6.2.4 Per-worker Feedback

The second argument returned by the function `starpu_profiling_worker_get_info()` is a structure `starpu_profiling_worker.info` that gives statistics about the specified worker. This structure specifies when StarPU started collecting profiling information for that worker (`starpu_profiling_worker.info::start_time`), the duration of the profiling measurement interval (`starpu_profiling_worker.info::total_time`), the time spent executing kernels (`starpu_profiling_worker.info::executing_time`), the time spent sleeping because there is no task to execute at all (`starpu_profiling_worker.info::sleeping_time`), and the number of tasks that were executed while profiling was enabled. These values give an estimation of the proportion of time spent do real work, and the time spent either sleeping because there are not enough executable tasks or simply wasted in pure StarPU overhead.

Calling `starpu_profiling_worker_get_info()` resets the profiling information associated to a worker.

When an FxT trace is generated (see [Generating Traces With FxT](#)), it is also possible to use the tool `starpu_workers_activity` (see [Monitoring Activity](#)) to generate a graphic showing the evolution of these values during the time, for the different workers.

### 6.2.5 Bus-related Feedback

TODO: ajouter `STARPU_BUS_STATS`

# Chapter 7

## Tips and Tricks To Know About

### 7.1 How To Initialize A Computation Library Once For Each Worker?

Some libraries need to be initialized once for each concurrent instance that may run on the machine. For instance, a C++ computation class which is not thread-safe by itself, but for which several instanciated objects of that class can be used concurrently. This can be used in StarPU by initializing one such object per worker. For instance, the libstarpufft example does the following to be able to use FFTW on CPUs.

Some global array stores the instanciated objects:

```
fftw_plan plan_cpu[STARPU_NMAXWORKERS];
```

At initialisation time of libstarpuf, the objects are initialized:

```
int workerid;
for (workerid = 0; workerid < starpu_worker.get_count(); workerid++) {
    switch (starpu_worker.get_type(workerid)) {
        case STARPU_CPU_WORKER:
            plan_cpu[workerid] = fftw_plan(...);
            break;
    }
}
```

And in the codelet body, they are used:

```
static void fft(void *descr[], void *_args)
{
    int workerid = starpu_worker.get_id();
    fftw_plan plan = plan_cpu[workerid];
    ...

    fftw_execute(plan, ...);
}
```

This however is not sufficient for FFT on CUDA: initialization has to be done from the workers themselves. This can be done thanks to `starpu.execute_on_each_worker()`. For instance libstarpufft does the following.

```
static void fft_plan_gpu(void *args)
{
    plan plan = args;
    int n2 = plan->n2[0];
    int workerid = starpu_worker.get_id();

    cufftPlanId(&plan->plans[workerid].plan_cuda, n, _CUFFT_C2C, 1);
    cufftSetStream(plan->plans[workerid].plan_cuda, starpu_cuda_get_local_stream
                  ());
}
void starpufft_plan(void)
{
    starpu_execute_on_each_worker(fft_plan_gpu, plan,
                                 STARPU_CUDA);
}
```

### 7.2 How to limit memory per node

TODO

Talk about `STARPU_LIMIT_CUDA_devid_MEM`, `STARPU_LIMIT_CUDA_MEM`, `STARPU_LIMIT_OPENCL_devid_MEM`, `STARPU_LIMIT_OPENCL_MEM` and `STARPU_LIMIT_CPU_MEM`  
`starpu_memory_get_available()`

### 7.3 Thread Binding on NetBSD

When using StarPU on a NetBSD machine, if the topology discovery library `hwloc` is used, thread binding will fail. To prevent the problem, you should at least use the version 1.7 of `hwloc`, and also issue the following call:

```
$ sysctl -w security.models.extensions.user_set_cpu_affinity=1
```

Or add the following line in the file `/etc/sysctl.conf`

```
security.models.extensions.user_set_cpu_affinity=1
```

### 7.4 Using StarPU With MKL 11 (Intel Composer XE 2013)

Some users had issues with MKL 11 and StarPU (versions 1.1rc1 and 1.0.5) on Linux with MKL, using 1 thread for MKL and doing all the parallelism using StarPU (no multithreaded tasks), setting the environment variable `MKL_NUM_THREADS` to 1, and using the threaded MKL library, with `iomp5`.

Using this configuration, StarPU uses only 1 core, no matter the value of `STARPU_NCPU`. The problem is actually a thread pinning issue with MKL.

The solution is to set the environment variable `KMP_AFFINITY` to disabled ([http://software.intel.com/sites/products/documentation/studio/composer/en-us/2011Update/compiler-c/optaps/common/optaps\\_openmp\\_thread\\_affinity.htm](http://software.intel.com/sites/products/documentation/studio/composer/en-us/2011Update/compiler-c/optaps/common/optaps_openmp_thread_affinity.htm)).

# Chapter 8

## MPI Support

The integration of MPI transfers within task parallelism is done in a very natural way by the means of asynchronous interactions between the application and StarPU. This is implemented in a separate libstarpumpi library which basically provides "StarPU" equivalents of `MPI_*` functions, where `void *` buffers are replaced with `starpu.data_handle_t`, and all GPU-RAM-NIC transfers are handled efficiently by StarPU-MPI. The user has to use the usual `mpirun` command of the MPI implementation to start StarPU on the different MPI nodes.

An MPI Insert Task function provides an even more seamless transition to a distributed application, by automatically issuing all required data transfers according to the task graph and an application-provided distribution.

### 8.1 Simple Example

The flags required to compile or link against the MPI layer are accessible with the following commands:

```
$ pkg-config --cflags starpumpi-1.1    # options for the compiler  
$ pkg-config --libs starpumpi-1.1       # options for the linker
```

You also need pass the option `--static` if the application is to be linked statically.

```
void increment_token(void)  
{  
    struct starpu_task *task = starpu_task_create();  
  
    task->cl = &increment_cl;  
    task->handles[0] = token.handle;  
  
    starpu_task_submit(task);  
}  
  
int main(int argc, char **argv)  
{  
    int rank, size;  
  
    starpu_init(NULL);  
    starpu_mpi_initialize_extended(&rank, &size);  
  
    starpu_vector_data_register(&token.handle, 0, (uintptr_t)&token, 1, sizeof(  
        unsigned));  
  
    unsigned nloops = NITER;  
    unsigned loop;  
  
    unsigned last_loop = nloops - 1;  
    unsigned last_rank = size - 1;  
  
    for (loop = 0; loop < nloops; loop++) {  
        int tag = loop*size + rank;  
  
        if (loop == 0 && rank == 0)  
        {  
            token = 0;  
            fprintf(stdout, "Start with token value %d\n", token);  
        }  
        else  
        {  
            starpu_mpi_irecv_detached(token.handle, (rank+size-1)*size, tag,  
                MPI_COMM_WORLD, NULL, NULL);  
        }  
  
        increment_token();  
    }  
}
```

```

    if (loop == last_loop && rank == last_rank)
    {
        starpu_data_acquire(token_handle, STARPU_R);
        fprintf(stdout, "Finished: token value %d\n", token);
        starpu_data_release(token_handle);
    }
    else
    {
        starpu_mpi_isend_detached(token_handle, (rank+1)%size, tag+1,
                                    MPI_COMM_WORLD, NULL, NULL);
    }
}

starpu_task_wait_for_all();

starpu_mpi_shutdown();
starpu_shutdown();

if (rank == last_rank)
{
    fprintf(stderr, "[%d] token = %d == %d * %d ?\n", rank, token, nloops, size);
    STARPU_ASSERT(token == nloops*size);
}

```

## 8.2 Point To Point Communication

The standard point to point communications of MPI have been implemented. The semantic is similar to the MPI one, but adapted to the DSM provided by StarPU. A MPI request will only be submitted when the data is available in the main memory of the node submitting the request.

There are two types of asynchronous communications: the classic asynchronous communications and the detached communications. The classic asynchronous communications (`starpu_mpi_isend()` and `starpu_mpi_irecv()`) need to be followed by a call to `starpu_mpi_wait()` or to `starpu_mpi_test()` to wait for or to test the completion of the communication. Waiting for or testing the completion of detached communications is not possible, this is done internally by StarPU-MPI, on completion, the resources are automatically released. This mechanism is similar to the pthread detach state attribute which determines whether a thread will be created in a joinable or a detached state.

For any communication, the call of the function will result in the creation of a StarPU-MPI request, the function `starpu_data_acquire_cb()` is then called to asynchronously request StarPU to fetch the data in main memory; when the data is available in main memory, a StarPU-MPI function is called to put the new request in the list of the ready requests.

The StarPU-MPI progression thread regularly polls this list of ready requests. For each new ready request, the appropriate function is called to post the corresponding MPI call. For example, calling `starpu_mpi_isend()` will result in posting `MPI_Isend`. If the request is marked as detached, the request will be put in the list of detached requests.

The StarPU-MPI progression thread also polls the list of detached requests. For each detached request, it regularly tests the completion of the MPI request by calling `MPI_Test`. On completion, the data handle is released, and if a callback was defined, it is called.

[Communication](#) gives the list of all the point to point communications defined in StarPU-MPI.

## 8.3 Exchanging User Defined Data Interface

New data interfaces defined as explained in [Defining A New Data Interface](#) can also be used within StarPU-MPI and exchanged between nodes. Two functions need to be defined through the type `starpu_data_interface_ops`. The function `starpu_data_interface_ops::pack_data` takes a handle and returns a contiguous memory buffer along with its size where data to be conveyed to another node should be copied. The reversed operation is implemented in the function `starpu_data_interface_ops::unpack_data` which takes a contiguous memory buffer and recreates the data handle.

```

static int complex_pack_data(starpus_data_handle_t handle, unsigned node, void **ptr,
                           ssize_t *count)
{
    STARPU_ASSERT(starpus_data_test_if_allocated_on_node(handle, node));

    struct starpu_complex_interface *complex_interface =
        (struct starpu_complex_interface *) starpu_data_get_interface_on_node(
            handle, node);

    *count = complex_get_size(handle);

```

```

    *ptr = malloc(*count);
    memcpy(*ptr, complex_interface->real, complex_interface->nx*sizeof(double));
    memcpy(*ptr+complex_interface->nx*sizeof(double), complex_interface->imaginary,
           complex_interface->nx*sizeof(double));

    return 0;
}

static int complex_unpack_data(starpu_data_handle_t handle, unsigned node, void *ptr,
                               size_t count)
{
    STARPUL_ASSERT(starpu_data_test_if_allocated_on_node(handle, node));

    struct starpu_complex_interface *complex_interface =
        (struct starpu_complex_interface *) starpu_data_get_interface_on_node(
            handle, node);

    memcpy(complex_interface->real, ptr, complex_interface->nx*sizeof(double));
    memcpy(complex_interface->imaginary, ptr+complex_interface->nx*sizeof(double),
           complex_interface->nx*sizeof(double));

    return 0;
}

static struct starpu_data_interface_ops interface_complex_ops =
{
    ...
    .pack_data = complex_pack_data,
    .unpack_data = complex_unpack_data
};

```

## 8.4 MPI Insert Task Utility

To save the programmer from having to explicit all communications, StarPU provides an "MPI Insert Task Utility". The principle is that the application decides a distribution of the data over the MPI nodes by allocating it and notifying StarPU of that decision, i.e. tell StarPU which MPI node "owns" which data. It also decides, for each handle, an MPI tag which will be used to exchange the content of the handle. All MPI nodes then process the whole task graph, and StarPU automatically determines which node actually execute which task, and trigger the required MPI transfers.

The list of functions is described in [MPI Insert Task](#).

Here an stencil example showing how to use `starpu_mpi_insert_task()`. One first needs to define a distribution function which specifies the locality of the data. Note that that distribution information needs to be given to StarPU by calling `starpu_data_set_rank()`. A MPI tag should also be defined for each data handle by calling `starpu_data_set_tag()`.

```

/* Returns the MPI node number where data is */
int my_distrib(int x, int y, int nb_nodes) {
    /* Block distrib */
    return ((int)(x / sqrt(nb_nodes)) + (y / sqrt(nb_nodes)) * sqrt(nb_nodes)) % nb_nodes;

    // /* Other examples useful for other kinds of computations */
    // /* / distrib */
    // return (x+y) % nb_nodes;

    // /* Block cyclic distrib */
    // unsigned side = sqrt(nb_nodes);
    // return x % side + (y % side) * size;
}

```

Now the data can be registered within StarPU. Data which are not owned but will be needed for computations can be registered through the lazy allocation mechanism, i.e. with a `home_node` set to `-1`. StarPU will automatically allocate the memory when it is used for the first time.

One can note an optimization here (the `else if` test): we only register data which will be needed by the tasks that we will execute.

```

unsigned matrix[X][Y];
starpu_data_handle_t data_handles[X][Y];

for(x = 0; x < X; x++) {
    for(y = 0; y < Y; y++) {
        int mpi_rank = my_distrib(x, y, size);
        if (mpi_rank == my_rank)
            /* Owning data */
            starpu_variable_data_register(&data_handles[x][y], 0,
                                         (uintptr_t)&(matrix[x][y]), sizeof(unsigned));
        else if (my_rank == my_distrib(x+1, y, size) || my_rank == my_distrib(x-1, y, size)
                 || my_rank == my_distrib(x, y+1, size) || my_rank == my_distrib(x, y-1, size))
            /* I don't own that index, but will need it for my computations */
    }
}

```

```

        starpu.variable.data.register(&data_handles[x][y], -1,
                                      (uintptr_t)NULL, sizeof(unsigned));
    else
        /* I know it's useless to allocate anything for this */
        data_handles[x][y] = NULL;
    if (data_handles[x][y]) {
        starpu.data.set_rank(data_handles[x][y], mpi_rank);
        starpu.data.set_tag(data_handles[x][y], x*x+y);
    }
}

```

Now `starpu_mpi_insert_task()` can be called for the different steps of the application.

```

for(loop=0 ; loop<niter; loop++)
    for (x = 1; x < X-1; x++)
        for (y = 1; y < Y-1; y++)
            starpu_mpi.insert.task(MPI_COMM_WORLD, &stencil5_c1,
                STARPU_RW, data_handles[x][y],
                STARPU_R, data_handles[x-1][y],
                STARPU_R, data_handles[x+1][y],
                STARPU_R, data_handles[x][y-1],
                STARPU_R, data_handles[x][y+1],
                0);
starpu_task.wait_for_all();

```

I.e. all MPI nodes process the whole task graph, but as mentioned above, for each task, only the MPI node which owns the data being written to (here, `data_handles[x][y]`) will actually run the task. The other MPI nodes will automatically send the required data.

This can be a concern with a growing number of nodes. To avoid this, the application can prune the task for loops according to the data distribution, so as to only submit tasks on nodes which have to care about them (either to execute them, or to send the required data).

## 8.5 MPI Data migration

The application can dynamically change its mind about the data distribution, to balance the load over MPI nodes for instance. This can be done very simply by requesting an explicit move and then change the registered rank. For instance, we here switch to a new distribution function `my_distrib2`: we first register any data that wasn't registered already and will be needed, then migrate the data, and register the new location.

```

for(x = 0; x < X; x++) {
    for (y = 0; y < Y; y++) {
        int my_rank = my_distrib2(x, y, size);
        if (!data_handles[x][y] && (mpi.rank == my_rank
            || my_rank == my_distrib(x+1, y, size) || my_rank == my_distrib(x-1, y, size)
            || my_rank == my_distrib(x, y+1, size) || my_rank == my_distrib(x, y-1, size)))
            /* Register newly-needed data */
            starpu_variable_data_register(&data_handles[x][y], -1,
                (uintptr_t)NULL, sizeof(unsigned));
        if (data_handles[x][y]) {
            /* Migrate the data */
            starpu_mpi_get_data_on_node_detached(MPI_COMM_WORLD,
                data_handles[x][y], mpi_rank, NULL, NULL);
            /* And register the new rank of the matrix */
            starpu_data_set_rank(data_handles[x][y], mpi_rank);
        }
    }
}

```

From then on, further tasks submissions will use the new data distribution, which will thus change both MPI communications and task assignments.

Very importantly, since all nodes have to agree on which node owns which data so as to determine MPI communications and task assignments the same way, all nodes have to perform the same data migration, and at the same point among task submissions. It thus does not require a strict synchronization, just a clear separation of task submissions before and after the data redistribution.

Before data unregistration, it has to be migrated back to its original home node (the value, at least), since that is where the user-provided buffer resides. Otherwise the unregistration will complain that it does not have the latest value on the original home node.

```
for(x = 0; x < X; x++) {
    for (y = 0; y < Y; y++) {
        if (data_handles[x][y]) {
            int mpi_rank = my_distrib(x, y, size);
            /* Get back data to original place where the user-provided buffer is. */
        }
    }
}
```

```

        starpu_mpi_get_data_on_node_detached(MPI_COMM_WORLD,
data_handles[x][y], mpi.rank, NULL, NULL);
        /* And unregister it */
        starpu_data_unregister(data_handles[x][y]);
    }
}

```

## 8.6 MPI Collective Operations

The functions are described in [MPI Collective Operations](#).

```

if (rank == root)
{
    /* Allocate the vector */
    vector = malloc(nbblocks * sizeof(float *));
    for(x=0 ; x<nbblocks ; x++)
    {
        starpu_malloc((void **) &vector[x], block_size*sizeof(float));
    }

    /* Allocate data handles and register data to StarPU */
    data.handles = malloc(nbblocks*sizeof(starpu_data.handle_t *));
    for(x = 0; x < nbblocks ; x++)
    {
        int mpi_rank = my_distrib(x, nodes);
        if (rank == root) {
            starpu_vector_data_register(&data.handles[x], 0, (uintptr_t)vector[x],
                                         block_size, sizeof(float));
        }
        else if ((mpi_rank == rank) || ((rank == mpi_rank+1 || rank == mpi_rank-1))) {
            /* I own that index, or i will need it for my computations */
            starpu_vector_data_register(&data.handles[x], -1, (uintptr_t)NULL,
                                         block_size, sizeof(float));
        }
        else {
            /* I know it's useless to allocate anything for this */
            data.handles[x] = NULL;
        }
        if (data.handles[x]) {
            starpu_data_set_rank(data.handles[x], mpi_rank);
            starpu_data_set_tag(data.handles[x], x*nbblocks+y);
        }
    }

    /* Scatter the matrix among the nodes */
    starpu_mpi_scatter_detached(data.handles, nbblocks, root, MPI_COMM_WORLD);

    /* Calculation */
    for(x = 0; x < nbblocks ; x++) {
        if (data.handles[x]) {
            int owner = starpu_data_get_rank(data.handles[x]);
            if (owner == rank) {
                starpu_insert_task(&c1, STARPU_RW, data.handles[x], 0);
            }
        }
    }

    /* Gather the matrix on main node */
    starpu_mpi_gather_detached(data.handles, nbblocks, 0, MPI_COMM_WORLD);
}

```



# Chapter 9

## FFT Support

StarPU provides `libstarpufft`, a library whose design is very similar to both `fftw` and `cuffft`, the difference being that it takes benefit from both CPUs and GPUs. It should however be noted that GPUs do not have the same precision as CPUs, so the results may differ by a negligible amount.

Different precisions are available, namely float, double and long double precisions, with the following `fftw` naming conventions:

- double precision structures and functions are named e.g. `starpufft_execute()`
- float precision structures and functions are named e.g. `starpufftf_execute()`
- long double precision structures and functions are named e.g. `starpufftl_execute()`

The documentation below is given with names for double precision, replace `starpufft` with `starpufftf` or `starpufftl` as appropriate.

Only complex numbers are supported at the moment.

The application has to call `starpu_init()` before calling `starpufft` functions.

Either main memory pointers or data handles can be provided.

- To provide main memory pointers, use `starpufft_start()` or `starpufft_execute()`. Only one FFT can be performed at a time, because StarPU will have to register the data on the fly. In the `starpufft_start()` case, `starpufft_cleanup()` needs to be called to unregister the data.
- To provide data handles (which is preferable), use `starpufft_start_handle()` (preferred) or `starpufft_execute_handle()`. Several FFTs tasks can be submitted for a given plan, which permits e.g. to start a series of FFT with just one plan. `starpufft_start_handle()` is preferable since it does not wait for the task completion, and thus permits to enqueue a series of tasks.

All functions are defined in [FFT Support](#).

### 9.1 Compilation

The flags required to compile or link against the FFT library are accessible with the following commands:

```
$ pkg-config --cflags starpufft-1.1    # options for the compiler  
$ pkg-config --libs starpufft-1.1      # options for the linker
```

Also pass the option `--static` if the application is to be linked statically.



# Chapter 10

## C Extensions

When GCC plug-in support is available, StarPU builds a plug-in for the GNU Compiler Collection (GCC), which defines extensions to languages of the C family (C, C++, Objective-C) that make it easier to write StarPU code. This feature is only available for GCC 4.5 and later; it is known to work with GCC 4.5, 4.6, and 4.7. You may need to install a specific `-dev` package of your distro, such as `gcc-4.6-plugin-dev` on Debian and derivatives. In addition, the plug-in's test suite is only run when `GNU Guile` is found at configure-time. Building the GCC plug-in can be disabled by configuring with `--disable-gcc-extensions`.

Those extensions include syntactic sugar for defining tasks and their implementations, invoking a task, and manipulating data buffers. Use of these extensions can be made conditional on the availability of the plug-in, leading to valid C sequential code when the plug-in is not used ([Using C Extensions Conditionally](#)).

When StarPU has been installed with its GCC plug-in, programs that use these extensions can be compiled this way:

```
$ gcc -c -fplugin='pkg-config starpu-1.1 --variable=gccplugin' foo.c
```

When the plug-in is not available, the above `pkg-config` command returns the empty string.

In addition, the `-fplugin-arg-starpu-verbose` flag can be used to obtain feedback from the compiler as it analyzes the C extensions used in source files.

This section describes the C extensions implemented by StarPU's GCC plug-in. It does not require detailed knowledge of the StarPU library.

Note: this is still an area under development and subject to change.

### 10.1 Defining Tasks

The StarPU GCC plug-in views tasks as “extended” C functions:

- tasks may have several implementations—e.g., one for CPUs, one written in OpenCL, one written in CUDA;
- tasks may have several implementations of the same target—e.g., several CPU implementations;
- when a task is invoked, it may run in parallel, and StarPU is free to choose any of its implementations.

Tasks and their implementations must be *declared*. These declarations are annotated with attributes (<http://gcc.gnu.org/onlinedocs/gcc/Attribute-Syntax.html#Attribute-Syntax>): the declaration of a task is a regular C function declaration with an additional `task` attribute, and task implementations are declared with a `task_implementation` attribute.

The following function attributes are provided:

**task** Declare the given function as a StarPU task. Its return type must be `void`. When a function declared as `task` has a user-defined body, that body is interpreted as the implicit definition of the task's CPU implementation (see example below). In all cases, the actual definition of a task's body is automatically generated by the compiler.

Under the hood, declaring a task leads to the declaration of the corresponding `codelet` ([Codelet and Tasks](#)). If one or more task implementations are declared in the same compilation unit, then the codelet and the function itself are also defined; they inherit the scope of the task.

Scalar arguments to the task are passed by value and copied to the target device if need be—technically, they are passed as the buffer `starpu.task::cl.arg` ([Codelet and Tasks](#)).

Pointer arguments are assumed to be registered data buffers—the `handles` argument of a task (`starpu-task::handles`) ; `const`-qualified pointer arguments are viewed as read-only buffers (`STARPU_R`), and non-`const`-qualified buffers are assumed to be used read-write (`STARPU_RW`). In addition, the `output` type attribute can be as a type qualifier for output pointer or array parameters (`STARPU_W`).

**`task_implementation (target, task)`** Declare the given function as an implementation of `task` to run on `target`. `target` must be a string, currently one of "cpu", "opencl", or "cuda".

# Chapter 11

## SOCL OpenCL Extensions

SOCL is an OpenCL implementation based on StarPU. It gives a unified access to every available OpenCL device: applications can now share entities such as Events, Contexts or Command Queues between several OpenCL implementations.

In addition, command queues that are created without specifying a device provide automatic scheduling of the submitted commands on OpenCL devices contained in the context to which the command queue is attached.

Note: this is still an area under development and subject to change.

When compiling StarPU, SOCL will be enabled if a valid OpenCL implementation is found on your system. To be able to run the SOCL test suite, the environment variable `SOCL_OCL_LIB_OPENCL` needs to be defined to the location of the file `libOpenCL.so` of the OCL ICD implementation. You should for example add the following line in your file `.bashrc`

```
export SOCL_OCL_LIB_OPENCL=/usr/lib/x86_64-linux-gnu/libOpenCL.so
```

You can then run the test suite in the directory `socl/examples`.

```
$ make check
...
PASS: basic/basic
PASS: testmap/testmap
PASS: clinfo/clinfo
PASS: matmul/matmul
PASS: mansched/mansched
=====
All 5 tests passed
=====
```

The environment variable `OCL_ICD_VENDORS` has to point to the directory where the ICD files are installed. When compiling StarPU, the files are in the directory `socl/vendors`. With an installed version of StarPU, the files are installed in the directory `$prefix/share/starpu/opencl/vendors`.

To run the tests by hand, you have to call for example,

```
$ LD_PRELOAD=$SOCL_OCL_LIB_OPENCL OCL_ICD_VENDORS=socl/vendors/ socl/examples/clinfo/clinfo
Number of platforms: 2
Plaform Profile: FULL_PROFILE
Plaform Version: OpenCL 1.1 CUDA 4.2.1
Plaform Name: NVIDIA CUDA
Plaform Vendor: NVIDIA Corporation
Plaform Extensions: cl_khr_byte_addressable_store cl_khr_icd cl_khr_gl_sharing cl_nv_compiler_options cl_nv_
Plaform Profile: FULL_PROFILE
Plaform Version: OpenCL 1.0 SOCL Edition (0.1.0)
Plaform Name: SOCL Platform
Plaform Vendor: INRIA
Plaform Extensions: cl_khr_icd
...
$
```



# Chapter 12

## Scheduling Contexts

TODO: improve!

### 12.1 General Ideas

Scheduling contexts represent abstract sets of workers that allow the programmers to control the distribution of computational resources (i.e. CPUs and GPUs) to concurrent parallel kernels. The main goal is to minimize interferences between the execution of multiple parallel kernels, by partitioning the underlying pool of workers using contexts.

### 12.2 Creating A Context

By default, the application submits tasks to an initial context, which disposes of all the computation resources available to StarPU (all the workers). If the application programmer plans to launch several parallel kernels simultaneously, by default these kernels will be executed within this initial context, using a single scheduler policy(see [Task Scheduling Policy](#)). Meanwhile, if the application programmer is aware of the demands of these kernels and of the specificity of the machine used to execute them, the workers can be divided between several contexts. These scheduling contexts will isolate the execution of each kernel and they will permit the use of a scheduling policy proper to each one of them.

Scheduling Contexts may be created in two ways: either the programmers indicates the set of workers corresponding to each context (providing he knows the identifiers of the workers running within StarPU), or the programmer does not provide any worker list and leaves the Hypervisor assign workers to each context according to their needs ([Scheduling Context Hypervisor](#))

Both cases require a call to the function `starpu_sched_ctx_create`, which requires as input the worker list (the exact list or a NULL pointer) and the scheduling policy. The latter one can be a character list corresponding to the name of a StarPU predefined policy or the pointer to a custom policy. The function returns an identifier of the context created which you will use to indicate the context you want to submit the tasks to.

```
/* the list of resources the context will manage */
int workerids[3] = {1, 3, 10};

/* indicate the scheduling policy to be used within the context, the list of
   workers assigned to it, the number of workers, the name of the context */
int id_ctx = starpu_sched_ctx_create("dmdu", workerids, 3, "my_ctx");

/* let StarPU know that the following tasks will be submitted to this context */
starpu_sched_ctx_set_task_context(id);

/* submit the task to StarPU */
starpu_task_submit(task);
```

Note: Parallel greedy and parallel heft scheduling policies do not support the existence of several disjoint contexts on the machine. Combined workers are constructed depending on the entire topology of the machine, not only the one belonging to a context.

## 12.3 Modifying A Context

A scheduling context can be modified dynamically. The applications may change its requirements during the execution and the programmer can add additional workers to a context or remove if no longer needed. In the following example we have two scheduling contexts `sched_ctx1` and `sched_ctx2`. After executing a part of the tasks some of the workers of `sched_ctx1` will be moved to context `sched_ctx2`.

```
/* the list of resources that context 1 will give away */
int workerids[3] = {1, 3, 10};

/* add the workers to context 1 */
starpu_sched.ctx.add_workers(workerids, 3, sched_ctx2);

/* remove the workers from context 2 */
starpu_sched.ctx.remove_workers(workerids, 3, sched_ctx1);
```

## 12.4 Submitting Tasks To A Context

The application may submit tasks to several contexts either simultaneously or sequentially. If several threads of submission are used the function `starpu_sched_ctx_set_context` may be called just before `starpu_task_submit`. Thus StarPU considers that the current thread will submit tasks to the corresponding context.

When the application may not assign a thread of submission to each context, the id of the context must be indicated by using the function `starpu_task_submit_to_ctx` or the field `STARPU_SCHED_CTX` for `starpu_insert_task`.

## 12.5 Deleting A Context

When a context is no longer needed it must be deleted. The application can indicate which context should keep the resources of a deleted one. All the tasks of the context should be executed before doing this. Thus, the programmer may use either a barrier and then delete the context directly, or just indicate that other tasks will not be submitted later on to the context (such that when the last task is executed its workers will be moved to the inheritor) and delete the context at the end of the execution (when a barrier will be used eventually).

```
/* when the context 2 is deleted context 1 inherits its resources */
starpu_sched.ctx.set_inheritor(sched_ctx2, sched_ctx1);

/* submit tasks to context 2 */
for (i = 0; i < ntasks; i++)
    starpu_task_submit_to_ctx(task[i], sched_ctx2);

/* indicate that context 2 finished submitting and that */
/* as soon as the last task of context 2 finished executing */
/* its workers can be moved to the inheritor context */
starpu_sched.ctx.finished_submit(sched_ctx1);

/* wait for the tasks of both contexts to finish */
starpu_task.wait_for_all();

/* delete context 2 */
starpu_sched.ctx.delete(sched_ctx2);

/* delete context 1 */
starpu_sched.ctx.delete(sched_ctx1);
```

## 12.6 Emptying A Context

A context may have no resources at the begining or at a certain moment of the execution. Task can still be submitted to these contexts and they will be executed as soon as the contexts will have resources. A list of tasks pending to be executed is kept and when workers are added to the contexts these tasks start being submitted. However, if resources are never allocated to the context the program will not terminate. If these tasks have low priority the programmer can forbid the application to submit them by calling the function `starpu_sched_ctx_stop_task_submission()`.

## 12.7 Contexts Sharing Workers

Contexts may share workers when a single context cannot execute efficiently enough alone on these workers or when the application decides to express a hierarchy of contexts. The workers apply an algorithm of “Round-Robin” to chose the context on which they will “pop” next. By using the function `starpusched_ctx_set_turn_to_other_ctx`, the programmer can impose the `workerid` to “pop” in the context `sched_ctx_id` next.



# Chapter 13

## Scheduling Context Hypervisor

### 13.1 What Is The Hypervisor

StarPU proposes a platform to construct Scheduling Contexts, to delete and modify them dynamically. A parallel kernel, can thus be isolated into a scheduling context and interferences between several parallel kernels are avoided. If the user knows exactly how many workers each scheduling context needs, he can assign them to the contexts at their creation time or modify them during the execution of the program.

The Scheduling Context Hypervisor Plugin is available for the users who do not dispose of a regular parallelism, who cannot know in advance the exact size of the context and need to resize the contexts according to the behavior of the parallel kernels.

The Hypervisor receives information from StarPU concerning the execution of the tasks, the efficiency of the resources, etc. and it decides accordingly when and how the contexts can be resized. Basic strategies of resizing scheduling contexts already exist but a platform for implementing additional custom ones is available.

### 13.2 Start the Hypervisor

The Hypervisor must be initialized once at the beginning of the application. At this point a resizing policy should be indicated. This strategy depends on the information the application is able to provide to the hypervisor as well as on the accuracy needed for the resizing procedure. For example, the application may be able to provide an estimation of the workload of the contexts. In this situation the hypervisor may decide what resources the contexts need. However, if no information is provided the hypervisor evaluates the behavior of the resources and of the application and makes a guess about the future. The hypervisor resizes only the registered contexts.

### 13.3 Interrogate The Runtime

The runtime provides the hypervisor with information concerning the behavior of the resources and the application. This is done by using the `performance_counters` which represent callbacks indicating when the resources are idle or not efficient, when the application submits tasks or when it becomes to slow.

### 13.4 Trigger the Hypervisor

The resizing is triggered either when the application requires it (`sc_hypervisor_resize_ctxs`) or when the initials distribution of resources alters the performance of the application (the application is to slow or the resource are idle for too long time). If the environment variable `SC_HYPERVERSOR_TRIGGER_RESIZE` is set to `speed` the monitored speed of the contexts is compared to a theoretical value computed with a linear program, and the resizing is triggered whenever the two values do not correspond. Otherwise, if the environment variable is set to `idle` the hypervisor triggers the resizing algorithm whenever the workers are idle for a period longer than the threshold indicated by the programmer. When this happens different resizing strategy are applied that target minimizing the total execution of the application, the instant speed or the idle time of the resources.

## 13.5 Resizing Strategies

The plugin proposes several strategies for resizing the scheduling context.

The **Application driven** strategy uses the user's input concerning the moment when he wants to resize the contexts. Thus, the users tags the task that should trigger the resizing process. We can set directly the field `starpu_task::hypervisor_tag` or use the macro `::STARPU_HYPERVISOR_TAG` in the function `starpu_insert_task()`.

```
task.hypervisor.tag = 2;
```

or

```
starpu_insert_task(&codelet,
    ...,
    STARPU_HYPERVISOR_TAG, 2,
    0);
```

Then the user has to indicate that when a task with the specified tag is executed the contexts should resize.

```
sc.hypervisor.resize(sched_ctx, 2);
```

The user can use the same tag to change the resizing configuration of the contexts if he considers it necessary.

```
sc.hypervisor.ctl(sched_ctx,
    SC_HYPERVISOR_MIN_WORKERS, 6,
    SC_HYPERVISOR_MAX_WORKERS, 12,
    SC_HYPERVISOR_TIME_TO_APPLY, 2,
    NULL);
```

The **Idleness** based strategy moves workers unused in a certain context to another one needing them. (see Users' Input In The Resizing Process)

```
int workerids[3] = {1, 3, 10};
int workerids2[9] = {0, 2, 4, 5, 6, 7, 8, 9, 11};
sc.hypervisor.ctl(sched_ctx_id,
    SC_HYPERVISOR_MAX_IDLE, workerids, 3, 10000.0,
    SC_HYPERVISOR_MAX_IDLE, workerids2, 9, 50000.0,
    NULL);
```

The **Gflops rate** based strategy resizes the scheduling contexts such that they all finish at the same time. The speed of each of them is computed and once one of them is significantly slower the resizing process is triggered. In order to do these computations the user has to input the total number of instructions needed to be executed by the parallel kernels and the number of instruction to be executed by each task.

The number of flops to be executed by a context are passed as parameter when they are registered to the hypervisor, (`sc.hypervisor.register_ctx(sched_ctx_id, flops)`) and the one to be executed by each task are passed when the task is submitted. The corresponding field is `starpu_task::flops` and the corresponding macro in the function `starpu_insert_task()` is `STARPU_FLOPS` (**Caution:** but take care of passing a double, not an integer, otherwise parameter passing will be bogus). When the task is executed the resizing process is triggered.

```
task.flops = 100;
```

or

```
starpu_insert_task(&codelet,
    ...,
    STARPU_FLOPS, (double) 100,
    0);
```

The **Feft** strategy uses a linear program to predict the best distribution of resources such that the application finishes in a minimum amount of time. As for the **Gflops rate** strategy the programmers has to indicate the total number of flops to be executed when registering the context. This number of flops may be updated dynamically during the execution of the application whenever this information is not very accurate from the beginning. The function `sc.hypervisor.update_diff_total_flop` is called in order add or remove a difference to the flops left to be executed. Tasks are provided also the number of flops corresponding to each one of them. During the execution of the application the hypervisor monitors the consumed flops and recomputes the time left and the number of resources to use. The speed of each type of resource is (re)evaluated and inserter in the linear program in order to better adapt to the needs of the application.

The **Feft** strategy uses a linear program too, that considers all the types of tasks and the number of each of them and it tries to allocates resources such that the application finishes in a minimum amount of time. A previous calibration of StarPU would be useful in order to have good predictions of the execution time of each type of task.

The types of tasks may be determined directly by the hypervisor when they are submitted. However there are applications that do not expose all the graph of tasks from the beginning. In this case in order to let the hypervisor know about all the tasks the function `sc_hypervisor_set_type_of_task` will just inform the hypervisor about future tasks without submitting them right away.

The **Ispeed** strategy divides the execution of the application in several frames. For each frame the hypervisor computes the speed of the contexts and tries making them run at the same speed. The strategy requires less contribution from the user as the hypervisor requires only the size of the frame in terms of flops.

```
int workerids[3] = {1, 3, 10};
int workerids2[9] = {0, 2, 4, 5, 6, 7, 8, 9, 11};
sc_hypervisor_ctl(sched_ctx_id,
    SC_HYPERVISOR_ISPEED_W_SAMPLE, workerids, 3, 2000000000.0,
    SC_HYPERVISOR_ISPEED_W_SAMPLE, workerids2, 9, 200000000000.0
    ,
    SC_HYPERVISOR_ISPEED_CTX_SAMPLE, 6000000000.0,
    NULL);
```

The **Throughput** strategy focuses on maximizing the throughput of the resources and resizes the contexts such that the machine is running at its maximum efficiency (maximum instant speed of the workers).

## 13.6 a new hypervisor policy

While Scheduling Context Hypervisor Plugin comes with a variety of resizing policies (see [Resizing Strategies](#)), it may sometimes be desirable to implement custom policies to address specific problems. The API described below allows users to write their own resizing policy.

Here an example of how to define a new policy

```
struct sc_hypervisor_policy dummy_policy =
{
    .handle.poped_task = dummy_handle_poped_task,
    .handle.pushed.task = dummy_handle_pushed_task,
    .handle.idle_cycle = dummy_handle_idle_cycle,
    .handle.idle_end = dummy_handle_idle_end,
    .handle.post_exec_hook = dummy_handle_post_exec_hook,
    .custom = 1,
    .name = "dummy"
};
```



## **Part II**

# **Inside StarPU**



## Chapter 14

# Execution Configuration Through Environment Variables

The behavior of the StarPU library and tools may be tuned thanks to the following environment variables.

### 14.1 Configuring Workers

**STARPU\_NCPU** Specify the number of CPU workers (thus not including workers dedicated to control accelerators). Note that by default, StarPU will not allocate more CPU workers than there are physical CPUs, and that some CPUs are used to control the accelerators.

**STARPU\_NCPUS** This variable is deprecated. You should use [STARPU\\_NCPU](#).

**STARPU\_NCUDA** Specify the number of CUDA devices that StarPU can use. If [STARPU\\_NCUDA](#) is lower than the number of physical devices, it is possible to select which CUDA devices should be used by the means of the environment variable [STARPU\\_WORKERS\\_CUDAID](#). By default, StarPU will create as many CUDA workers as there are CUDA devices.

**STARPU\_NOPENCL** OpenCL equivalent of the environment variable [STARPU\\_NCUDA](#).

**STARPU\_OPENCL\_ON\_CPUS** By default, the OpenCL driver only enables GPU and accelerator devices. By setting the environment variable [STARPU\\_OPENCL\\_ON\\_CPUS](#) to 1, the OpenCL driver will also enable CPU devices.

**STARPU\_OPENCL\_ONLY\_ON\_CPUS** By default, the OpenCL driver enables GPU and accelerator devices. By setting the environment variable [STARPU\\_OPENCL\\_ONLY\\_ON\\_CPUS](#) to 1, the OpenCL driver will ONLY enable CPU devices.

**STARPU\_WORKERS\_NOBIND** Setting it to non-zero will prevent StarPU from binding its threads to CPUs. This is for instance useful when running the testsuite in parallel.

**STARPU\_WORKERS\_CPUID** Passing an array of integers (starting from 0) in [STARPU\\_WORKERS\\_CPUID](#) specifies on which logical CPU the different workers should be bound. For instance, if [STARPU\\_WORKERS\\_CPUID](#) = "0 1 4 5", the first worker will be bound to logical CPU #0, the second CPU worker will be bound to logical CPU #1 and so on. Note that the logical ordering of the CPUs is either determined by the OS, or provided by the library `hwloc` in case it is available.

Note that the first workers correspond to the CUDA workers, then come the OpenCL workers, and finally the CPU workers. For example if we have [STARPU\\_NCUDA](#)=1, [STARPU\\_NOPENCL](#)=1, [STARPU\\_NCPU](#)=2 and [STARPU\\_WORKERS\\_CPUID](#) = "0 2 1 3", the CUDA device will be controlled by logical CPU #0, the OpenCL device will be controlled by logical CPU #2, and the logical CPUs #1 and #3 will be used by the CPU workers.

If the number of workers is larger than the array given in [STARPU\\_WORKERS\\_CPUID](#), the workers are bound to the logical CPUs in a round-robin fashion: if [STARPU\\_WORKERS\\_CPUID](#) = "0 1", the first and the third (resp. second and fourth) workers will be put on CPU #0 (resp. CPU #1).

This variable is ignored if the field `starpu_conf::use_explicit_workers_bindid` passed to `starpu_init()` is set.

**STARPU\_WORKERS\_CUDAID** Similarly to the `STARPU_WORKERS_CPUID` environment variable, it is possible to select which CUDA devices should be used by StarPU. On a machine equipped with 4 GPUs, setting `S-TARPU_WORKERS_CUDAID = "1 3"` and `STARPU_NCUDA=2` specifies that 2 CUDA workers should be created, and that they should use CUDA devices #1 and #3 (the logical ordering of the devices is the one reported by CUDA).

This variable is ignored if the field `starpu_conf::use_explicit_workers_cuda_gpid` passed to `starpu_init()` is set.

**STARPU\_WORKERS\_OPENCLID** OpenCL equivalent of the `STARPU_WORKERS_CUDAID` environment variable.

This variable is ignored if the field `starpu_conf::use_explicit_workers_opencl_gpid` passed to `starpu_init()` is set.

**STARPU\_SINGLE\_COMBINED\_WORKER** If set, StarPU will create several workers which won't be able to work concurrently. It will by default create combined workers which size goes from 1 to the total number of CPU workers in the system. `STARPU_MIN_WORKERSIZE` and `STARPU_MAX_WORKERSIZE` can be used to change this default.

**STARPU\_MIN\_WORKERSIZE** `STARPU_MIN_WORKERSIZE` permits to specify the minimum size of the combined workers (instead of the default 2)

**STARPU\_MAX\_WORKERSIZE** `STARPU_MAX_WORKERSIZE` permits to specify the minimum size of the combined workers (instead of the number of CPU workers in the system)

**STARPU\_SYNTHESIZE\_ARITY\_COMBINED\_WORKER** Let the user decide how many elements are allowed between combined workers created from hwloc information. For instance, in the case of sockets with 6 cores without shared L2 caches, if `STARPU_SYNTHESIZE_ARITY_COMBINED_WORKER` is set to 6, no combined worker will be synthesized beyond one for the socket and one per core. If it is set to 3, 3 intermediate combined workers will be synthesized, to divide the socket cores into 3 chunks of 2 cores. If it set to 2, 2 intermediate combined workers will be synthesized, to divide the the socket cores into 2 chunks of 3 cores, and then 3 additional combined workers will be synthesized, to divide the former synthesized workers into a bunch of 2 cores, and the remaining core (for which no combined worker is synthesized since there is already a normal worker for it).

The default, 2, thus makes StarPU tend to building a binary trees of combined workers.

**STARPU\_DISABLE\_ASYNCNCHRONOUS\_COPY** Disable asynchronous copies between CPU and GPU devices. The AMD implementation of OpenCL is known to fail when copying data asynchronously. When using this implementation, it is therefore necessary to disable asynchronous data transfers.

**STARPU\_DISABLE\_ASYNCNCHRONOUS\_CUDA\_COPY** Disable asynchronous copies between CPU and CUDA devices.

**STARPU\_DISABLE\_ASYNCNCHRONOUS\_OPENCL\_COPY** Disable asynchronous copies between CPU and OpenCL devices. The AMD implementation of OpenCL is known to fail when copying data asynchronously. When using this implementation, it is therefore necessary to disable asynchronous data transfers.

**STARPU\_ENABLE\_CUDA\_GPU\_GPU\_DIRECT** Enable (1) or Disable (0) direct CUDA transfers from GPU to GPU, without copying through RAM. The default is Enabled. This permits to test the performance effect of GPU-Direct.

## 14.2 Configuring The Scheduling Engine

**STARPU\_SCHED** Choose between the different scheduling policies proposed by StarPU: work random, stealing, greedy, with performance models, etc.

Use `STARPU_SCHED=help` to get the list of available schedulers.

**STARPU\_CALIBRATE** If this variable is set to 1, the performance models are calibrated during the execution. If it is set to 2, the previous values are dropped to restart calibration from scratch. Setting this variable to 0 disable calibration, this is the default behaviour.

Note: this currently only applies to `dm` and `dmda` scheduling policies.

**STARPU\_BUS\_CALIBRATE** If this variable is set to 1, the bus is recalibrated during initialization.

**STARPU\_PREFETCH** This variable indicates whether data prefetching should be enabled (0 means that it is disabled). If prefetching is enabled, when a task is scheduled to be executed e.g. on a GPU, StarPU will request an asynchronous transfer in advance, so that data is already present on the GPU when the task starts. As a result, computation and data transfers are overlapped. Note that prefetching is enabled by default in StarPU.

**STARPU\_SCHED\_ALPHA** To estimate the cost of a task StarPU takes into account the estimated computation time (obtained thanks to performance models). The alpha factor is the coefficient to be applied to it before adding it to the communication part.

**STARPU\_SCHED\_BETA** To estimate the cost of a task StarPU takes into account the estimated data transfer time (obtained thanks to performance models). The beta factor is the coefficient to be applied to it before adding it to the computation part.

**STARPU\_SCHED\_GAMMA** Define the execution time penalty of a joule ([Power-based Scheduling](#)).

**STARPU\_IDLE\_POWER** Define the idle power of the machine ([Power-based Scheduling](#)).

**STARPU\_PROFILING** Enable on-line performance monitoring ([Enabling On-line Performance Monitoring](#)).

## 14.3 Extensions

**SOCL\_OCL\_LIB\_OPENCL** THE SOCL test suite is only run when the environment variable **SOCL\_OCL\_LIB\_OPE-NCL** is defined. It should contain the location of the file `libOpenCL.so` of the OCL ICD implementation.

**OCL\_ICD\_VENDORS** When using SOCL with OpenCL ICD (<https://forge.imag.fr/projects/ocl-icd/>), this variable may be used to point to the directory where ICD files are installed. The default directory is `/etc/OpenCL/vendors`. StarPU installs ICD files in the directory `$prefix/share/starpu/opencl/vendors`.

**STARPU\_COMM\_STATS** Communication statistics for starpumpi ([MPI Support](#)) will be enabled when the environment variable **STARPU\_COMM\_STATS** is defined to an value other than 0.

**STARPU\_MPI\_CACHE** Communication cache for starpumpi ([MPI Support](#)) will be disabled when the environment variable **STARPU\_MPI\_CACHE** is set to 0. It is enabled by default or for any other values of the variable **STARPU\_MPI\_CACHE**.

## 14.4 Miscellaneous And Debug

**STARPU\_HOME** This specifies the main directory in which StarPU stores its configuration files. The default is `$HOME` on Unix environments, and `$USERPROFILE` on Windows environments.

**STARPU\_HOSTNAME** When set, force the hostname to be used when dealing performance model files. Models are indexed by machine name. When running for example on a homogenous cluster, it is possible to share the models between machines by setting `export STARPU_HOSTNAME=some_global_name`.

**STARPU\_OPENCL\_PROGRAM\_DIR** This specifies the directory where the OpenCL codelet source files are located. The function `starpu_opencl_load_program_source()` looks for the codelet in the current directory, in the directory specified by the environment variable **STARPU\_OPENCL\_PROGRAM\_DIR**, in the directory `share/starpu/opencl` of the installation directory of StarPU, and finally in the source directory of StarPU.

**STARPU\_SILENT** This variable allows to disable verbose mode at runtime when StarPU has been configured with the option `--enable-verbose`. It also disables the display of StarPU information and warning messages.

**STARPU\_LOGFILENAME** This variable specifies in which file the debugging output should be saved to.

**STARPU\_FXT\_PREFIX** This variable specifies in which directory to save the trace generated if FxT is enabled. It needs to have a trailing '/' character.

**STARPU\_LIMIT\_CUDA\_devid\_MEM** This variable specifies the maximum number of megabytes that should be available to the application on the CUDA device with the identifier `devid`. This variable is intended to be used for experimental purposes as it emulates devices that have a limited amount of memory. When defined, the variable overwrites the value of the variable [STARPU\\_LIMIT\\_CUDA\\_MEM](#).

**STARPU\_LIMIT\_CUDA\_MEM** This variable specifies the maximum number of megabytes that should be available to the application on each CUDA devices. This variable is intended to be used for experimental purposes as it emulates devices that have a limited amount of memory.

**STARPU\_LIMIT\_OPENCL\_devid\_MEM** This variable specifies the maximum number of megabytes that should be available to the application on the OpenCL device with the identifier `devid`. This variable is intended to be used for experimental purposes as it emulates devices that have a limited amount of memory. When defined, the variable overwrites the value of the variable [STARPU\\_LIMIT\\_OPENCL\\_MEM](#).

**STARPU\_LIMIT\_OPENCL\_MEM** This variable specifies the maximum number of megabytes that should be available to the application on each OpenCL devices. This variable is intended to be used for experimental purposes as it emulates devices that have a limited amount of memory.

**STARPU\_LIMIT\_CPU\_MEM** This variable specifies the maximum number of megabytes that should be available to the application on each CPU device. This variable is intended to be used for experimental purposes as it emulates devices that have a limited amount of memory.

**STARPU\_GENERATE\_TRACE** When set to 1, this variable indicates that StarPU should automatically generate a Paje trace when [starpu\\_shutdown\(\)](#) is called.

**STARPU\_MEMORY\_STATS** When set to 0, disable the display of memory statistics on data which have not been unregistered at the end of the execution ([Memory Feedback](#)).

**STARPU\_BUS\_STATS** When defined, statistics about data transfers will be displayed when calling [starpu\\_shutdown\(\)](#) ([Profiling](#)).

**STARPU\_WORKER\_STATS** When defined, statistics about the workers will be displayed when calling [starpu\\_shutdown\(\)](#) ([Profiling](#)). When combined with the environment variable **STARPU\_PROFILING**, it displays the power consumption ([Power-based Scheduling](#)).

**STARPU\_STATS** When set to 0, data statistics will not be displayed at the end of the execution of an application ([Data Statistics](#)).

**STARPU\_WATCHDOG\_TIMEOUT** When set to a value other than 0, allows to make StarPU print an error message whenever StarPU does not terminate any task for 10ms. Should be used in combination with **STARPU\_WATCHDOG\_CRASH** (see [Detection Stuck Conditions](#)).

**STARPU\_WATCHDOG\_CRASH** When set to a value other than 0, it triggers a crash when the watch dog is reached, thus allowing to catch the situation in gdb, etc (see [Detection Stuck Conditions](#))

## 14.5 Configuring The Hypervisor

**SC\_HYPERVISOR\_POLICY** Choose between the different resizing policies proposed by StarPU for the hypervisor: idle, app\_driven, feft\_lp, teft\_lp; ispeed\_lp, throughput\_lp etc.

Use `SC_HYPERVISOR_POLICY=help` to get the list of available policies for the hypervisor

**SC\_HYPERVISOR\_TRIGGER\_RESIZE** Choose how should the hypervisor be triggered: `speed` if the resizing algorithm should be called whenever the speed of the context does not correspond to an optimal precomputed value, `idle` if the resizing algorithm should be called whenever the workers are idle for a period longer than the value indicated when configuring the hypervisor.

**SC\_HYPERVISOR\_START\_RESIZE** Indicate the moment when the resizing should be available. The value correspond to the percentage of the total time of execution of the application. The default value is the resizing frame.

**SC\_HYPERVISOR\_MAX\_SPEED\_GAP** Indicate the ratio of speed difference between contexts that should trigger the hypervisor. This situation may occur only when a theoretical speed could not be computed and the hypervisor has no value to compare the speed to. Otherwise the resizing of a context is not influenced by the speed of the other contexts, but only by the the value that a context should have.

**SC\_HYPERVISOR\_STOP\_PRINT** By default the values of the speed of the workers is printed during the execution of the application. If the value 1 is given to this environment variable this printing is not done.

**SC\_HYPERVISOR.LAZY.RESIZE** By default the hypervisor resizes the contexts in a lazy way, that is workers are firstly added to a new context before removing them from the previous one. Once this workers are clearly taken into account into the new context (a task was poped there) we remove them from the previous one. However if the application would like that the change in the distribution of workers should change right away this variable should be set to 0

**SC\_HYPERVISOR\_SAMPLE\_CRITERIA** By default the hypervisor uses a sample of flops when computing the speed of the contexts and of the workers. If this variable is set to time the hypervisor uses a sample of time (10% of an approximation of the total execution time of the application)



# Chapter 15

# Compilation Configuration

The behavior of the StarPU library and tools may be tuned thanks to the following configure options.

## 15.1 Common Configuration

- enable-debug** Enable debugging messages.
- enable-spinlock-check** Enable checking that spinlocks are taken and released properly.
- enable-fast** Disable assertion checks, which saves computation time.
- enable-verbose** Increase the verbosity of the debugging messages. This can be disabled at runtime by setting the environment variable `STARPU_SILENT` to any value.

```
$ STARPU_SILENT=1 ./vector_scal
```

- enable-coverage** Enable flags for the coverage tool `gcov`.
- enable-quick-check** Specify tests and examples should be run on a smaller data set, i.e allowing a faster execution time
- enable-long-check** Enable some exhaustive checks which take a really long time.
- with-hwloc** Specify `hwloc` should be used by StarPU. `hwloc` should be found by the means of the tool `pkg-config`.
- with-hwloc=prefix** Specify `hwloc` should be used by StarPU. `hwloc` should be found in the directory specified by `prefix`
- without-hwloc** Specify `hwloc` should not be used by StarPU.
- disable-build-doc** Disable the creation of the documentation. This should be done on a machine which does not have the tools `doxygen` and `latex` (plus the packages `latex-xcolor` and `texlive-latex-extra`).

Additionally, the script `configure` recognize many variables, which can be listed by typing `./configure --help`. For example, `./configure NVCCFLAGS="-arch sm_13"` adds a flag for the compilation of CUDA kernels.

## 15.2 Configuring Workers

- enable-maxcpus=count** Use at most `count` CPU cores. This information is then available as the macro `:STARPU_MAXCPUS`.
- disable-cpu** Disable the use of CPUs of the machine. Only GPUs etc. will be used.
- enable-maxcudadev=count** Use at most `count` CUDA devices. This information is then available as the macro `STARPU_MAXCUDADEVS`.

- disable-cuda** Disable the use of CUDA, even if a valid CUDA installation was detected.
- with-cuda-dir=prefix** Search for CUDA under `prefix`, which should notably contain the file `include/cuda.h`.
- with-cuda-include-dir=dir** Search for CUDA headers under `dir`, which should notably contain the file `cuda.h`. This defaults to `/include` appended to the value given to [--with-cuda-dir](#).
- with-cuda-lib-dir=dir** Search for CUDA libraries under `dir`, which should notably contain the CUDA shared libraries—e.g., `libcuda.so`. This defaults to `/lib` appended to the value given to [--with-cuda-dir](#).
- disable-cuda-memcpy-peer** Explicitly disable peer transfers when using CUDA 4.0.
- enable-maxopencldevs=count** Use at most `count` OpenCL devices. This information is then available as the macro [STARPU\\_MAXOPENCLDEVS](#).
- disable-opencl** Disable the use of OpenCL, even if the SDK is detected.
- with-opencl-dir=prefix** Search for an OpenCL implementation under `prefix`, which should notably contain `include/CL/cl.h` (or `include/OpenCL/cl.h` on Mac OS).
- with-opencl-include-dir=dir** Search for OpenCL headers under `dir`, which should notably contain `CL/cl.h` (or `OpenCL/cl.h` on Mac OS). This defaults to `/include` appended to the value given to [--with-opencl-dir](#).
- with-opencl-lib-dir=dir** Search for an OpenCL library under `dir`, which should notably contain the OpenCL shared libraries—e.g. `libOpenCL.so`. This defaults to `/lib` appended to the value given to [--with-opencl-dir](#).
- enable-opencl-simulator** Enable considering the provided OpenCL implementation as a simulator, i.e. use the kernel duration returned by OpenCL profiling information as wallclock time instead of the actual measured real time. This requires simgrid support.
- enable-maximplementations=count** Allow for at most `count` codelet implementations for the same target device. This information is then available as the macro `::STARPU_MAXIMPLEMENTATIONS` macro.
- enable-max-sched-ctxs=count** Allow for at most `count` scheduling contexts. This information is then available as the macro `::STARPU_NMAX_SCHED_CTXS`.
- disable-asynchronous-copy** Disable asynchronous copies between CPU and GPU devices. The AMD implementation of OpenCL is known to fail when copying data asynchronously. When using this implementation, it is therefore necessary to disable asynchronous data transfers.
- disable-asynchronous-cuda-copy** Disable asynchronous copies between CPU and CUDA devices.
- disable-asynchronous-opencl-copy** Disable asynchronous copies between CPU and OpenCL devices. The AMD implementation of OpenCL is known to fail when copying data asynchronously. When using this implementation, it is therefore necessary to disable asynchronous data transfers.

## 15.3 Extension Configuration

- disable-socl** Disable the SOCL extension ([SOCL OpenCL Extensions](#)). By default, it is enabled when an OpenCL implementation is found.
- disable-starpu-top** Disable the StarPU-Top interface ([StarPU-Top Interface](#)). By default, it is enabled when the required dependencies are found.
- disable-gcc-extensions** Disable the GCC plug-in ([C Extensions](#)). By default, it is enabled when the GCC compiler provides a plug-in support.
- with-mpicc=path** Use the compiler `mpicc` at `path`, for StarPU-MPI. ([MPI Support](#)).
- enable-mpi-progression-hook** Enable the activity polling method for StarPU-MPI.

## 15.4 Advanced Configuration

- enable-perf-debug** Enable performance debugging through gprof.
- enable-model-debug** Enable performance model debugging.
- enable-stats** (see `../../src/datawizard/datastats.c`) Enable gathering of various data statistics ([Data Statistics](#)).
- enable-maxbuffers** Define the maximum number of buffers that tasks will be able to take as parameters, then available as the macro `STARPU_NMAXBUFS`.
- enable-allocation-cache** Enable the use of a data allocation cache to avoid the cost of it with CUDA. Still experimental.
- enable-opengl-render** Enable the use of OpenGL for the rendering of some examples.



# Chapter 16

## Module Index

### 16.1 Modules

Here is a list of all modules:

Codelet And Tasks . . . . .	121
CUDA Extensions . . . . .	145
Data Interfaces . . . . .	95
Data Management . . . . .	91
Data Partition . . . . .	113
Expert Mode . . . . .	165
Explicit Dependencies . . . . .	133
FFT Support . . . . .	154
FxT Support . . . . .	152
Implicit Data Dependencies . . . . .	135
Initialization and Termination . . . . .	75
Insert_Task . . . . .	131
Theoretical Lower Bound on Execution Time . . . . .	144
Miscellaneous Helpers . . . . .	152
MPI Support . . . . .	155
Multiformat Data Interface . . . . .	118
OpenCL Extensions . . . . .	146
Parallel Tasks . . . . .	164
Performance Model . . . . .	136
Profiling . . . . .	141
Running Drivers . . . . .	165
Scheduling Contexts . . . . .	170
Scheduling Policy . . . . .	175
Standard Memory Library . . . . .	79
Task Bundles . . . . .	160
Task Lists . . . . .	161
Threads . . . . .	82
Toolbox . . . . .	80
StarPU-Top Interface . . . . .	166
Versioning . . . . .	75
Workers' Properties . . . . .	88
Scheduling Context Hypervisor - Building a new resizing policy . . . . .	181
Scheduling Context Hypervisor - Regular usage . . . . .	178



# Chapter 17

## Module Documentation a.k.a StarPU's API

### 17.1 Versioning

#### Macros

- `#define STARPU_MAJOR_VERSION`
- `#define STARPU_MINOR_VERSION`
- `#define STARPU_RELEASE_VERSION`

#### Functions

- `void starpu_get_version (int *major, int *minor, int *release)`

##### 17.1.1 Detailed Description

##### 17.1.2 Macro Definition Documentation

###### `#define STARPU_MAJOR_VERSION`

Define the major version of StarPU. This is the version used when compiling the application.

###### `#define STARPU_MINOR_VERSION`

Define the minor version of StarPU. This is the version used when compiling the application.

###### `#define STARPU_RELEASE_VERSION`

Define the release version of StarPU. This is the version used when compiling the application.

##### 17.1.3 Function Documentation

###### `void starpu_get_version ( int * major, int * minor, int * release )`

Return as 3 integers the version of StarPU used when running the application.

### 17.2 Initialization and Termination

#### Data Structures

- struct `starpu_driver`
- union `starpu_driver.id`
- struct `starpu_vector_interface`
- struct `starpu_conf`

## Functions

- int `starpu_init` (struct `starpu_conf` \*`conf`) `STARPU_WARN_UNUSED_RESULT`
- int `starpu_conf_init` (struct `starpu_conf` \*`conf`)
- void `starpu_shutdown` (void)
- int `starpu_asynchronous_copy_disabled` (void)
- int `starpu_asynchronous_cuda_copy_disabled` (void)
- int `starpu_asynchronous_opencl_copy_disabled` (void)
- void `starpu_topology_print` (FILE \*`f`)

### 17.2.1 Detailed Description

### 17.2.2 Data Structure Documentation

#### `struct starpu_driver`

structure for a driver

##### Data Fields

enum <code>starpu_worker_archtype</code>	<code>type</code>	The type of the driver. Only <code>STARPU_CPU_WORKER</code> , <code>STARPU_CUDA_WORKER</code> and <code>STARPU_OPENCL_WORKER</code> are currently supported.
union <code>starpu_driver</code>	<code>id</code>	The identifier of the driver.

#### `union starpu_driver.id`

##### Data Fields

unsigned	<code>cpu_id</code>	
unsigned	<code>cuda_id</code>	
<code>cl_device_id</code>	<code>opencl_id</code>	

#### `struct starpu_vector_interface`

Vector interface

vector interface for contiguous (non-strided) buffers

##### Data Fields

<code>uintptr_t</code>	<code>ptr</code>	local pointer of the vector
<code>uintptr_t</code>	<code>dev.handle</code>	device handle of the vector.
<code>size_t</code>	<code>offset</code>	offset in the vector
<code>uint32_t</code>	<code>nx</code>	number of elements on the x-axis of the vector
<code>size_t</code>	<code>elemsize</code>	size of the elements of the vector

#### `struct starpu_conf`

This structure is passed to the `starpu_init()` function in order to configure StarPU. It has to be initialized with `starpu_conf_init()`. When the default value is used, StarPU automatically selects the number of processing units and takes the default scheduling policy. The environment variables overwrite the equivalent parameters.

##### Data Fields

<code>int</code>	<code>magic</code>	Will be initialized by <code>starpu_conf_init()</code> . Should not be set by hand.
<code>const char *</code>	<code>sched_policy_name</code>	This is the name of the scheduling policy. This can also be specified with the environment variable <code>STARPU_SCHED</code> . (default = NULL).

struct <code>starpu_sched_policy</code> *	<code>sched_policy</code>	This is the definition of the scheduling policy. This field is ignored if <code>starpu_conf::sched_policy_name</code> is set. (default = NULL)
int	<code>ncpus</code>	This is the number of CPU cores that StarPU can use. This can also be specified with the environment variable <code>STARPU_NCPU</code> . (default = -1)
int	<code>ncuda</code>	This is the number of CUDA devices that StarPU can use. This can also be specified with the environment variable <code>STARPU_NCUD</code> A. (default = -1)
int	<code>nopencl</code>	This is the number of OpenCL devices that StarPU can use. This can also be specified with the environment variable <code>STARPU_NOOPENCL</code> . (default = -1)
unsigned	<code>use_explicit_workers_bindid</code>	If this flag is set, the <code>starpu_conf::workers_bindid</code> array indicates where the different workers are bound, otherwise StarPU automatically selects where to bind the different workers. This can also be specified with the environment variable <code>STARPU_WORKERS_CPUID</code> . (default = 0)
unsigned	<code>workers_bindid[STARPU_NMAXWORKERS]</code>	If the <code>starpu_conf::use_explicit_workers_bindid</code> flag is set, this array indicates where to bind the different workers. The i-th entry of the <code>starpu_conf::workers_bindid</code> indicates the logical identifier of the processor which should execute the i-th worker. Note that the logical ordering of the CPUs is either determined by the OS, or provided by the hwloc library in case it is available.
unsigned	<code>use_explicit_workers_cuda_gpid</code>	If this flag is set, the CUDA workers will be attached to the CUDA devices specified in the <code>starpu_conf::workers_cuda_gpid</code> array. Otherwise, StarPU affects the CUDA devices in a round-robin fashion. This can also be specified with the environment variable <code>STARPU_WORKERS_CUDAID</code> . (default = 0)
unsigned	<code>workers_cuda_gpid[STARPU_NMAXWORKERS]</code>	If the <code>starpu_conf::use_explicit_workers_cuda_gpid</code> flag is set, this array contains the logical identifiers of the CUDA devices (as used by <code>cudaGetDevice()</code> ).
unsigned	<code>use_explicit_workers_opencl_gpid</code>	If this flag is set, the OpenCL workers will be attached to the OpenCL devices specified in the <code>starpu_conf::workers_opencl_gpid</code> array. Otherwise, StarPU affects the OpenCL devices in a round-robin fashion. This can also be specified with the environment variable <code>STARPU_WORKERS_OPENCLID</code> . (default = 0)
unsigned	<code>workers_opencl_gpid[STARPU_NMAXWORKERS]</code>	If the <code>starpu_conf::use_explicit_workers_opencl_gpid</code> flag is set, this array contains the logical identifiers of the OpenCL devices to be used.
int	<code>bus_calibrate</code>	If this flag is set, StarPU will recalibrate the bus. If this value is equal to -1, the default value is used. This can also be specified with the environment variable <code>STARPU_BUS_CALIBRATE</code> . (default = 0)
int	<code>calibrate</code>	If this flag is set, StarPU will calibrate the performance models when executing tasks. If this value is equal to -1, the default value is used. If the value is equal to 1, it will force continuing calibration. If the value is equal to 2, the existing performance models will be overwritten. This can also be specified with the environment variable <code>STARPU_CALIBRATE</code> . (default = 0)

	int	single_-combined_-worker	By default, StarPU executes parallel tasks concurrently. Some parallel libraries (e.g. most OpenMP implementations) however do not support concurrent calls to parallel code. In such case, setting this flag makes StarPU only start one parallel task at a time (but other CPU and GPU tasks are not affected and can be run concurrently). The parallel task scheduler will however still try varying combined worker sizes to look for the most efficient ones. This can also be specified with the environment variable <a href="#">STARPU_SINGLE_COMBINED_WORKER</a> . (default = 0)
	int	disable_-asynchronous_-copy	This flag should be set to 1 to disable asynchronous copies between CPUs and all accelerators. This can also be specified with the environment variable <a href="#">STARPU_DISABLE_ASYNCHRONOUS_COPY</a> . The AMD implementation of OpenCL is known to fail when copying data asynchronously. When using this implementation, it is therefore necessary to disable asynchronous data transfers. This can also be specified at compilation time by giving to the configure script the option <a href="#">--disable-asynchronous-copy</a> . (default = 0)
	int	disable_-asynchronous_-cuda_copy	This flag should be set to 1 to disable asynchronous copies between CPUs and CUDA accelerators. This can also be specified with the environment variable <a href="#">STARPU_DISABLE_ASYNCHRONOUS_CUDA_COPY</a> . This can also be specified at compilation time by giving to the configure script the option <a href="#">--disable-asynchronous-cuda-copy</a> . (default = 0)
	int	disable_-asynchronous_-opencl_copy	This flag should be set to 1 to disable asynchronous copies between CPUs and OpenCL accelerators. This can also be specified with the environment variable <a href="#">STARPU_DISABLE_ASYNCHRONOUS_OPENCL_COPY</a> . The AMD implementation of OpenCL is known to fail when copying data asynchronously. When using this implementation, it is therefore necessary to disable asynchronous data transfers. This can also be specified at compilation time by giving to the configure script the option <a href="#">--disable-asynchronous-opencl-copy</a> . (default = 0)
	unsigned *	cuda_opengl_-interoperability	Enable CUDA/OpenGL interoperation on these CUDA devices. This can be set to an array of CUDA device identifiers for which cudaGLSetGLDevice() should be called instead of cudaSetDevice(). Its size is specified by the starpu.conf::n_cuda_opengl_interoperability field below (default = NULL)
	unsigned	n_cuda_opengl_-interoperability	
	struct <a href="#">starpu_driver</a> *	not_launched_drivers	Array of drivers that should not be launched by StarPU. The application will run in one of its own threads. (default = NULL)
	unsigned	n_not_launched_drivers	The number of StarPU drivers that should not be launched by StarPU. (default = 0)
	unsigned	trace_buffer_size	Specifies the buffer size used for FxT tracing. Starting from FxT version 0.2.12, the buffer will automatically be flushed when it fills in, but it may still be interesting to specify a bigger value to avoid any flushing (which would disturb the trace).

### 17.2.3 Function Documentation

`int starpu_init( struct starpu_conf * conf )`

This is StarPU initialization method, which must be called prior to any other StarPU call. It is possible to specify StarPU's configuration (e.g. scheduling policy, number of cores, ...) by passing a non-null argument. Default configuration is used if the passed argument is NULL. Upon successful completion, this function returns 0. Otherwise, -ENODEV indicates that no worker was available (so that StarPU was not initialized).

```
int starpu_conf_init ( struct starpu_conf * conf )
```

This function initializes the conf structure passed as argument with the default values. In case some configuration parameters are already specified through environment variables, `starpu_conf_init()` initializes the fields of the structure according to the environment variables. For instance if `STARPU_CALIBRATE` is set, its value is put in the field `starpu_conf::calibrate` of the structure passed as argument. Upon successful completion, this function returns 0. Otherwise, -EINVAL indicates that the argument was NULL.

```
void starpu_shutdown ( void )
```

This is StarPU termination method. It must be called at the end of the application: statistics and other post-mortem debugging information are not guaranteed to be available until this method has been called.

```
int starpu_asynchronous_copy_disabled ( void )
```

Return 1 if asynchronous data transfers between CPU and accelerators are disabled.

```
int starpu_asynchronous_cuda_copy_disabled ( void )
```

Return 1 if asynchronous data transfers between CPU and CUDA accelerators are disabled.

```
int starpu_asynchronous_opencl_copy_disabled ( void )
```

Return 1 if asynchronous data transfers between CPU and OpenCL accelerators are disabled.

```
void starpu_topology_print ( FILE * f )
```

Prints a description of the topology on f.

## 17.3 Standard Memory Library

### Macros

- `#define starpu_data_malloc_pinned_if_possible`
- `#define starpu_data_free_pinned_if_possible`
- `#define STARPU_MALLOC_PINNED`
- `#define STARPU_MALLOC_COUNT`

### Functions

- `int starpu_malloc_flags (void **A, size_t dim, int flags)`
- `void starpu_malloc_set_align (size_t align)`
- `int starpu_malloc (void **A, size_t dim)`
- `int starpu_free (void *A)`
- `int starpu_free_flags (void *A, size_t dim, int flags)`
- `starpu_size_t starpu_memory_get_available (unsigned node)`

#### 17.3.1 Detailed Description

#### 17.3.2 Macro Definition Documentation

`#define starpu_data_malloc_pinned_if_possible`

**Deprecated** Equivalent to `starpu_malloc()`. This macro is provided to avoid breaking old codes.

`#define starpu_data_free_pinned_if_possible`

**Deprecated** Equivalent to `starpu_free()`. This macro is provided to avoid breaking old codes.

**#define STARPU\_MALLOC\_PINNED**

Value passed to the function [starpu\\_malloc\\_flags\(\)](#) to indicate the memory allocation should be pinned.

**#define STARPU\_MALLOC\_COUNT**

Value passed to the function [starpu\\_malloc\\_flags\(\)](#) to indicate the memory allocation should be in the limit defined by the environment variables `STARPU_LIMIT_CUDA_devid_MEM`, `STARPU_LIMIT_CUDA_MEM`, `STARPU_LIMIT_OPENCL_devid_MEM`, `STARPU_LIMIT_OPENCL_MEM` and `STARPU_LIMIT_CPU_MEM` (see Section [How to limit memory per node](#)). If no memory is available, it tries to reclaim memory from StarPU. Memory allocated this way needs to be freed by calling the function [starpu\\_free\\_flags\(\)](#) with the same flag.

### 17.3.3 Function Documentation

**int starpu\_malloc\_flags ( void \*\* A, size\_t dim, int flags )**

Performs a memory allocation based on the constraints defined by the given flag.

**void starpu\_malloc\_set\_align ( size\_t align )**

This function sets an alignment constraints for [starpu\\_malloc\(\)](#) allocations. `align` must be a power of two. This is for instance called automatically by the OpenCL driver to specify its own alignment constraints.

**int starpu\_malloc ( void \*\* A, size\_t dim )**

This function allocates data of the given size in main memory. It will also try to pin it in CUDA or OpenCL, so that data transfers from this buffer can be asynchronous, and thus permit data transfer and computation overlapping. The allocated buffer must be freed thanks to the [starpu\\_free\(\)](#) function.

**int starpu\_free ( void \* A )**

This function frees memory which has previously been allocated with [starpu\\_malloc\(\)](#).

**int starpu\_free\_flags ( void \* A, size\_t dim, int flags )**

This function frees memory by specifying its size. The given flags should be consistent with the ones given to [starpu\\_malloc\\_flags\(\)](#) when allocating the memory.

**ssize\_t starpu\_memory\_get\_available ( unsigned node )**

If a memory limit is defined on the given node (see Section [How to limit memory per node](#)), return the amount of available memory on the node. Otherwise return -1.

## 17.4 Toolbox

The following macros allow to make GCC extensions portable, and to have a code which can be compiled with any C compiler.

### Macros

- `#define STARPU_GNUC_PREREQ(maj, min)`
- `#define STARPU_UNLIKELY(expr)`
- `#define STARPU_LIKELY(expr)`
- `#define STARPU_ATTRIBUTE_UNUSED`
- `#define STARPU_ATTRIBUTE_INTERNAL`
- `#define STARPU_ATTRIBUTE_MALLOC`
- `#define STARPU_ATTRIBUTE_WARN_UNUSED_RESULT`
- `#define STARPU_ATTRIBUTE_PURE`
- `#define STARPU_ATTRIBUTE_ALIGNED(size)`

- `#define STARPU_WARN_UNUSED_RESULT`
- `#define STARPU_POISON_PTR`
- `#define STARPU_MIN(a, b)`
- `#define STARPU_MAX(a, b)`
- `#define STARPU_ASSERT(x)`
- `#define STARPU_ASSERT_MSG(x, msg,...)`
- `#define STARPU_ABORT()`
- `#define STARPU_ABORT_MSG(msg,...)`
- `#define STARPU_CHECK_RETURN_VALUE(err, message,...)`
- `#define STARPU_CHECK_RETURN_VALUE_IS(err, value, message,...)`
- `#define STARPU_RMB()`
- `#define STARPU_WMB()`

## Functions

- static \_\_starpu\_inline int `starpu_get_env_number` (const char \*str)

### 17.4.1 Detailed Description

The following macros allow to make GCC extensions portable, and to have a code which can be compiled with any C compiler.

### 17.4.2 Macro Definition Documentation

**#define STARPU\_GNUC\_PREREQ( maj, min )**

Return true (non-zero) if GCC version MAJ.MIN or later is being used (macro taken from glibc.)

**#define STARPU\_UNLIKELY( expr )**

When building with a GNU C Compiler, this macro allows programmers to mark an expression as unlikely.

**#define STARPU\_LIKELY( expr )**

When building with a GNU C Compiler, this macro allows programmers to mark an expression as likely.

**#define STARPU\_ATTRIBUTE\_UNUSED**

When building with a GNU C Compiler, this macro is defined to `__attribute__((unused))`

**#define STARPU\_ATTRIBUTE\_INTERNAL**

When building with a GNU C Compiler, this macro is defined to `__attribute__((visibility ("internal")))`

**#define STARPU\_ATTRIBUTE\_MALLOC**

When building with a GNU C Compiler, this macro is defined to `__attribute__((malloc))`

**#define STARPU\_ATTRIBUTE\_WARN\_UNUSED\_RESULT**

When building with a GNU C Compiler, this macro is defined to `__attribute__((warn_unused_result))`

**#define STARPU\_ATTRIBUTE\_PURE**

When building with a GNU C Compiler, this macro is defined to `__attribute__((pure))`

**#define STARPU\_ATTRIBUTE\_ALIGNED( size )**

When building with a GNU C Compiler, this macro is defined to `__attribute__((aligned(size)))`

**#define STARPU\_WARN\_UNUSED\_RESULT**

When building with a GNU C Compiler, this macro is defined to `_attribute__((warn_unused_result))`

**#define STARPU\_POISON\_PTR**

This macro defines a value which can be used to mark pointers as invalid values.

**#define STARPU\_MIN( a, b )**

This macro returns the min of the two parameters.

**#define STARPU\_MAX( a, b )**

This macro returns the max of the two parameters.

**#define STARPU\_ASSERT( x )**

Unless StarPU has been configured with the option `--enable-fast`, this macro will abort if the expression is false.

**#define STARPU\_ASSERT\_MSG( x, msg, ... )**

Unless StarPU has been configured with the option `--enable-fast`, this macro will abort if the expression is false. The given message will be displayed.

**#define STARPU\_ABORT( )**

This macro aborts the program.

**#define STARPU\_ABORT\_MSG( msg, ... )**

This macro aborts the program, and displays the given message.

**#define STARPU\_CHECK\_RETURN\_VALUE( err, message, ... )**

If `err` has a value which is not 0, the given message is displayed before aborting.

**#define STARPU\_CHECK\_RETURN\_VALUE\_IS( err, value, message, ... )**

If `err` has a value which is not `value`, the given message is displayed before aborting.

**#define STARPU\_RMB( )**

This macro can be used to do a synchronization.

**#define STARPU\_WMB( )**

This macro can be used to do a synchronization.

### 17.4.3 Function Documentation

**int starpu\_get\_env\_number ( const char \* str ) [static]**

If `str` is the name of a existing environment variable which is defined to an integer, the function returns the value of the integer. It returns 0 otherwise.

## 17.5 Threads

This section describes the thread facilities provided by StarPU. The thread function are either implemented on top of the pthread library or the Simgrid library when the simulated performance mode is enabled ([Simulated Performance](#)).

## Macros

- #define STARPU\_PTHREAD\_CREATE\_ON(name, thread, attr, routine, arg, where)
- #define STARPU\_PTHREAD\_CREATE(thread, attr, routine, arg)
- #define STARPU\_PTHREAD\_MUTEX\_INIT(mutex, attr)
- #define STARPU\_PTHREAD\_MUTEX\_DESTROY(mutex)
- #define STARPU\_PTHREAD\_MUTEX\_LOCK(mutex)
- #define STARPU\_PTHREAD\_MUTEX\_UNLOCK(mutex)
- #define STARPU\_PTHREAD\_KEY\_CREATE(key, destr)
- #define STARPU\_PTHREAD\_KEY\_DELETE(key)
- #define STARPU\_PTHREAD\_SETSPECIFIC(key, ptr)
- #define STARPU\_PTHREAD\_GETSPECIFIC(key)
- #define STARPU\_PTHREAD\_RWLOCK\_INIT(rwlock, attr)
- #define STARPU\_PTHREAD\_RWLOCK\_RDLOCK(rwlock)
- #define STARPU\_PTHREAD\_RWLOCK\_WRLOCK(rwlock)
- #define STARPU\_PTHREAD\_RWLOCK\_UNLOCK(rwlock)
- #define STARPU\_PTHREAD\_RWLOCK\_DESTROY(rwlock)
- #define STARPU\_PTHREAD\_COND\_INIT(cond, attr)
- #define STARPU\_PTHREAD\_COND\_DESTROY(cond)
- #define STARPU\_PTHREAD\_COND\_SIGNAL(cond)
- #define STARPU\_PTHREAD\_COND\_BROADCAST(cond)
- #define STARPU\_PTHREAD\_COND\_WAIT(cond, mutex)
- #define STARPU\_PTHREAD\_BARRIER\_INIT(barrier, attr, count)
- #define STARPU\_PTHREAD\_BARRIER\_DESTROY(barrier)
- #define STARPU\_PTHREAD\_BARRIER\_WAIT(barrier)
- #define STARPU\_PTHREAD\_MUTEX\_INITIALIZER
- #define STARPU\_PTHREAD\_COND\_INITIALIZER

## Functions

- int **starpu\_pthread\_create\_on** (char \*name, starpu\_pthread\_t \*thread, const starpu\_pthread\_attr\_t \*attr, void \*(\*start\_routine)(void \*), void \*arg, int where)
- int **starpu\_pthread\_create** (starpu\_pthread\_t \*thread, const starpu\_pthread\_attr\_t \*attr, void \*(\*start\_routine)(void \*), void \*arg)
- int **starpu\_pthread\_join** (starpu\_pthread\_t thread, void \*\*retval)
- int **starpu\_pthread\_attr\_init** (starpu\_pthread\_attr\_t \*attr)
- int **starpu\_pthread\_attr\_destroy** (starpu\_pthread\_attr\_t \*attr)
- int **starpu\_pthread\_attr\_setdetachstate** (starpu\_pthread\_attr\_t \*attr, int detachstate)
- int **starpu\_pthread\_mutex\_init** (starpu\_pthread\_mutex\_t \*mutex, const starpu\_pthread\_mutexattr\_t \*mutexattr)
- int **starpu\_pthread\_mutex\_destroy** (starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_mutex\_lock** (starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_mutex\_unlock** (starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_mutex\_trylock** (starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_key\_create** (starpu\_pthread\_key\_t \*key, void(\*destr\_function)(void \*))
- int **starpu\_pthread\_key\_delete** (starpu\_pthread\_key\_t key)
- int **starpu\_pthread\_setspecific** (starpu\_pthread\_key\_t key, const void \*pointer)
- void \* **starpu\_pthread\_getspecific** (starpu\_pthread\_key\_t key)
- int **starpu\_pthread\_cond\_init** (starpu\_pthread\_cond\_t \*cond, starpu\_pthread\_condattr\_t \*cond\_attr)
- int **starpu\_pthread\_cond\_signal** (starpu\_pthread\_cond\_t \*cond)
- int **starpu\_pthread\_cond\_broadcast** (starpu\_pthread\_cond\_t \*cond)
- int **starpu\_pthread\_cond\_wait** (starpu\_pthread\_cond\_t \*cond, starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_cond\_timedwait** (starpu\_pthread\_cond\_t \*cond, starpu\_pthread\_mutex\_t \*mutex, const struct timespec \*abstime)
- int **starpu\_pthread\_cond\_destroy** (starpu\_pthread\_cond\_t \*cond)
- int **starpu\_pthread\_rwlock\_init** (starpu\_pthread\_rwlock\_t \*rwlock, const starpu\_pthread\_rwlockattr\_t \*attr)
- int **starpu\_pthread\_rwlock\_destroy** (starpu\_pthread\_rwlock\_t \*rwlock)
- int **starpu\_pthread\_rwlock\_rdlock** (starpu\_pthread\_rwlock\_t \*rwlock)
- int **starpu\_pthread\_rwlock\_wrlock** (starpu\_pthread\_rwlock\_t \*rwlock)
- int **starpu\_pthread\_rwlock\_unlock** (starpu\_pthread\_rwlock\_t \*rwlock)

### 17.5.1 Detailed Description

This section describes the thread facilities provided by StarPU. The thread function are either implemented on top of the pthread library or the Simgrid library when the simulated performance mode is enabled ([Simulated Performance](#)).

### 17.5.2 Macro Definition Documentation

```
#define STARPU_PTHREAD_CREATE_ON( name, thread, attr, routine, arg, where )
```

This macro calls the function [starpu\\_pthread\\_create\\_on\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_CREATE( thread, attr, routine, arg )
```

This macro calls the function [starpu\\_pthread\\_create\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_MUTEX_INIT( mutex, attr )
```

This macro calls the function [starpu\\_pthread\\_mutex\\_init\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_MUTEX_DESTROY( mutex )
```

This macro calls the function [starpu\\_pthread\\_mutex\\_destroy\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_MUTEX_LOCK( mutex )
```

This macro calls the function [starpu\\_pthread\\_mutex\\_lock\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_MUTEX_UNLOCK( mutex )
```

This macro calls the function [starpu\\_pthread\\_mutex\\_unlock\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_KEY_CREATE( key, destr )
```

This macro calls the function [starpu\\_pthread\\_key\\_create\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_KEY_DELETE( key )
```

This macro calls the function [starpu\\_pthread\\_key\\_delete\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_SETSPECIFIC( key, ptr )
```

This macro calls the function [starpu\\_pthread\\_setspecific\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_GETSPECIFIC( key )
```

This macro calls the function [starpu\\_pthread\\_getspecific\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_RWLOCK_INIT( rwlock, attr )
```

This macro calls the function [starpu\\_pthread\\_rwlock\\_init\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_RWLOCK_RDLOCK( rwlock )
```

This macro calls the function [starpu\\_pthread\\_rwlock\\_rdlock\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_RWLOCK_WRLock( rwlock )
```

This macro calls the function [starpu\\_pthread\\_rwlock\\_wrlock\(\)](#) and aborts on error.

```
#define STARPU_PTHREAD_RWLOCK_UNLOCK( rwlock )
```

This macro calls the function `starpu_pthread_rwlock_unlock()` and aborts on error.

```
#define STARPU_PTHREAD_RWLOCK_DESTROY( rwlock )
```

This macro calls the function `starpu_pthread_rwlock_destroy()` and aborts on error.

```
#define STARPU_PTHREAD_COND_INIT( cond, attr )
```

This macro calls the function `starpu_pthread_cond_init()` and aborts on error.

```
#define STARPU_PTHREAD_COND_DESTROY( cond )
```

This macro calls the function `starpu_pthread_cond_destroy()` and aborts on error.

```
#define STARPU_PTHREAD_COND_SIGNAL( cond )
```

This macro calls the function `starpu_pthread_cond_signal()` and aborts on error.

```
#define STARPU_PTHREAD_COND_BROADCAST( cond )
```

This macro calls the function `starpu_pthread_cond_broadcast()` and aborts on error.

```
#define STARPU_PTHREAD_COND_WAIT( cond, mutex )
```

This macro calls the function `starpu_pthread_cond_wait()` and aborts on error.

```
#define STARPU_PTHREAD_BARRIER_INIT( barrier, attr, count )
```

This macro calls the function `starpu_pthread_barrier_init()` and aborts on error.

```
#define STARPU_PTHREAD_BARRIER_DESTROY( barrier )
```

This macro calls the function `starpu_pthread_barrier_destroy()` and aborts on error.

```
#define STARPU_PTHREAD_BARRIER_WAIT( barrier )
```

This macro calls the function `starpu_pthread_barrier_wait()` and aborts on error.

## **STARPU\_PTHREAD\_MUTEX\_INITIALIZER**

This macro initializes the mutex given in parameter.

## **STARPU\_PTHREAD\_COND\_INITIALIZER**

This macro initializes the condition variable given in parameter.

### **17.5.3 Function Documentation**

```
int starpu_pthread_create ( starpu_pthread_t * thread, const starpu_pthread_attr_t * attr, void *(*)(void *) start_routine, void * arg )
```

This function starts a new thread in the calling process. The new thread starts execution by invoking `start_routine`; `arg` is passed as the sole argument of `start_routine`.

```
int starpu_pthread_join ( starpu_pthread_t thread, void ** retval )
```

This function waits for the thread specified by `thread` to terminate. If that thread has already terminated, then the function returns immediately. The thread specified by `thread` must be joinable.

**int starpu\_pthread\_attr\_init ( starpu\_pthread\_attr\_t \* attr )**

This function initializes the thread attributes object pointed to by `attr` with default attribute values.

It does not do anything when the simulated performance mode is enabled ([Simulated Performance](#)).

**int starpu\_pthread\_attr\_destroy ( starpu\_pthread\_attr\_t \* attr )**

This function destroys a thread attributes object which is no longer required. Destroying a thread attributes object has no effect on threads that were created using that object.

It does not do anything when the simulated performance mode is enabled ([Simulated Performance](#)).

**int starpu\_pthread\_attr\_setdetachstate ( starpu\_pthread\_attr\_t \* attr, int detachstate )**

This function sets the detach state attribute of the thread attributes object referred to by `attr` to the value specified in `detachstate`. The detach state attribute determines whether a thread created using the thread attributes object `attr` will be created in a joinable or a detached state.

It does not do anything when the simulated performance mode is enabled ([Simulated Performance](#)).

**int starpu\_pthread\_mutex\_init ( starpu\_pthread\_mutex\_t \* mutex, const starpu\_pthread\_mutexattr\_t \* mutexattr )**

This function initializes the mutex object pointed to by `mutex` according to the mutex attributes specified in `mutexattr`. If `mutexattr` is NULL, default attributes are used instead.

**int starpu\_pthread\_mutex\_destroy ( starpu\_pthread\_mutex\_t \* mutex )**

This function destroys a mutex object, freeing the resources it might hold. The mutex must be unlocked on entrance.

**int starpu\_pthread\_mutex\_lock ( starpu\_pthread\_mutex\_t \* mutex )**

This function locks the given mutex. If the mutex is currently unlocked, it becomes locked and owned by the calling thread, and the function returns immediately. If the mutex is already locked by another thread, the function suspends the calling thread until the mutex is unlocked.

**int starpu\_pthread\_mutex\_unlock ( starpu\_pthread\_mutex\_t \* mutex )**

This function unlocks the given mutex. The mutex is assumed to be locked and owned by the calling thread on entrance to `starpu_pthread_mutex_unlock()`.

**int starpu\_pthread\_mutex\_trylock ( starpu\_pthread\_mutex\_t \* mutex )**

This function behaves identically to `starpu_pthread_mutex_lock()`, except that it does not block the calling thread if the mutex is already locked by another thread (or by the calling thread in the case of a “fast” mutex). Instead, the function returns immediately with the error code EBUSY.

**int starpu\_pthread\_key\_create ( starpu\_pthread\_key\_t \* key, void(\*)(void \*) destr\_function )**

This function allocates a new TSD key. The key is stored in the location pointed to by `key`.

**int starpu\_pthread\_key\_delete ( starpu\_pthread\_key\_t key )**

This function deallocates a TSD key. It does not check whether non-NULL values are associated with that key in the currently executing threads, nor call the destructor function associated with the key.

**int starpu\_pthread\_setspecific ( starpu\_pthread\_key\_t key, const void \* pointer )**

This function changes the value associated with `key` in the calling thread, storing the given `pointer` instead.

**\* `starpu_pthread_getspecific ( starpu_pthread_key_t key )`**

This function returns the value associated with `key` on success, and `NULL` on error.

**`starpu_pthread_cond_init ( starpu_pthread_cond_t * cond, starpu_pthread_condattr_t * cond_attr )`**

This function initializes the condition variable `cond`, using the condition attributes specified in `cond_attr`, or default attributes if `cond_attr` is `NULL`.

**`starpu_pthread_cond_signal ( starpu_pthread_cond_t * cond )`**

This function restarts one of the threads that are waiting on the condition variable `cond`. If no threads are waiting on `cond`, nothing happens. If several threads are waiting on `cond`, exactly one is restarted, but it not specified which.

**`starpu_pthread_cond_broadcast ( starpu_pthread_cond_t * cond )`**

This function restarts all the threads that are waiting on the condition variable `cond`. Nothing happens if no threads are waiting on `cond`.

**`starpu_pthread_cond_wait ( starpu_pthread_cond_t * cond, starpu_pthread_mutex_t * mutex )`**

This function atomically unlocks the `mutex` (as per [starpu\\_pthread\\_mutex\\_unlock\(\)](#)) and waits for the condition variable `cond` to be signaled. The thread execution is suspended and does not consume any CPU time until the condition variable is signaled. The `mutex` must be locked by the calling thread on entrance to [starpu\\_pthread\\_cond\\_wait\(\)](#). Before returning to the calling thread, the function re-acquires `mutex` (as per [starpu\\_pthread\\_mutex\\_lock\(\)](#)).

**`starpu_pthread_cond_timedwait ( starpu_pthread_cond_t * cond, starpu_pthread_mutex_t * mutex, const struct timespec * abstime )`**

This function atomically unlocks `mutex` and waits on `cond`, as [starpu\\_pthread\\_cond\\_wait\(\)](#) does, but it also bounds the duration of the wait.

**`starpu_pthread_cond_destroy ( starpu_pthread_cond_t * cond )`**

This function destroys a condition variable, freeing the resources it might hold. No threads must be waiting on the condition variable on entrance to the function.

**`starpu_pthread_rwlock_init ( starpu_pthread_rwlock_t * rwlock, const starpu_pthread_rwlockattr_t * attr )`**

This function is the same as [starpu\\_pthread\\_mutex\\_init\(\)](#).

**`starpu_pthread_rwlock_destroy ( starpu_pthread_rwlock_t * rwlock )`**

This function is the same as [starpu\\_pthread\\_mutex\\_destroy\(\)](#).

**`starpu_pthread_rwlock_rdlock ( starpu_pthread_rwlock_t * rwlock )`**

This function is the same as [starpu\\_pthread\\_mutex\\_lock\(\)](#).

**`starpu_pthread_rwlock_wrlock ( starpu_pthread_rwlock_t * rwlock )`**

This function is the same as [starpu\\_pthread\\_mutex\\_lock\(\)](#).

**`starpu_pthread_rwlock_unlock ( starpu_pthread_rwlock_t * rwlock )`**

This function is the same as [starpu\\_pthread\\_mutex\\_unlock\(\)](#).

## 17.6 Workers' Properties

### Data Structures

- struct `starpu_worker_collection`
- struct `starpu_sched_ctx_iterator`

### Macros

- `#define STARPU_NMAXWORKERS`

### Enumerations

- enum `starpu_node_kind` { `STARPU_UNUSED`, `STARPU_CPU_RAM`, `STARPU_CUDA_RAM`, `STARPU_OPENCL_RAM` }
- enum `starpu_worker_archtype` { `STARPU_ANY_WORKER`, `STARPU_CPU_WORKER`, `STARPU_CUDA_WORKER`, `STARPU_OPENCL_WORKER` }
- enum `starpu_worker_collection_type` { `STARPU_WORKER_LIST` }

### Functions

- unsigned `starpu_worker_get_count` (void)
- int `starpu_worker_get_count_by_type` (enum `starpu_worker_archtype` type)
- unsigned `starpu_cpu_worker_get_count` (void)
- unsigned `starpu_cuda_worker_get_count` (void)
- unsigned `starpu_opencl_worker_get_count` (void)
- int `starpu_worker_get_id` (void)
- int `starpu_worker_get_ids_by_type` (enum `starpu_worker_archtype` type, int \*workerids, int maxsize)
- int `starpu_worker_get_by_type` (enum `starpu_worker_archtype` type, int num)
- int `starpu_worker_get_by_devid` (enum `starpu_worker_archtype` type, int devid)
- int `starpu_worker_get_devid` (int id)
- enum `starpu_worker_archtype starpu_worker_get_type` (int id)
- void `starpu_worker_get_name` (int id, char \*dst, size\_t maxlen)
- unsigned `starpu_worker_get_memory_node` (unsigned workerid)
- enum `starpu_node_kind starpu_node_get_kind` (unsigned node)

#### 17.6.1 Detailed Description

#### 17.6.2 Data Structure Documentation

##### `struct starpu_worker_collection`

A scheduling context manages a collection of workers that can be memorized using different data structures. Thus, a generic structure is available in order to simplify the choice of its type. Only the list data structure is available but further data structures (like tree) implementations are foreseen.

##### Data Fields

- void \* `workerids`
- unsigned `nworkers`
- enum `starpu_worker_collection_type` type
- unsigned(\* `has_next`) (struct `starpu_worker_collection` \*workers, struct `starpu_sched_ctx_iterator` \*it)
- int(\* `get_next`) (struct `starpu_worker_collection` \*workers, struct `starpu_sched_ctx_iterator` \*it)
- int(\* `add`) (struct `starpu_worker_collection` \*workers, int worker)
- int(\* `remove`) (struct `starpu_worker_collection` \*workers, int worker)
- void(\* `init`) (struct `starpu_worker_collection` \*workers)
- void(\* `deinit`) (struct `starpu_worker_collection` \*workers)
- void(\* `init_iterator`) (struct `starpu_worker_collection` \*workers, struct `starpu_sched_ctx_iterator` \*it)

### Field Documentation

**starpu\_worker\_collection::workerids** The workerids managed by the collection

**starpu\_worker\_collection::nworkers** The number of workers in the collection

**starpu\_worker\_collection::type** The type of structure (currently [STARPU\\_WORKER\\_LIST](#) is the only one available)

**starpu\_worker\_collection::has\_next** Checks if there is another element in collection

**starpu\_worker\_collection::get\_next** return the next element in the collection

**starpu\_worker\_collection::add** add a new element in the collection

**starpu\_worker\_collection::remove** remove an element from the collection

**starpu\_worker\_collection::init** Initialize the collection

**starpu\_worker\_collection::deinit** Deinitialize the collection

**starpu\_worker\_collection::init\_iterator** Initialize the cursor if there is one

### struct starpu\_sched\_ctx.iterator

Structure needed to iterate on the collection

#### Data Fields

int	cursor	The index of the current worker in the collection, needed when iterating on the collection.
-----	--------	---

### 17.6.3 Macro Definition Documentation

**#define STARPU\_NMAXWORKERS**

Define the maximum number of workers managed by StarPU.

### 17.6.4 Enumeration Type Documentation

**enum starpu\_node\_kind**

TODO

Enumerator

**STARPU\_UNUSED** TODO

**STARPU\_CPU\_RAM** TODO

**STARPU\_CUDA\_RAM** TODO

**STARPU\_OPENCL\_RAM** TODO

**enum starpu\_worker\_archtype**

Worker Architecture Type

Enumerator

**STARPU\_ANY\_WORKER** any worker, used in the hypervisor

**STARPU\_CPU\_WORKER** CPU core

**STARPU\_CUDA\_WORKER** NVIDIA CUDA device

**STARPU\_OPENCL\_WORKER** OpenCL device

**enum starpu\_worker\_collection\_type**

Types of structures the worker collection can implement

Enumerator

**STARPU\_WORKER\_LIST** The collection is an array

**17.6.5 Function Documentation****unsigned starpu\_worker\_get\_count( void )**

This function returns the number of workers (i.e. processing units executing StarPU tasks). The returned value should be at most [STARPU\\_NMAXWORKERS](#).

**int starpu\_worker\_get\_count\_by\_type( enum starpu\_worker\_archtype type )**

Returns the number of workers of the given type. A positive (or NULL) value is returned in case of success, -EINVAL indicates that the type is not valid otherwise.

**unsigned starpu\_cpu\_worker\_get\_count( void )**

This function returns the number of CPUs controlled by StarPU. The returned value should be at most [STARPU\\_MAXCPUS](#).

**unsigned starpu\_cuda\_worker\_get\_count( void )**

This function returns the number of CUDA devices controlled by StarPU. The returned value should be at most [STARPU\\_MAXCUDADEVS](#).

**unsigned starpu\_opencl\_worker\_get\_count( void )**

This function returns the number of OpenCL devices controlled by StarPU. The returned value should be at most [STARPU\\_MAXOPENCLDEVS](#).

**int starpu\_worker\_get\_id( void )**

This function returns the identifier of the current worker, i.e the one associated to the calling thread. The returned value is either -1 if the current context is not a StarPU worker (i.e. when called from the application outside a task or a callback), or an integer between 0 and [starpu\\_worker\\_get\\_count\(\)](#) - 1.

**int starpu\_worker\_get\_ids\_by\_type( enum starpu\_worker\_archtype type, int \* workerids, int maxsize )**

This function gets the list of identifiers of workers with the given type. It fills the array `workerids` with the identifiers of the workers that have the type indicated in the first argument. The argument `maxsize` indicates the size of the array `workerids`. The returned value gives the number of identifiers that were put in the array. -ERANGE is returned if `maxsize` is lower than the number of workers with the appropriate type: in that case, the array is filled with the `maxsize` first elements. To avoid such overflows, the value of `maxsize` can be chosen by the means of the function [starpu\\_worker\\_get\\_count\\_by\\_type\(\)](#), or by passing a value greater or equal to [STARPU\\_NMAXWORKERS](#).

**int starpu\_worker\_get\_by\_type ( enum starpu\_worker\_archtype type, int num )**

This returns the identifier of the num-th worker that has the specified type type. If there are no such worker, -1 is returned.

**int starpu\_worker\_get\_by\_devid ( enum starpu\_worker\_archtype type, int devid )**

This returns the identifier of the worker that has the specified type type and device id devid (which may not be the n-th, if some devices are skipped for instance). If there are no such worker, -1 is returned.

**int starpu\_worker\_get\_devid ( int id )**

This function returns the device id of the given worker. The worker should be identified with the value returned by the [starpu\\_worker\\_get\\_id\(\)](#) function. In the case of a CUDA worker, this device identifier is the logical device identifier exposed by CUDA (used by the function `cudaGetDevice()` for instance). The device identifier of a CPU worker is the logical identifier of the core on which the worker was bound; this identifier is either provided by the OS or by the library `hwloc` in case it is available.

**enum starpu\_worker\_archtype starpu\_worker\_get\_type ( int id )**

This function returns the type of processing unit associated to a worker. The worker identifier is a value returned by the function [starpu\\_worker\\_get\\_id\(\)](#). The returned value indicates the architecture of the worker: **STARPU\_CPU\_WORKER** for a CPU core, **STARPU\_CUDA\_WORKER** for a CUDA device, and **STARPU\_OPENCL\_WORKER** for a OpenCL device. The value returned for an invalid identifier is unspecified.

**void starpu\_worker\_get\_name ( int id, char \* dst, size\_t maxlen )**

This function allows to get the name of a given worker. StarPU associates a unique human readable string to each processing unit. This function copies at most the maxlen first bytes of the unique string associated to a worker identified by its identifier id into the dst buffer. The caller is responsible for ensuring that dst is a valid pointer to a buffer of maxlen bytes at least. Calling this function on an invalid identifier results in an unspecified behaviour.

**unsigned starpu\_worker\_get\_memory\_node ( unsigned workerid )**

This function returns the identifier of the memory node associated to the worker identified by workerid.

**enum starpu\_node\_kind starpu\_node\_get\_kind ( unsigned node )**

Returns the type of the given node as defined by [starpu\\_node\\_kind](#). For example, when defining a new data interface, this function should be used in the allocation function to determine on which device the memory needs to be allocated.

## 17.7 Data Management

This section describes the data management facilities provided by StarPU. We show how to use existing data interfaces in [Data Interfaces](#), but developers can design their own data interfaces if required.

### Typedefs

- `typedef struct _starpu_data_state * starpu_data_handle_t`

### Enumerations

- `enum starpu_data_access_mode {  
 STARPU_NONE, STARPU_R, STARPU_W, STARPU_RW,  
 STARPU_SCRATCH, STARPU_REDUX }  
}`

## Basic Data Management API

Data management is done at a high-level in StarPU: rather than accessing a mere list of contiguous buffers, the tasks may manipulate data that are described by a high-level construct which we call data interface.

An example of data interface is the "vector" interface which describes a contiguous data array on a specific memory node. This interface is a simple structure containing the number of elements in the array, the size of the elements, and the address of the array in the appropriate address space (this address may be invalid if there is no valid copy of the array in the memory node). More informations on the data interfaces provided by StarPU are given in [Data Interfaces](#).

When a piece of data managed by StarPU is used by a task, the task implementation is given a pointer to an interface describing a valid copy of the data that is accessible from the current processing unit.

Every worker is associated to a memory node which is a logical abstraction of the address space from which the processing unit gets its data. For instance, the memory node associated to the different CPU workers represents main memory (RAM), the memory node associated to a GPU is DRAM embedded on the device. Every memory node is identified by a logical index which is accessible from the function `starpu_worker_get_memory_node()`. When registering a piece of data to StarPU, the specified memory node indicates where the piece of data initially resides (we also call this memory node the home node of a piece of data).

- void `starpu_data_register (starpu_data_handle_t *handleptr, unsigned home_node, void *data_interface, struct starpu_data_interface_ops *ops)`
- void `starpu_data_register_same (starpu_data_handle_t *handledst, starpu_data_handle_t handlesrc)`
- void `starpu_data_unregister (starpu_data_handle_t handle)`
- void `starpu_data_unregister_no_coherency (starpu_data_handle_t handle)`
- void `starpu_data_unregister_submit (starpu_data_handle_t handle)`
- void `starpu_data_invalidate (starpu_data_handle_t handle)`
- void `starpu_data_invalidate_submit (starpu_data_handle_t handle)`
- void `starpu_data_set_wt_mask (starpu_data_handle_t handle, uint32_t wt_mask)`
- int `starpu_data_prefetch_on_node (starpu_data_handle_t handle, unsigned node, unsigned async)`
- `starpu_data_handle_t starpu_data_lookup (const void *ptr)`
- int `starpu_data_request_allocation (starpu_data_handle_t handle, unsigned node)`
- void `starpu_data_query_status (starpu_data_handle_t handle, int memory_node, int *is_allocated, int *is_valid, int *is_requested)`
- void `starpu_data_advise_as_important (starpu_data_handle_t handle, unsigned is_important)`
- void `starpu_data_set_reduction_methods (starpu_data_handle_t handle, struct starpu_codelet *redux_cl, struct starpu_codelet *init_cl)`

## Access registered data from the application

- #define `STARPU_DATA_ACQUIRE_CB(handle, mode, code)`
- int `starpu_data_acquire (starpu_data_handle_t handle, enum starpu_data_access_mode mode)`
- int `starpu_data_acquire_cb (starpu_data_handle_t handle, enum starpu_data_access_mode mode, void(*callback)(void *), void *arg)`
- int `starpu_data_acquire_on_node (starpu_data_handle_t handle, unsigned node, enum starpu_data_access_mode mode)`
- int `starpu_data_acquire_on_node_cb (starpu_data_handle_t handle, unsigned node, enum starpu_data_access_mode mode, void(*callback)(void *), void *arg)`
- void `starpu_data_release (starpu_data_handle_t handle)`
- void `starpu_data_release_on_node (starpu_data_handle_t handle, unsigned node)`

### 17.7.1 Detailed Description

This section describes the data management facilities provided by StarPU. We show how to use existing data interfaces in [Data Interfaces](#), but developers can design their own data interfaces if required.

### 17.7.2 Macro Definition Documentation

```
#define STARPU_DATA_ACQUIRE_CB( handle, mode, code )
```

`STARPU_DATA_ACQUIRE_CB()` is the same as `starpu_data_acquire_cb()`, except that the code to be executed in a callback is directly provided as a macro parameter, and the data handle is automatically released after it. This permits to easily execute code which depends on the value of some registered data. This is non-blocking too and may be called from task callbacks.

### 17.7.3 Typedef Documentation

`starpu_data_handle_t`

StarPU uses `starpu_data_handle_t` as an opaque handle to manage a piece of data. Once a piece of data has been registered to StarPU, it is associated to a `starpu_data_handle_t` which keeps track of the state of the piece of data over the entire machine, so that we can maintain data consistency and locate data replicates for instance.

### 17.7.4 Enumeration Type Documentation

`enum starpu_data_access_mode`

This datatype describes a data access mode.

Enumerator

**`STARPU_NONE`** TODO

**`STARPU_R`** read-only mode.

**`STARPU_W`** write-only mode.

**`STARPU_RW`** read-write mode. This is equivalent to `STARPU_R|STARPU_W`

**`STARPU_SCRATCH`** A temporary buffer is allocated for the task, but StarPU does not enforce data consistency—i.e. each device has its own buffer, independently from each other (even for CPUs), and no data transfer is ever performed. This is useful for temporary variables to avoid allocating/freeing buffers inside each task. Currently, no behavior is defined concerning the relation with the `STARPU_R` and `STARPU_W` modes and the value provided at registration — i.e., the value of the scratch buffer is undefined at entry of the codelet function. It is being considered for future extensions at least to define the initial value. For now, data to be used in `STARPU_SCRATCH` mode should be registered with node `-1` and a `NULL` pointer, since the value of the provided buffer is simply ignored for now.

**`STARPU_REDUX`** todo

### 17.7.5 Function Documentation

```
void starpu_data_register ( starpu_data_handle_t * handleptr, unsigned home_node, void * data_interface,
                           struct starpu_data_interface_ops * ops )
```

Register a piece of data into the handle located at the `handleptr` address. The `data_interface` buffer contains the initial description of the data in the `home_node`. The `ops` argument is a pointer to a structure describing the different methods used to manipulate this type of interface. See `starpu_data_interface_ops` for more details on this structure. If `home_node` is `-1`, StarPU will automatically allocate the memory when it is used for the first time in write-only mode. Once such data handle has been automatically allocated, it is possible to access it using any access mode. Note that StarPU supplies a set of predefined types of interface (e.g. vector or matrix) which can be registered by the means of helper functions (e.g. `starpu_vector_data_register()` or `starpu_matrix_data_register()`).

```
void starpu_data_register_same ( starpu_data_handle_t * handledst, starpu_data_handle_t handlesrc )
```

Register a new piece of data into the handle `handledst` with the same interface as the handle `handlesrc`.

**void starpu\_data\_unregister ( starpu\_data\_handle\_t handle )**

This function unregisters a data handle from StarPU. If the data was automatically allocated by StarPU because the home node was -1, all automatically allocated buffers are freed. Otherwise, a valid copy of the data is put back into the home node in the buffer that was initially registered. Using a data handle that has been unregistered from StarPU results in an undefined behaviour. In case we do not need to update the value of the data in the home node, we can use the function [starpu\\_data\\_unregister\\_no\\_coherency\(\)](#) instead.

**void starpu\_data\_unregister\_no\_coherency ( starpu\_data\_handle\_t handle )**

This is the same as [starpu\\_data\\_unregister\(\)](#), except that StarPU does not put back a valid copy into the home node, in the buffer that was initially registered.

**void starpu\_data\_unregister\_submit ( starpu\_data\_handle\_t handle )**

Destroy the data handle once it is not needed anymore by any submitted task. No coherency is assumed.

**void starpu\_data\_invalidate ( starpu\_data\_handle\_t handle )**

Destroy all replicates of the data handle immediately. After data invalidation, the first access to the handle must be performed in write-only mode. Accessing an invalidated data in read-mode results in undefined behaviour.

**void starpu\_data\_invalidate\_submit ( starpu\_data\_handle\_t handle )**

Submits invalidation of the data handle after completion of previously submitted tasks.

**void starpu\_data\_set\_wt\_mask ( starpu\_data\_handle\_t handle, uint32\_t wt\_mask )**

This function sets the write-through mask of a given data (and its children), i.e. a bitmask of nodes where the data should be always replicated after modification. It also prevents the data from being evicted from these nodes when memory gets scarce. When the data is modified, it is automatically transferred into those memory node. For instance a  $1 << 0$  write-through mask means that the CUDA workers will commit their changes in main memory (node 0).

**int starpu\_data\_prefetch\_on\_node ( starpu\_data\_handle\_t handle, unsigned node, unsigned async )**

Issue a prefetch request for a given data to a given node, i.e. requests that the data be replicated to the given node, so that it is available there for tasks. If the `async` parameter is 0, the call will block until the transfer is achieved, else the call will return as soon as the request is scheduled (which may however have to wait for a task completion).

**starpu\_data\_handle\_t starpu\_data\_lookup ( const void \* ptr )**

Return the handle corresponding to the data pointed to by the `ptr` host pointer.

**int starpu\_data\_request\_allocation ( starpu\_data\_handle\_t handle, unsigned node )**

Explicitly ask StarPU to allocate room for a piece of data on the specified memory node.

**void starpu\_data\_query\_status ( starpu\_data\_handle\_t handle, int \* memory\_node, int \* is\_allocated, int \* is\_valid, int \* is\_requested )**

Query the status of `handle` on the specified `memory_node`.

**void starpu\_data\_advise\_as\_important ( starpu\_data\_handle\_t handle, unsigned is\_important )**

This function allows to specify that a piece of data can be discarded without impacting the application.

```
void starpu_data_set_reduction_methods ( starpu_data_handle_t handle, struct starpu_codelet * redux_cl,
struct starpu_codelet * init_cl )
```

This sets the codelets to be used for `handle` when it is accessed in the mode `STARPU_REDUX`. Per-worker buffers will be initialized with the codelet `init_cl`, and reduction between per-worker buffers will be done with the codelet `redux_cl`.

```
int starpu_data_acquire ( starpu_data_handle_t handle, enum starpu_data_access_mode mode )
```

The application must call this function prior to accessing registered data from main memory outside tasks. StarPU ensures that the application will get an up-to-date copy of the data in main memory located where the data was originally registered, and that all concurrent accesses (e.g. from tasks) will be consistent with the access mode specified in the mode argument. `starpu_data_release()` must be called once the application does not need to access the piece of data anymore. Note that implicit data dependencies are also enforced by `starpu_data_acquire()`, i.e. `starpu_data_acquire()` will wait for all tasks scheduled to work on the data, unless they have been disabled explicitly by calling `starpu_data_set_default_sequential_consistency_flag()` or `starpu_data_set_sequential_consistency_flag()`. `starpu_data_acquire()` is a blocking call, so that it cannot be called from tasks or from their callbacks (in that case, `starpu_data_acquire()` returns `-EDEADLK`). Upon successful completion, this function returns 0.

```
int starpu_data_acquire_cb ( starpu_data_handle_t handle, enum starpu_data_access_mode mode,
void(*)(void *) callback, void * arg )
```

Asynchronous equivalent of `starpu_data_acquire()`. When the data specified in `handle` is available in the appropriate access mode, the `callback` function is executed. The application may access the requested data during the execution of this `callback`. The `callback` function must call `starpu_data_release()` once the application does not need to access the piece of data anymore. Note that implicit data dependencies are also enforced by `starpu_data_acquire_cb()` in case they are not disabled. Contrary to `starpu_data_acquire()`, this function is non-blocking and may be called from task callbacks. Upon successful completion, this function returns 0.

```
int starpu_data_acquire_on_node ( starpu_data_handle_t handle, unsigned node, enum starpu_data_access_mode mode )
```

This is the same as `starpu_data_acquire()`, except that the data will be available on the given memory node instead of main memory.

```
int starpu_data_acquire_on_node_cb ( starpu_data_handle_t handle, unsigned node, enum starpu_data_access_mode mode, void(*)(void *) callback, void * arg )
```

This is the same as `starpu_data_acquire_cb()`, except that the data will be available on the given memory node instead of main memory.

```
void starpu_data_release ( starpu_data_handle_t handle )
```

This function releases the piece of data acquired by the application either by `starpu_data_acquire()` or by `starpu_data_acquire_cb()`.

```
void starpu_data_release_on_node ( starpu_data_handle_t handle, unsigned node )
```

This is the same as `starpu_data_release()`, except that the data will be available on the given memory node instead of main memory.

## 17.8 Data Interfaces

### Data Structures

- struct `starpu_data_interface_ops`
- struct `starpu_data_copy_methods`
- struct `starpu_variable_interface`
- struct `starpu_vector_interface`

- struct `starpu_matrix_interface`
- struct `starpu_block_interface`
- struct `starpu_bcsr_interface`
- struct `starpu_csr_interface`
- struct `starpu_coo_interface`

## Enumerations

- enum `starpu_data_interface_id` {
 `STARPU_UNKNOWN_INTERFACE_ID`, `STARPU_MATRIX_INTERFACE_ID`, `STARPU_BLOCK_INTERFACE_ID`, `STARPU_VECTOR_INTERFACE_ID`, `STARPU_CSR_INTERFACE_ID`, `STARPU_BCSR_INTERFACE_ID`, `STARPU_VARIABLE_INTERFACE_ID`, `STARPU_VOID_INTERFACE_ID`, `STARPU_MULTIFORMAT_INTERFACE_ID`, `STARPU_COO_INTERFACE_ID`, `STARPU_MAX_INTERFACE_ID` }

## Registering Data

There are several ways to register a memory region so that it can be managed by StarPU. The functions below allow the registration of vectors, 2D matrices, 3D matrices as well as BCSR and CSR sparse matrices.

- void `starpu_void_data_register` (`starpu_data_handle_t` \*handle)
- void `starpu_variable_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, `uintptr_t` ptr, `size_t` size)
- void `starpu_vector_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, `uintptr_t` ptr, `uint32_t` nx, `size_t` elemsize)
- void `starpu_matrix_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, `uintptr_t` ptr, `uint32_t` ld, `uint32_t` nx, `uint32_t` ny, `size_t` elemsize)
- void `starpu_block_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, `uintptr_t` ptr, `uint32_t` ldy, `uint32_t` ldz, `uint32_t` nx, `uint32_t` ny, `uint32_t` nz, `size_t` elemsize)
- void `starpu_bcsr_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, `uint32_t` nnz, `uint32_t` nrow, `uintptr_t` nzval, `uint32_t` \*colind, `uint32_t` \*rowptr, `uint32_t` firstentry, `uint32_t` r, `uint32_t` c, `size_t` elemsize)
- void `starpu_csr_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, `uint32_t` nnz, `uint32_t` nrow, `uintptr_t` nzval, `uint32_t` \*colind, `uint32_t` \*rowptr, `uint32_t` firstentry, `size_t` elemsize)
- void `starpu_coo_data_register` (`starpu_data_handle_t` \*handleptr, unsigned home\_node, `uint32_t` nx, `uint32_t` ny, `uint32_t` n\_values, `uint32_t` \*columns, `uint32_t` \*rows, `uintptr_t` values, `size_t` elemsize)
- void \* `starpu_data_get_interface_on_node` (`starpu_data_handle_t` handle, unsigned memory\_node)

## Accessing Data Interfaces

Each data interface is provided with a set of field access functions. The ones using a void \* parameter aimed to be used in codelet implementations (see for example the code in [Vector Scaling Using StarPU's API](#)).

- void \* `starpu_data_handle_to_pointer` (`starpu_data_handle_t` handle, unsigned node)
- void \* `starpu_data_get_local_ptr` (`starpu_data_handle_t` handle)
- enum `starpu_data_interface_id` `starpu_data_get_interface_id` (`starpu_data_handle_t` handle)
- `size_t` `starpu_data_get_size` (`starpu_data_handle_t` handle)
- int `starpu_data_pack` (`starpu_data_handle_t` handle, void \*\*ptr, `starpu_ssize_t` \*count)
- int `starpu_data_unpack` (`starpu_data_handle_t` handle, void \*ptr, `size_t` count)

## Accessing Variable Data Interfaces

- #define `STARPU_VARIABLE_GET_PTR`(interface)
- #define `STARPU_VARIABLE_GET_ELEMSIZE`(interface)
- #define `STARPU_VARIABLE_GET_DEV_HANDLE`(interface)
- #define `STARPU_VARIABLE_GET_OFFSET`
- `size_t` `starpu_variable_get_elemsize` (`starpu_data_handle_t` handle)
- `uintptr_t` `starpu_variable_get_local_ptr` (`starpu_data_handle_t` handle)

## Accessing Vector Data Interfaces

- #define STARPU\_VECTOR\_GET\_PTR(interface)
- #define STARPU\_VECTOR\_GET\_DEV\_HANDLE(interface)
- #define STARPU\_VECTOR\_GET\_OFFSET(interface)
- #define STARPU\_VECTOR\_GET\_NX(interface)
- #define STARPU\_VECTOR\_GET\_ELEMSIZE(interface)
- uint32\_t starpu\_vector\_get\_nx (starpu\_data\_handle\_t handle)
- size\_t starpu\_vector\_get\_elemsize (starpu\_data\_handle\_t handle)
- uintptr\_t starpu\_vector\_get\_local\_ptr (starpu\_data\_handle\_t handle)

## Accessing Matrix Data Interfaces

- #define STARPU\_MATRIX\_GET\_PTR(interface)
- #define STARPU\_MATRIX\_GET\_DEV\_HANDLE(interface)
- #define STARPU\_MATRIX\_GET\_OFFSET(interface)
- #define STARPU\_MATRIX\_GET\_NX(interface)
- #define STARPU\_MATRIX\_GET\_NY(interface)
- #define STARPU\_MATRIX\_GET\_LD(interface)
- #define STARPU\_MATRIX\_GET\_ELEMSIZE(interface)
- uint32\_t starpu\_matrix\_get\_nx (starpu\_data\_handle\_t handle)
- uint32\_t starpu\_matrix\_get\_ny (starpu\_data\_handle\_t handle)
- uint32\_t starpu\_matrix\_get\_local\_ld (starpu\_data\_handle\_t handle)
- uintptr\_t starpu\_matrix\_get\_local\_ptr (starpu\_data\_handle\_t handle)
- size\_t starpu\_matrix\_get\_elemsize (starpu\_data\_handle\_t handle)

## Accessing Block Data Interfaces

- #define STARPU\_BLOCK\_GET\_PTR(interface)
- #define STARPU\_BLOCK\_GET\_DEV\_HANDLE(interface)
- #define STARPU\_BLOCK\_GET\_OFFSET(interface)
- #define STARPU\_BLOCK\_GET\_NX(interface)
- #define STARPU\_BLOCK\_GET\_NY(interface)
- #define STARPU\_BLOCK\_GET\_NZ(interface)
- #define STARPU\_BLOCK\_GET\_LDY(interface)
- #define STARPU\_BLOCK\_GET\_LDZ(interface)
- #define STARPU\_BLOCK\_GET\_ELEMSIZE(interface)
- uint32\_t starpu\_block\_get\_nx (starpu\_data\_handle\_t handle)
- uint32\_t starpu\_block\_get\_ny (starpu\_data\_handle\_t handle)
- uint32\_t starpu\_block\_get\_nz (starpu\_data\_handle\_t handle)
- uint32\_t starpu\_block\_get\_local\_ldy (starpu\_data\_handle\_t handle)
- uint32\_t starpu\_block\_get\_local\_ldz (starpu\_data\_handle\_t handle)
- uintptr\_t starpu\_block\_get\_local\_ptr (starpu\_data\_handle\_t handle)
- size\_t starpu\_block\_get\_elemsize (starpu\_data\_handle\_t handle)

## Accessing BCSR Data Interfaces

- #define STARPU\_BCSR\_GET\_NNZ(interface)
- #define STARPU\_BCSR\_GET\_NZVAL(interface)
- #define STARPU\_BCSR\_GET\_NZVAL\_DEV\_HANDLE(interface)
- #define STARPU\_BCSR\_GET\_COLIND(interface)
- #define STARPU\_BCSR\_GET\_COLIND\_DEV\_HANDLE(interface)
- #define STARPU\_BCSR\_GET\_ROWPTR(interface)
- #define STARPU\_BCSR\_GET\_ROWPTR\_DEV\_HANDLE(interface)
- #define STARPU\_BCSR\_GET\_OFFSET
- uint32\_t starpu\_bcsr\_get\_nnz (starpu\_data\_handle\_t handle)

- `uint32_t starpu_bcsr_get_nrow (starpu_data_handle_t handle)`
- `uint32_t starpu_bcsr_get_firstentry (starpu_data_handle_t handle)`
- `uintptr_t starpu_bcsr_get_local_nzval (starpu_data_handle_t handle)`
- `uint32_t * starpu_bcsr_get_local_colind (starpu_data_handle_t handle)`
- `uint32_t * starpu_bcsr_get_local_rowptr (starpu_data_handle_t handle)`
- `uint32_t starpu_bcsr_get_r (starpu_data_handle_t handle)`
- `uint32_t starpu_bcsr_get_c (starpu_data_handle_t handle)`
- `size_t starpu_bcsr_get_elemsize (starpu_data_handle_t handle)`

## Accessing CSR Data Interfaces

- `#define STARPU_CSR_GET_NNZ(interface)`
- `#define STARPU_CSR_GET_NROW(interface)`
- `#define STARPU_CSR_GET_NZVAL(interface)`
- `#define STARPU_CSR_GET_NZVAL_DEV_HANDLE(interface)`
- `#define STARPU_CSR_GET_COLIND(interface)`
- `#define STARPU_CSR_GET_COLIND_DEV_HANDLE(interface)`
- `#define STARPU_CSR_GET_ROWPTR(interface)`
- `#define STARPU_CSR_GET_ROWPTR_DEV_HANDLE(interface)`
- `#define STARPU_CSR_GET_OFFSET`
- `#define STARPU_CSR_GET_FIRSTENTRY(interface)`
- `#define STARPU_CSR_GET_ELEMSIZE(interface)`
- `uint32_t starpu_csr_get_nnz (starpu_data_handle_t handle)`
- `uint32_t starpu_csr_get_nrow (starpu_data_handle_t handle)`
- `uint32_t starpu_csr_get_firstentry (starpu_data_handle_t handle)`
- `uintptr_t starpu_csr_get_local_nzval (starpu_data_handle_t handle)`
- `uint32_t * starpu_csr_get_local_colind (starpu_data_handle_t handle)`
- `uint32_t * starpu_csr_get_local_rowptr (starpu_data_handle_t handle)`
- `size_t starpu_csr_get_elemsize (starpu_data_handle_t handle)`

## Accessing COO Data Interfaces

- `#define STARPU_COO_GET_COLUMNS(interface)`
- `#define STARPU_COO_GET_COLUMNS_DEV_HANDLE(interface)`
- `#define STARPU_COO_GET_ROWS(interface)`
- `#define STARPU_COO_GET_ROWS_DEV_HANDLE(interface)`
- `#define STARPU_COO_GET_VALUES(interface)`
- `#define STARPU_COO_GET_VALUES_DEV_HANDLE(interface)`
- `#define STARPU_COO_GET_OFFSET`
- `#define STARPU_COO_GET_NX(interface)`
- `#define STARPU_COO_GET_NY(interface)`
- `#define STARPU_COO_GET_NVALUES(interface)`
- `#define STARPU_COO_GET_ELEMSIZE(interface)`

## Defining Interface

Applications can provide their own interface as shown in [Defining A New Data Interface](#).

- `uintptr_t starpu_malloc_on_node (unsigned dst_node, size_t size)`
- `void starpu_free_on_node (unsigned dst_node, uintptr_t addr, size_t size)`
- `int starpu_interface_copy (uintptr_t src, size_t src_offset, unsigned src_node, uintptr_t dst, size_t dst_offset, unsigned dst_node, size_t size, void *async_data)`
- `uint32_t starpu_hash_crc32c_be_n (const void *input, size_t n, uint32_t inputcrc)`
- `uint32_t starpu_hash_crc32c_be (uint32_t input, uint32_t inputcrc)`
- `uint32_t starpu_hash_crc32c_string (const char *str, uint32_t inputcrc)`
- `int starpu_data_interface_get_next_id (void)`

### 17.8.1 Detailed Description

### 17.8.2 Data Structure Documentation

#### `struct starpu_data_interface_ops`

Per-interface data transfer methods.

#### Data Fields

- `void(* register_data_handle )(starpu_data_handle_t handle, unsigned home_node, void *data_interface)`
- `starpu_ssize_t(* allocate_data_on_node )(void *data_interface, unsigned node)`
- `void(* free_data_on_node )(void *data_interface, unsigned node)`
- `struct starpu_data_copy_methods * copy_methods`
- `void **( handle_to_pointer )(starpu_data_handle_t handle, unsigned node)`
- `size_t(* get_size )(starpu_data_handle_t handle)`
- `uint32_t(* footprint )(starpu_data_handle_t handle)`
- `int(* compare )(void *data_interface_a, void *data_interface_b)`
- `void(* display )(starpu_data_handle_t handle, FILE *f)`
- `enum starpu_data_interface_id interfaceid`
- `size_t interface_size`
- `int is_multiformat`
- `struct starpu_multiformat_data_interface_ops *(* get_mf_ops )(void *data_interface)`
- `int(* pack_data )(starpu_data_handle_t handle, unsigned node, void **ptr, starpu_ssize_t *count)`
- `int(* unpack_data )(starpu_data_handle_t handle, unsigned node, void *ptr, size_t count)`

#### Field Documentation

**`starpu_data_interface_ops::register_data_handle`** Register an existing interface into a data handle.

**`starpu_data_interface_ops::allocate_data_on_node`** Allocate data for the interface on a given node.

**`starpu_data_interface_ops::free_data_on_node`** Free data of the interface on a given node.

**`starpu_data_interface_ops::copy_methods`** ram/cuda/opencl synchronous and asynchronous transfer methods.

**`starpu_data_interface_ops::handle_to_pointer`** Return the current pointer (if any) for the handle on the given node.

**`starpu_data_interface_ops::get_size`** Return an estimation of the size of data, for performance models.

**`starpu_data_interface_ops::footprint`** Return a 32bit footprint which characterizes the data size.

**`starpu_data_interface_ops::compare`** Compare the data size of two interfaces.

**`starpu_data_interface_ops::display`** Dump the sizes of a handle to a file.

**`starpu_data_interface_ops::interfaceid`** An identifier that is unique to each interface.

**`starpu_data_interface_ops::interface_size`** The size of the interface data descriptor.

**`starpu_data_interface_ops::is_multiformat`** todo

**starpu\_data\_interface\_ops::get\_mf\_ops** todo

**starpu\_data\_interface\_ops::pack\_data** Pack the data handle into a contiguous buffer at the address ptr and set the size of the newly created buffer in count. If ptr is NULL, the function should not copy the data in the buffer but just set count to the size of the buffer which would have been allocated. The special value -1 indicates the size is yet unknown.

**starpu\_data\_interface\_ops::unpack\_data** Unpack the data handle from the contiguous buffer at the address ptr of size count

### struct starpu\_data\_copy\_methods

Defines the per-interface methods. If the any\_to\_any method is provided, it will be used by default if no more specific method is provided. It can still be useful to provide more specific method in case of e.g. available particular CUDA or OpenCL support.

#### Data Fields

- int(\* **ram\_to\_ram** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node)
- int(\* **ram\_to\_cuda** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node)
- int(\* **ram\_to\_opencl** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node)
- int(\* **cuda\_to\_ram** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node)
- int(\* **cuda\_to\_cuda** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node)
- int(\* **cuda\_to\_opencl** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node)
- int(\* **opencl\_to\_ram** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node)
- int(\* **opencl\_to\_cuda** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node)
- int(\* **opencl\_to\_opencl** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node)
- int(\* **ram\_to\_cuda\_async** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node, cudaStream\_t stream)
- int(\* **cuda\_to\_ram\_async** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node, cudaStream\_t stream)
- int(\* **cuda\_to\_cuda\_async** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node, cudaStream\_t stream)
- int(\* **ram\_to\_opencl\_async** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node, cl\_event \*event)
- int(\* **opencl\_to\_ram\_async** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node, cl\_event \*event)
- int(\* **opencl\_to\_opencl\_async** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node, cl\_event \*event)
- int(\* **any\_to\_any** )(void \*src\_interface, unsigned src\_node, void \*dst\_interface, unsigned dst\_node, void \*async\_data)

#### Field Documentation

**starpu\_data\_copy\_methods::ram\_to\_ram** Define how to copy data from the src\_interface interface on the src\_node CPU node to the dst\_interface interface on the dst\_node CPU node. Return 0 on success.

**starpu\_data\_copy\_methods::ram\_to\_cuda** Define how to copy data from the src\_interface interface on the src\_node CPU node to the dst\_interface interface on the dst\_node CUDA node. Return 0 on success.

**starpu\_data\_copy\_methods::ram\_to\_opencl** Define how to copy data from the src\_interface interface on the src\_node CPU node to the dst\_interface interface on the dst\_node OpenCL node. Return 0 on success.

**starpu\_data\_copy\_methods::cuda\_to\_ram** Define how to copy data from the src\_interface interface on the src\_node CUDA node to the dst\_interface interface on the dst\_node CPU node. Return 0 on success.

**starpu\_data\_copy\_methods::cuda\_to\_cuda** Define how to copy data from the src\_interface interface on the src\_node CUDA node to the dst\_interface interface on the dst\_node CUDA node. Return 0 on success.

**starpu\_data\_copy\_methods::cuda\_to\_opencl** Define how to copy data from the src\_interface interface on the src\_node CUDA node to the dst\_interface interface on the dst\_node OpenCL node. Return 0 on success.

**starpu\_data\_copy\_methods::opencl\_to\_ram** Define how to copy data from the src\_interface interface on the src\_node OpenCL node to the dst\_interface interface on the dst\_node CPU node. Return 0 on success.

**starpu\_data\_copy\_methods::opencl\_to\_cuda** Define how to copy data from the src\_interface interface on the src\_node OpenCL node to the dst\_interface interface on the dst\_node CUDA node. Return 0 on success.

**starpu\_data\_copy\_methods::opencl\_to\_opencl** Define how to copy data from the src\_interface interface on the src\_node OpenCL node to the dst\_interface interface on the dst\_node OpenCL node. Return 0 on success.

**starpu\_data\_copy\_methods::ram\_to\_cuda\_async** Define how to copy data from the src\_interface interface on the src\_node CPU node to the dst\_interface interface on the dst\_node CUDA node, using the given stream. Must return 0 if the transfer was actually completed completely synchronously, or -EAGAIN if at least some transfers are still ongoing and should be awaited for by the core.

**starpu\_data\_copy\_methods::cuda\_to\_ram\_async** Define how to copy data from the src\_interface interface on the src\_node CUDA node to the dst\_interface interface on the dst\_node CPU node, using the given stream. Must return 0 if the transfer was actually completed completely synchronously, or -EAGAIN if at least some transfers are still ongoing and should be awaited for by the core.

**starpu\_data\_copy\_methods::cuda\_to\_cuda\_async** Define how to copy data from the src\_interface interface on the src\_node CUDA node to the dst\_interface interface on the dst\_node CUDA node, using the given stream. Must return 0 if the transfer was actually completed completely synchronously, or -EAGAIN if at least some transfers are still ongoing and should be awaited for by the core.

**starpu\_data\_copy\_methods::ram\_to\_opencl\_async** Define how to copy data from the src\_interface interface on the src\_node CPU node to the dst\_interface interface on the dst\_node OpenCL node, by recording in event, a pointer to a cl\_event, the event of the last submitted transfer. Must return 0 if the transfer was actually completed completely synchronously, or -EAGAIN if at least some transfers are still ongoing and should be awaited for by the core.

**starpu\_data\_copy\_methods::opencl\_to\_ram\_async** Define how to copy data from the src\_interface interface on the src\_node OpenCL node to the dst\_interface interface on the dst\_node CPU node, by recording in event, a pointer to a cl\_event, the event of the last submitted transfer. Must return 0 if the transfer was actually completed completely synchronously, or -EAGAIN if at least some transfers are still ongoing and should be awaited for by the core.

**starpu\_data\_copy\_methods::opencl\_to\_opencl\_async** Define how to copy data from the src\_interface interface on the src\_node OpenCL node to the dst\_interface interface on the dst\_node OpenCL node, by recording in event, a pointer to a cl\_event, the event of the last submitted transfer. Must return 0 if the transfer was actually completed completely synchronously, or -EAGAIN if at least some transfers are still ongoing and should be awaited for by the core.

**starpu\_data\_copy\_methods::any\_to\_any** Define how to copy data from the src\_interface interface on the src\_node node to the dst\_interface interface on the dst\_node node. This is meant to be implemented through the [starpu\\_interface\\_copy\(\)](#) helper, to which `async_data` should be passed as such, and will be used to manage asynchronicity. This must return -EAGAIN if any of the [starpu\\_interface\\_copy\(\)](#) calls has returned -EAGAIN (i.e. at least some transfer is still ongoing), and return 0 otherwise.

**struct starpu\_variable\_interface**

Variable interface for a single data (not a vector, a matrix, a list, ...)

**Data Fields**

<code>uintptr_t</code>	<code>ptr</code>	local pointer of the variable
<code>size_t</code>	<code>elemsize</code>	size of the variable

**struct starpu\_vector\_interface**

Vector interface

vector interface for contiguous (non-strided) buffers

**Data Fields**

<code>uintptr_t</code>	<code>ptr</code>	local pointer of the vector
<code>uintptr_t</code>	<code>dev_handle</code>	device handle of the vector.
<code>size_t</code>	<code>offset</code>	offset in the vector
<code>uint32_t</code>	<code>nx</code>	number of elements on the x-axis of the vector
<code>size_t</code>	<code>elemsize</code>	size of the elements of the vector

**struct starpu\_matrix\_interface**

Matrix interface for dense matrices

**Data Fields**

<code>uintptr_t</code>	<code>ptr</code>	local pointer of the matrix
<code>uintptr_t</code>	<code>dev_handle</code>	device handle of the matrix.
<code>size_t</code>	<code>offset</code>	offset in the matrix
<code>uint32_t</code>	<code>nx</code>	number of elements on the x-axis of the matrix
<code>uint32_t</code>	<code>ny</code>	number of elements on the y-axis of the matrix
<code>uint32_t</code>	<code>ld</code>	number of elements between each row of the matrix. Maybe be equal to <code>starpu_matrix_interface::nx</code> when there is no padding.
<code>size_t</code>	<code>elemsize</code>	size of the elements of the matrix

**struct starpu\_block\_interface**

Block interface for 3D dense blocks

**Data Fields**

<code>uintptr_t</code>	<code>ptr</code>	local pointer of the block
<code>uintptr_t</code>	<code>dev_handle</code>	device handle of the block.
<code>size_t</code>	<code>offset</code>	offset in the block.
<code>uint32_t</code>	<code>nx</code>	number of elements on the x-axis of the block.
<code>uint32_t</code>	<code>ny</code>	number of elements on the y-axis of the block.
<code>uint32_t</code>	<code>nz</code>	number of elements on the z-axis of the block.
<code>uint32_t</code>	<code>ldy</code>	number of elements between two lines
<code>uint32_t</code>	<code>ldz</code>	number of elements between two planes
<code>size_t</code>	<code>elemsize</code>	size of the elements of the block.

**struct starpu\_bcsr\_interface**

BCSR interface for sparse matrices (blocked compressed sparse row representation)

**Data Fields**

<code>uint32_t</code>	<code>nnz</code>	number of non-zero BLOCKS
<code>uint32_t</code>	<code>nrow</code>	number of rows (in terms of BLOCKS)

uintptr_t	nzval	non-zero values
uint32_t *	colind	position of non-zero entries on the row
uint32_t *	rowptr	index (in nzval) of the first entry of the row
uint32_t	firstentry	k for k-based indexing (0 or 1 usually). Also useful when partitioning the matrix.
uint32_t	r	size of the blocks
uint32_t	c	size of the blocks
size_t	elemsize	size of the elements of the matrix

**struct starpu\_csr\_interface**

CSR interface for sparse matrices (compressed sparse row representation)

## Data Fields

uint32_t	nnz	number of non-zero entries
uint32_t	nrow	number of rows
uintptr_t	nzval	non-zero values
uint32_t *	colind	position of non-zero entries on the row
uint32_t *	rowptr	index (in nzval) of the first entry of the row
uint32_t	firstentry	k for k-based indexing (0 or 1 usually). also useful when partitioning the matrix.
size_t	elemsize	size of the elements of the matrix

**struct starpu\_coo\_interface**

COO Matrices

## Data Fields

uint32_t *	columns	column array of the matrix
uint32_t *	rows	row array of the matrix
uintptr_t	values	values of the matrix
uint32_t	nx	number of elements on the x-axis of the matrix
uint32_t	ny	number of elements on the y-axis of the matrix
uint32_t	n_values	number of values registered in the matrix
size_t	elemsize	size of the elements of the matrix

**17.8.3 Macro Definition Documentation****#define STARPU\_VARIABLE\_GET\_PTR( *interface* )**Return a pointer to the variable designated by *interface*.**#define STARPU\_VARIABLE\_GET\_ELEMSIZE( *interface* )**Return the size of the variable designated by *interface*.**#define STARPU\_VARIABLE\_GET\_DEV\_HANDLE( *interface* )**Return a device handle for the variable designated by *interface*, to be used on OpenCL. The offset documented below has to be used in addition to this.**#define STARPU\_VARIABLE\_GET\_OFFSET**Return the offset in the variable designated by *interface*, to be used with the device handle.**#define STARPU\_VECTOR\_GET\_PTR( *interface* )**Return a pointer to the array designated by *interface*, valid on CPUs and CUDA only. For OpenCL, the device handle and offset need to be used instead.

```
#define STARPU_VECTOR_GET_DEV_HANDLE( interface )
```

Return a device handle for the array designated by *interface*, to be used on OpenCL. The offset documented below has to be used in addition to this.

```
#define STARPU_VECTOR_GET_OFFSET( interface )
```

Return the offset in the array designated by *interface*, to be used with the device handle.

```
#define STARPU_VECTOR_GET_NX( interface )
```

Return the number of elements registered into the array designated by *interface*.

```
#define STARPU_VECTOR_GET_ELEMSIZE( interface )
```

Return the size of each element of the array designated by *interface*.

```
#define STARPU_MATRIX_GET_PTR( interface )
```

Return a pointer to the matrix designated by *interface*, valid on CPUs and CUDA devices only. For OpenCL devices, the device handle and offset need to be used instead.

```
#define STARPU_MATRIX_GET_DEV_HANDLE( interface )
```

Return a device handle for the matrix designated by *interface*, to be used on OpenCL. The offset documented below has to be used in addition to this.

```
#define STARPU_MATRIX_GET_OFFSET( interface )
```

Return the offset in the matrix designated by *interface*, to be used with the device handle.

```
#define STARPU_MATRIX_GET_NX( interface )
```

Return the number of elements on the x-axis of the matrix designated by *interface*.

```
#define STARPU_MATRIX_GET_NY( interface )
```

Return the number of elements on the y-axis of the matrix designated by *interface*.

```
#define STARPU_MATRIX_GET_LD( interface )
```

Return the number of elements between each row of the matrix designated by *interface*. May be equal to nx when there is no padding.

```
#define STARPU_MATRIX_GET_ELEMSIZE( interface )
```

Return the size of the elements registered into the matrix designated by *interface*.

```
#define STARPU_BLOCK_GET_PTR( interface )
```

Return a pointer to the block designated by *interface*.

```
#define STARPU_BLOCK_GET_DEV_HANDLE( interface )
```

Return a device handle for the block designated by *interface*, to be used on OpenCL. The offset documented below has to be used in addition to this.

```
#define STARPU_BLOCK_GET_OFFSET( interface )
```

Return the offset in the block designated by *interface*, to be used with the device handle.

**#define STARPU\_BLOCK\_GET\_NX( *interface* )**

Return the number of elements on the x-axis of the block designated by *interface*.

**#define STARPU\_BLOCK\_GET\_NY( *interface* )**

Return the number of elements on the y-axis of the block designated by *interface*.

**#define STARPU\_BLOCK\_GET\_NZ( *interface* )**

Return the number of elements on the z-axis of the block designated by *interface*.

**#define STARPU\_BLOCK\_GET\_LDY( *interface* )**

Return the number of elements between each row of the block designated by *interface*. May be equal to nx when there is no padding.

**#define STARPU\_BLOCK\_GET\_LDZ( *interface* )**

Return the number of elements between each z plane of the block designated by *interface*. May be equal to nx\*ny when there is no padding.

**#define STARPU\_BLOCK\_GET\_ELEMSIZE( *interface* )**

Return the size of the elements of the block designated by *interface*.

**#define STARPU\_BCSR\_GET\_NNZ( *interface* )**

Return the number of non-zero values in the matrix designated by *interface*.

**#define STARPU\_BCSR\_GET\_NZVAL( *interface* )**

Return a pointer to the non-zero values of the matrix designated by *interface*.

**#define STARPU\_BCSR\_GET\_NZVAL\_DEV\_HANDLE( *interface* )**

Return a device handle for the array of non-zero values in the matrix designated by *interface*. The offset documented below has to be used in addition to this.

**#define STARPU\_BCSR\_GET\_COLIND( *interface* )**

Return a pointer to the column index of the matrix designated by *interface*.

**#define STARPU\_BCSR\_GET\_COLIND\_DEV\_HANDLE( *interface* )**

Return a device handle for the column index of the matrix designated by *interface*. The offset documented below has to be used in addition to this.

**#define STARPU\_BCSR\_GET\_ROWPTR( *interface* )**

Return a pointer to the row pointer array of the matrix designated by *interface*.

**#define STARPU\_BCSR\_GET\_ROWPTR\_DEV\_HANDLE( *interface* )**

Return a device handle for the row pointer array of the matrix designated by *interface*. The offset documented below has to be used in addition to this.

**#define STARPU\_BCSR\_GET\_OFFSET**

Return the offset in the arrays (colind, rowptr, nzval) of the matrix designated by *interface*, to be used with the device handles.

**#define STARPU\_CSR\_GET\_NNZ( *interface* )**

Return the number of non-zero values in the matrix designated by *interface*.

**#define STARPU\_CSR\_GET\_NROW( *interface* )**

Return the size of the row pointer array of the matrix designated by *interface*.

**#define STARPU\_CSR\_GET\_NZVAL( *interface* )**

Return a pointer to the non-zero values of the matrix designated by *interface*.

**#define STARPU\_CSR\_GET\_NZVAL\_DEV\_HANDLE( *interface* )**

Return a device handle for the array of non-zero values in the matrix designated by *interface*. The offset documented below has to be used in addition to this.

**#define STARPU\_CSR\_GET\_COLIND( *interface* )**

Return a pointer to the column index of the matrix designated by *interface*.

**#define STARPU\_CSR\_GET\_COLIND\_DEV\_HANDLE( *interface* )**

Return a device handle for the column index of the matrix designated by *interface*. The offset documented below has to be used in addition to this.

**#define STARPU\_CSR\_GET\_ROWPTR( *interface* )**

Return a pointer to the row pointer array of the matrix designated by *interface*.

**#define STARPU\_CSR\_GET\_ROWPTR\_DEV\_HANDLE( *interface* )**

Return a device handle for the row pointer array of the matrix designated by *interface*. The offset documented below has to be used in addition to this.

**#define STARPU\_CSR\_GET\_OFFSET**

Return the offset in the arrays (colind, rowptr, nzval) of the matrix designated by *interface*, to be used with the device handles.

**#define STARPU\_CSR\_GET\_FIRSTENTRY( *interface* )**

Return the index at which all arrays (the column indexes, the row pointers...) of the *interface* start.

**#define STARPU\_CSR\_GET\_ELEMSIZE( *interface* )**

Return the size of the elements registered into the matrix designated by *interface*.

**#define STARPU\_COO\_GET\_COLUMNS( *interface* )**

Return a pointer to the column array of the matrix designated by *interface*.

**#define STARPU\_COO\_GET\_COLUMNS\_DEV\_HANDLE( *interface* )**

Return a device handle for the column array of the matrix designated by *interface*, to be used on OpenCL. The offset documented below has to be used in addition to this.

**#define STARPU\_COO\_GET\_ROWS( *interface* )**

Return a pointer to the rows array of the matrix designated by *interface*.

**#define STARPU\_COO\_GET\_ROWS\_DEV\_HANDLE( *interface* )**

Return a device handle for the row array of the matrix designated by *interface*, to be used on OpenCL. The offset documented below has to be used in addition to this.

**#define STARPU\_COO\_GET\_VALUES( *interface* )**

Return a pointer to the values array of the matrix designated by *interface*.

**#define STARPU\_COO\_GET\_VALUES\_DEV\_HANDLE( *interface* )**

Return a device handle for the value array of the matrix designated by *interface*, to be used on OpenCL. The offset documented below has to be used in addition to this.

**#define STARPU\_COO\_GET\_OFFSET**

Return the offset in the arrays of the COO matrix designated by *interface*.

**#define STARPU\_COO\_GET\_NX( *interface* )**

Return the number of elements on the x-axis of the matrix designated by *interface*.

**#define STARPU\_COO\_GET\_NY( *interface* )**

Return the number of elements on the y-axis of the matrix designated by *interface*.

**#define STARPU\_COO\_GET\_NVALUES( *interface* )**

Return the number of values registered in the matrix designated by *interface*.

**#define STARPU\_COO\_GET\_ELEMSIZE( *interface* )**

Return the size of the elements registered into the matrix designated by *interface*.

#### 17.8.4 Enumeration Type Documentation

**enum starpu\_data\_interface\_id**

Identifier for all predefined StarPU data interfaces

Enumerator

**STARPU\_UNKNOWN\_INTERFACE\_ID** Unknown interface

**STARPU\_MATRIX\_INTERFACE\_ID** Identifier for the matrix data interface

**STARPU\_BLOCK\_INTERFACE\_ID** Identifier for block data interface

**STARPU\_VECTOR\_INTERFACE\_ID** Identifier for the vector data interface

**STARPU\_CSR\_INTERFACE\_ID** Identifier for the csr data interface

**STARPU\_BCSR\_INTERFACE\_ID** Identifier for the bcsr data interface

**STARPU\_VARIABLE\_INTERFACE\_ID** Identifier for the variable data interface

**STARPU\_VOID\_INTERFACE\_ID** Identifier for the void data interface

**STARPU\_MULTIFORMAT\_INTERFACE\_ID** Identifier for the multiformat data interface

**STARPU\_COO\_INTERFACE\_ID** Identifier for the coo data interface

**STARPU\_MAX\_INTERFACE\_ID** Maximum number of data interfaces

### 17.8.5 Function Documentation

**void starpu\_void\_data\_register ( starpu\_data\_handle\_t \* handle )**

Register a void interface. There is no data really associated to that interface, but it may be used as a synchronization mechanism. It also permits to express an abstract piece of data that is managed by the application internally: this makes it possible to forbid the concurrent execution of different tasks accessing the same void data in read-write concurrently.

**void starpu\_variable\_data\_register ( starpu\_data\_handle\_t \* handle, unsigned home\_node, uintptr\_t ptr, size\_t size )**

Register the `size` byte element pointed to by `ptr`, which is typically a scalar, and initialize `handle` to represent this data item.

Here an example of how to use the function.

```
float var;
starpu_data_handle_t var_handle;
starpu_variable_data_register(&var.handle, 0, (uintptr_t)&var, sizeof(var));
```

**void starpu\_vector\_data\_register ( starpu\_data\_handle\_t \* handle, unsigned home\_node, uintptr\_t ptr, uint32\_t nx, size\_t elemsize )**

Register the `nx` `elemsize`-byte elements pointed to by `ptr` and initialize `handle` to represent it.

Here an example of how to use the function.

```
float vector[NX];
starpu_data_handle_t vector_handle;
starpu_vector_data_register(&vector.handle, 0, (uintptr_t)vector, NX, sizeof(
    vector[0]));
```

**void starpu\_matrix\_data\_register ( starpu\_data\_handle\_t \* handle, unsigned home\_node, uintptr\_t ptr, uint32\_t ld, uint32\_t nx, uint32\_t ny, size\_t elemsize )**

Register the `nx` x `ny` 2D matrix of `elemsize`-byte elements pointed by `ptr` and initialize `handle` to represent it. `ld` specifies the number of elements between rows. a value greater than `nx` adds padding, which can be useful for alignment purposes.

Here an example of how to use the function.

```
float *matrix;
starpu_data_handle_t matrix_handle;
matrix = (float*)malloc(width * height * sizeof(float));
starpu_matrix_data_register(&matrix.handle, 0, (uintptr_t)matrix, width, width,
    height, sizeof(float));
```

**void starpu\_block\_data\_register ( starpu\_data\_handle\_t \* handle, unsigned home\_node, uintptr\_t ptr, uint32\_t ldy, uint32\_t ldz, uint32\_t nx, uint32\_t ny, uint32\_t nz, size\_t elemsize )**

Register the `nx` x `ny` x `nz` 3D matrix of `elemsize` byte elements pointed by `ptr` and initialize `handle` to represent it. Again, `ldy` and `ldz` specify the number of elements between rows and between z planes.

Here an example of how to use the function.

```
float *block;
starpu_data_handle_t block_handle;
block = (float*)malloc(nx*ny*nz*sizeof(float));
starpu_block_data_register(&block.handle, 0, (uintptr_t)block, nx, nx*ny, nx, ny,
    nz, sizeof(float));
```

```
void starpu_bcsr_data_register ( starpu_data_handle_t * handle, unsigned home_node, uint32_t nnz,
uint32_t nrow, uintptr_t nzval, uint32_t * colind, uint32_t * rowptr, uint32_t firstentry, uint32_t r, uint32_t c,
size_t elemsize )
```

This variant of [starpu\\_data\\_register\(\)](#) uses the BCSR (Blocked Compressed Sparse Row Representation) sparse matrix interface. Register the sparse matrix made of `nnz` non-zero blocks of elements of size `elemsize` stored in `nzval` and initializes `handle` to represent it. Blocks have size `r * c`. `nrow` is the number of rows (in terms of blocks), `colind[i]` is the block-column index for block `i` in `nzval`, `rowptr[i]` is the block-index (in `nzval`) of the first block of row `i`. `firstentry` is the index of the first entry of the given arrays (usually 0 or 1).

```
void starpu_csr_data_register ( starpu_data_handle_t * handle, unsigned home_node, uint32_t nnz,
uint32_t nrow, uintptr_t nzval, uint32_t * colind, uint32_t * rowptr, uint32_t firstentry, size_t elemsize )
```

This variant of [starpu\\_data\\_register\(\)](#) uses the CSR (Compressed Sparse Row Representation) sparse matrix interface. TODO

```
void starpu_coo_data_register ( starpu_data_handle_t * handleptr, unsigned home_node, uint32_t nx,
uint32_t ny, uint32_t n_values, uint32_t * columns, uint32_t * rows, uintptr_t values, size_t elemsize )
```

Register the `nx x ny` 2D matrix given in the COO format, using the `columns`, `rows`, `values` arrays, which must have `n_values` elements of size `elemsize`. Initialize `handleptr`.

```
void * starpu_data_get_interface_on_node ( starpu_data_handle_t handle, unsigned memory_node )
```

Return the interface associated with `handle` on `memory_node`.

```
void * starpu_data_handle_to_pointer ( starpu_data_handle_t handle, unsigned node )
```

Return the pointer associated with `handle` on node `node` or `NULL` if `handle`'s interface does not support this operation or data for this `handle` is not allocated on that `node`.

```
void * starpu_data_get_local_ptr ( starpu_data_handle_t handle )
```

Return the local pointer associated with `handle` or `NULL` if `handle`'s interface does not have data allocated locally

```
enum starpu_data_interface_id starpu_data_get_interface_id ( starpu_data_handle_t handle )
```

Return the unique identifier of the interface associated with the given `handle`.

```
size_t starpu_data_get_size ( starpu_data_handle_t handle )
```

Return the size of the data associated with `handle`.

```
int starpu_data_pack ( starpu_data_handle_t handle, void ** ptr, starpu_ssize_t * count )
```

Execute the packing operation of the interface of the data registered at `handle` (see [starpu\\_data\\_interface\\_ops](#)). This packing operation must allocate a buffer large enough at `ptr` and copy into the newly allocated buffer the data associated to `handle`. `count` will be set to the size of the allocated buffer. If `ptr` is `NULL`, the function should not copy the data in the buffer but just set `count` to the size of the buffer which would have been allocated. The special value `-1` indicates the size is yet unknown.

```
int starpu_data_unpack ( starpu_data_handle_t handle, void * ptr, size_t count )
```

Unpack in `handle` the data located at `ptr` of size `count` as described by the interface of the data. The interface registered at `handle` must define a unpacking operation (see [starpu\\_data\\_interface\\_ops](#)). The memory at the address `ptr` is freed after calling the data unpacking operation.

**size\_t starpu\_variable\_get\_elemsize ( starpu\_data\_handle\_t handle )**

Return the size of the variable designated by handle.

**uintptr\_t starpu\_variable\_get\_local\_ptr ( starpu\_data\_handle\_t handle )**

Return a pointer to the variable designated by handle.

**uint32\_t starpu\_vector\_get\_nx ( starpu\_data\_handle\_t handle )**

Return the number of elements registered into the array designated by handle.

**size\_t starpu\_vector\_get\_elemsize ( starpu\_data\_handle\_t handle )**

Return the size of each element of the array designated by handle.

**uintptr\_t starpu\_vector\_get\_local\_ptr ( starpu\_data\_handle\_t handle )**

Return the local pointer associated with handle.

**uint32\_t starpu\_matrix\_get\_nx ( starpu\_data\_handle\_t handle )**

Return the number of elements on the x-axis of the matrix designated by handle.

**uint32\_t starpu\_matrix\_get\_ny ( starpu\_data\_handle\_t handle )**

Return the number of elements on the y-axis of the matrix designated by handle.

**uint32\_t starpu\_matrix\_get\_local\_id ( starpu\_data\_handle\_t handle )**

Return the number of elements between each row of the matrix designated by handle. Maybe be equal to nx when there is no padding.

**uintptr\_t starpu\_matrix\_get\_local\_ptr ( starpu\_data\_handle\_t handle )**

Return the local pointer associated with handle.

**size\_t starpu\_matrix\_get\_elemsize ( starpu\_data\_handle\_t handle )**

Return the size of the elements registered into the matrix designated by handle.

**uint32\_t starpu\_block\_get\_nx ( starpu\_data\_handle\_t handle )**

Return the number of elements on the x-axis of the block designated by handle.

**uint32\_t starpu\_block\_get\_ny ( starpu\_data\_handle\_t handle )**

Return the number of elements on the y-axis of the block designated by handle.

**uint32\_t starpu\_block\_get\_nz ( starpu\_data\_handle\_t handle )**

Return the number of elements on the z-axis of the block designated by handle.

**uint32\_t starpu\_block\_get\_local\_idy ( starpu\_data\_handle\_t handle )**

Return the number of elements between each row of the block designated by handle, in the format of the current memory node.

**uint32\_t starpu\_block\_get\_local\_ldz ( starpu\_data\_handle\_t handle )**

Return the number of elements between each z plane of the block designated by `handle`, in the format of the current memory node.

**uintptr\_t starpu\_block\_get\_local\_ptr ( starpu\_data\_handle\_t handle )**

Return the local pointer associated with `handle`.

**size\_t starpu\_block\_get\_elemsize ( starpu\_data\_handle\_t handle )**

Return the size of the elements of the block designated by `handle`.

**uint32\_t starpu\_bcsr\_get\_nnz ( starpu\_data\_handle\_t handle )**

Return the number of non-zero elements in the matrix designated by `handle`.

**uint32\_t starpu\_bcsr\_get\_nrow ( starpu\_data\_handle\_t handle )**

Return the number of rows (in terms of blocks of size r\*c) in the matrix designated by `handle`.

**uint32\_t starpu\_bcsr\_get\_firstentry ( starpu\_data\_handle\_t handle )**

Return the index at which all arrays (the column indexes, the row pointers...) of the matrix designated by `handle`.

**uintptr\_t starpu\_bcsr\_get\_local\_nzval ( starpu\_data\_handle\_t handle )**

Return a pointer to the non-zero values of the matrix designated by `handle`.

**uint32\_t \* starpu\_bcsr\_get\_local\_colind ( starpu\_data\_handle\_t handle )**

Return a pointer to the column index, which holds the positions of the non-zero entries in the matrix designated by `handle`.

**uint32\_t \* starpu\_bcsr\_get\_local\_rowptr ( starpu\_data\_handle\_t handle )**

Return the row pointer array of the matrix designated by `handle`.

**uint32\_t starpu\_bcsr\_get\_r ( starpu\_data\_handle\_t handle )**

Return the number of rows in a block.

**uint32\_t starpu\_bcsr\_get\_c ( starpu\_data\_handle\_t handle )**

Return the number of columns in a block.

**size\_t starpu\_bcsr\_get\_elemsize ( starpu\_data\_handle\_t handle )**

Return the size of the elements in the matrix designated by `handle`.

**uint32\_t starpu\_csr\_get\_nnz ( starpu\_data\_handle\_t handle )**

Return the number of non-zero values in the matrix designated by `handle`.

**uint32\_t starpu\_csr\_get\_nrow ( starpu\_data\_handle\_t handle )**

Return the size of the row pointer array of the matrix designated by `handle`.

**uint32\_t starpu\_csr\_get\_firstentry ( starpu\_data\_handle\_t handle )**

Return the index at which all arrays (the column indexes, the row pointers...) of the matrix designated by `handle`.

**uintptr\_t starpu\_csr\_get\_local\_nzval ( starpu\_data\_handle\_t handle )**

Return a local pointer to the non-zero values of the matrix designated by `handle`.

**uint32\_t \* starpu\_csr\_get\_local\_colind ( starpu\_data\_handle\_t handle )**

Return a local pointer to the column index of the matrix designated by `handle`.

**uint32\_t \* starpu\_csr\_get\_local\_rowptr ( starpu\_data\_handle\_t handle )**

Return a local pointer to the row pointer array of the matrix designated by `handle`.

**size\_t starpu\_csr\_get\_elemsize ( starpu\_data\_handle\_t handle )**

Return the size of the elements registered into the matrix designated by `handle`.

**uintptr\_t starpu\_malloc\_on\_node ( unsigned dst\_node, size\_t size )**

Allocate `size` bytes on node `dst_node`. This returns 0 if allocation failed, the allocation method should then return -ENOMEM as allocated size.

**void starpu\_free\_on\_node ( unsigned dst\_node, uintptr\_t addr, size\_t size )**

Free `addr` of `size` bytes on node `dst_node`.

**int starpu\_interface\_copy ( uintptr\_t src, size\_t src\_offset, unsigned src\_node, uintptr\_t dst, size\_t dst\_offset, unsigned dst\_node, size\_t size, void \* async\_data )**

Copy `size` bytes from byte offset `src_offset` of `src` on `src_node` to byte offset `dst_offset` of `dst` on `dst_node`. This is to be used in the `any_to_any()` copy method, which is provided with the `async_data` to be passed to `starpu_interface_copy()`. This returns -EAGAIN if the transfer is still ongoing, or 0 if the transfer is already completed.

**uint32\_t starpu\_hash\_crc32c\_be\_n ( const void \* input, size\_t n, uint32\_t inputcrc )**

Compute the CRC of a byte buffer seeded by the `inputcrc` *current state*. The return value should be considered as the new *current state* for future CRC computation. This is used for computing data size footprint.

**uint32\_t starpu\_hash\_crc32c\_be ( uint32\_t input, uint32\_t inputcrc )**

Compute the CRC of a 32bit number seeded by the `inputcrc` *current state*. The return value should be considered as the new *current state* for future CRC computation. This is used for computing data size footprint.

**uint32\_t starpu\_hash\_crc32c\_string ( const char \* str, uint32\_t inputcrc )**

Compute the CRC of a string seeded by the `inputcrc` *current state*. The return value should be considered as the new *current state* for future CRC computation. This is used for computing data size footprint.

**int starpu\_data\_interface\_get\_next\_id ( void )**

Return the next available id for a newly created data interface ([Defining A New Data Interface](#)).

## 17.9 Data Partition

### Data Structures

- struct `starpu_data_filter`

## Basic API

- void `starpu_data_partition` (`starpu_data_handle_t` initial\_handle, struct `starpu_data_filter` \*f)
- void `starpu_data_unpartition` (`starpu_data_handle_t` root\_data, unsigned gathering\_node)
- int `starpu_data_get_nb_children` (`starpu_data_handle_t` handle)
- `starpu_data_handle_t starpu_data_get_child` (`starpu_data_handle_t` handle, unsigned i)
- `starpu_data_handle_t starpu_data_get_sub_data` (`starpu_data_handle_t` root\_data, unsigned depth,...)
- `starpu_data_handle_t starpu_data_vget_sub_data` (`starpu_data_handle_t` root\_data, unsigned depth, va\_list pa)
- void `starpu_data_map_filters` (`starpu_data_handle_t` root\_data, unsigned nfilters,...)
- void `starpu_data_vmap_filters` (`starpu_data_handle_t` root\_data, unsigned nfilters, va\_list pa)

## Predefined Vector Filter Functions

This section gives a partial list of the predefined partitioning functions for vector data. Examples on how to use them are shown in [Partitioning Data](#). The complete list can be found in the file `starpu_data_filters.h`.

- void `starpu_vector_filter_block` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_vector_filter_block_shadow` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_vector_filter_list` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_vector_filter_divide_in_2` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)

## Predefined Matrix Filter Functions

This section gives a partial list of the predefined partitioning functions for matrix data. Examples on how to use them are shown in [Partitioning Data](#). The complete list can be found in the file `starpu_data_filters.h`.

- void `starpu_matrix_filter_block` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_matrix_filter_block_shadow` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_matrix_filter_vertical_block` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_matrix_filter_vertical_block_shadow` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)

## Predefined Block Filter Functions

This section gives a partial list of the predefined partitioning functions for block data. Examples on how to use them are shown in [Partitioning Data](#). The complete list can be found in the file `starpu_data_filters.h`. A usage example is available in `examples/filters/shadow3d.c`

- void `starpu_block_filter_block` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_block_filter_block_shadow` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_block_filter_vertical_block` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_block_filter_vertical_block_shadow` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_block_filter_depth_block` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_block_filter_depth_block_shadow` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)

## Predefined BCSR Filter Functions

This section gives a partial list of the predefined partitioning functions for BCSR data. Examples on how to use them are shown in [Partitioning Data](#). The complete list can be found in the file `starpu_data_filters.h`.

- void `starpu_bcsr_filter_canonical_block` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_csr_filter_vertical_block` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)

### 17.9.1 Detailed Description

### 17.9.2 Data Structure Documentation

#### `struct starpu_data_filter`

The filter structure describes a data partitioning operation, to be given to the `starpu_data_partition()` function.

#### Data Fields

- void(\* `filter_func` )(void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*, unsigned id, unsigned nparts)
- unsigned `nchildren`
- unsigned(\* `get_nchildren` )(struct `starpu_data_filter` \*, `starpu_data_handle_t` initial\_handle)
- struct  
  `starpu_data_interface_ops` \*(\* `get_child_ops` )(struct `starpu_data_filter` \*, unsigned id)
- unsigned `filter_arg`
- void \* `filter_arg_ptr`

#### Field Documentation

**`starpu_data_filter::filter_func`** This function fills the child\_interface structure with interface information for the id-th child of the parent father\_interface (among nparts).

**`starpu_data_filter::nchildren`** This is the number of parts to partition the data into.

**`starpu_data_filter::get_nchildren`** This returns the number of children. This can be used instead of nchildren when the number of children depends on the actual data (e.g. the number of blocks in a sparse matrix).

**`starpu_data_filter::get_child_ops`** In case the resulting children use a different data interface, this function returns which interface is used by child number id.

**`starpu_data_filter::filter_arg`** Allow to define an additional parameter for the filter function.

**`starpu_data_filter::filter_arg_ptr`** Allow to define an additional pointer parameter for the filter function, such as the sizes of the different parts.

### 17.9.3 Function Documentation

#### `void starpu_data_partition ( starpu_data_handle_t initial_handle, struct starpu_data_filter * f )`

This requests partitioning one StarPU data initial\_handle into several subdata according to the filter f.

Here an example of how to use the function.

```
struct starpu_data_filter f = {
    .filter_func = starpu_matrix.filter.block,
    .nchildren = nslicesx,
    .get_nchildren = NULL,
    .get_child_ops = NULL
};  
starpu_data_partition(A_handle, &f);
```

**void starpu\_data\_unpartition ( starpu\_data\_handle\_t root\_data, unsigned gathering\_node )**

This unapplies one filter, thus unpartitioning the data. The pieces of data are collected back into one big piece in the `gathering_node` (usually 0). Tasks working on the partitioned data must be already finished when calling `starpu_data_unpartition()`.

Here an example of how to use the function.

```
starpu_data_unpartition(A_handle, 0);
```

**int starpu\_data\_get\_nb\_children ( starpu\_data\_handle\_t handle )**

This function returns the number of children.

**starpu\_data\_handle\_t starpu\_data\_get\_child ( starpu\_data\_handle\_t handle, unsigned i )**

Return the `i`th child of the given `handle`, which must have been partitionned beforehand.

**starpu\_data\_handle\_t starpu\_data\_get\_sub\_data ( starpu\_data\_handle\_t root\_data, unsigned depth, ... )**

After partitioning a StarPU data by applying a filter, `starpu_data_get_sub_data()` can be used to get handles for each of the data portions. `root_data` is the parent data that was partitioned. `depth` is the number of filters to traverse (in case several filters have been applied, to e.g. partition in row blocks, and then in column blocks), and the subsequent parameters are the indexes. The function returns a handle to the subdata.

Here an example of how to use the function.

```
h = starpu_data_get_sub_data(A_handle, 1, taskx);
```

**starpu\_data\_handle\_t starpu\_data\_vget\_sub\_data ( starpu\_data\_handle\_t root\_data, unsigned depth, va\_list pa )**

This function is similar to `starpu_data_get_sub_data()` but uses a `va_list` for the parameter list.

**void starpu\_data\_map\_filters ( starpu\_data\_handle\_t root\_data, unsigned nfilters, ... )**

Applies `nfilters` filters to the handle designated by `root_data` recursively. `nfilters` pointers to variables of the type `starpu_data_filter` should be given.

**void starpu\_data\_vmap\_filters ( starpu\_data\_handle\_t root\_data, unsigned nfilters, va\_list pa )**

Applies `nfilters` filters to the handle designated by `root_data` recursively. It uses a `va_list` of pointers to variables of the type `starpu_data_filter`.

**void starpu\_vector\_filter\_block ( void \* father\_interface, void \* child\_interface, struct starpu\_data\_filter \* f, unsigned id, unsigned nparts )**

Return in `child_interface` the `id`th element of the vector represented by `father_interface` once partitioned in `nparts` chunks of equal size.

**void starpu\_vector\_filter\_block\_shadow ( void \* father\_interface, void \* child\_interface, struct starpu\_data\_filter \* f, unsigned id, unsigned nparts )**

Return in `child_interface` the `id`th element of the vector represented by `father_interface` once partitioned in `nparts` chunks of equal size with a shadow border `filter_arg_ptr`, thus getting a vector of size  $(n-2*\text{shadow})/\text{nparts}+2*\text{shadow}$ . The `filter_arg_ptr` field of `f` must be the shadow size casted into `void*`. **IMPORTANT:** This can only be used for read-only access, as no coherency is enforced for the shadowed parts. An usage example is available in `examples/filters/shadow.c`

```
void starpu_vector_filter_list ( void * father_interface, void * child_interface, struct starpu_data_filter * f,
unsigned id, unsigned nparts )
```

Return in `child_interface` the `id` th element of the vector represented by `father_interface` once partitioned into `nparts` chunks according to the `filter_arg_ptr` field of `f`. The `filter_arg_ptr` field must point to an array of `nparts` `uint32_t` elements, each of which specifies the number of elements in each chunk of the partition.

```
void starpu_vector_filter_divide_in_2 ( void * father_interface, void * child_interface, struct
starpu_data_filter * f, unsigned id, unsigned nparts )
```

Return in `child_interface` the `id` th element of the vector represented by `father_interface` once partitioned in 2 chunks of equal size, ignoring `nparts`. Thus, `id` must be 0 or 1.

```
void starpu_matrix_filter_block ( void * father_interface, void * child_interface, struct starpu_data_filter * f,
unsigned id, unsigned nparts )
```

This partitions a dense Matrix along the x dimension, thus getting  $(x/nparts, y)$  matrices. If `nparts` does not divide `x`, the last submatrix contains the remainder.

```
void starpu_matrix_filter_block_shadow ( void * father_interface, void * child_interface, struct
starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a dense Matrix along the x dimension, with a shadow border `filter_arg_ptr`, thus getting  $((x-2*shadow)/nparts +2*shadow, y)$  matrices. If `nparts` does not divide `x-2*shadow`, the last submatrix contains the remainder. **IMPORTANT:** This can only be used for read-only access, as no coherency is enforced for the shadowed parts. A usage example is available in `examples/filters/shadow2d.c`

```
void starpu_matrix_filter_vertical_block ( void * father_interface, void * child_interface, struct
starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a dense Matrix along the y dimension, thus getting  $(x,y/nparts)$  matrices. If `nparts` does not divide `y`, the last submatrix contains the remainder.

```
void starpu_matrix_filter_vertical_block_shadow ( void * father_interface, void * child_interface, struct
starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a dense Matrix along the y dimension, with a shadow border `filter_arg_ptr`, thus getting  $(x,(y-2*shadow)/nparts +2*shadow)$  matrices. If `nparts` does not divide `y-2*shadow`, the last submatrix contains the remainder. **IMPORTANT:** This can only be used for read-only access, as no coherency is enforced for the shadowed parts. A usage example is available in `examples/filters/shadow2d.c`

```
void starpu_block_filter_block ( void * father_interface, void * child_interface, struct starpu_data_filter * f,
unsigned id, unsigned nparts )
```

This partitions a block along the X dimension, thus getting  $(x/nparts, y, z)$  3D matrices. If `nparts` does not divide `x`, the last submatrix contains the remainder.

```
void starpu_block_filter_block_shadow ( void * father_interface, void * child_interface, struct
starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a block along the X dimension, with a shadow border `filter_arg_ptr`, thus getting  $((x-2*shadow)/nparts +2*shadow, y, z)$  blocks. If `nparts` does not divide `x`, the last submatrix contains the remainder. **IMPORTANT:** This can only be used for read-only access, as no coherency is enforced for the shadowed parts.

```
void starpu_block_filter_vertical_block ( void * father_interface, void * child_interface, struct starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a block along the Y dimension, thus getting (x,y/nparts,z) blocks. If nparts does not divide y, the last submatrix contains the remainder.

```
void starpu_block_filter_vertical_block_shadow ( void * father_interface, void * child_interface, struct starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a block along the Y dimension, with a shadow border filter\_arg\_ptr, thus getting (x,(y-2\*shadow)/nparts +2\*shadow,z) 3D matrices. If nparts does not divide y, the last submatrix contains the remainder. **IMPORTANT:** This can only be used for read-only access, as no coherency is enforced for the shadowed parts.

```
void starpu_block_filter_depth_block ( void * father_interface, void * child_interface, struct starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a block along the Z dimension, thus getting (x,y,z/nparts) blocks. If nparts does not divide z, the last submatrix contains the remainder.

```
void starpu_block_filter_depth_block_shadow ( void * father_interface, void * child_interface, struct starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a block along the Z dimension, with a shadow border filter\_arg\_ptr, thus getting (x,y,(z-2\*shadow)/nparts +2\*shadow) blocks. If nparts does not divide z, the last submatrix contains the remainder. **IMPORTANT:** This can only be used for read-only access, as no coherency is enforced for the shadowed parts.

```
void starpu_bcsr_filter_canonical_block ( void * father_interface, void * child_interface, struct starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a block-sparse matrix into dense matrices.

```
void starpu_csr_filter_vertical_block ( void * father_interface, void * child_interface, struct starpu_data_filter * f, unsigned id, unsigned nparts )
```

This partitions a block-sparse matrix into vertical block-sparse matrices.

## 17.10 Multiformat Data Interface

### Data Structures

- struct [starpu\\_multiformat\\_data\\_interface\\_ops](#)
- struct [starpu\\_multiformat\\_interface](#)

### Macros

- #define [STARPU\\_MULTIFORMAT\\_GET\\_CPU\\_PTR](#)(interface)
- #define [STARPU\\_MULTIFORMAT\\_GET\\_CUDA\\_PTR](#)(interface)
- #define [STARPU\\_MULTIFORMAT\\_GET\\_OPENCL\\_PTR](#)(interface)
- #define [STARPU\\_MULTIFORMAT\\_GET\\_NX](#)(interface)

### Functions

- void [starpu\\_multiformat\\_data\\_register](#) (starpu\_data\_handle\_t \*handle, unsigned home\_node, void \*ptr, uint32\_t nobjects, struct [starpu\\_multiformat\\_data\\_interface\\_ops](#) \*format\_ops)

**17.10.1 Detailed Description****17.10.2 Data Structure Documentation****struct starpu\_multiformat\_data\_interface\_ops**

The different fields are:

**Data Fields**

<code>size_t</code>	<code>cpu_elemsize</code>	the size of each element on CPUs
<code>size_t</code>	<code>opencl_elemsize</code>	the size of each element on OpenCL devices
<code>struct starpu_codelet *</code>	<code>cpu_to_opencl_cl</code>	pointer to a codelet which converts from CPU to OpenCL
<code>struct starpu_codelet *</code>	<code>opencl_to_cpu_cl</code>	pointer to a codelet which converts from OpenCL to CPU
<code>size_t</code>	<code>cuda_elemsize</code>	the size of each element on CUDA devices
<code>struct starpu_codelet *</code>	<code>cpu_to_cuda_cl</code>	pointer to a codelet which converts from CPU to CUDA
<code>struct starpu_codelet *</code>	<code>cuda_to_cpu_cl</code>	pointer to a codelet which converts from CUDA to CPU

**struct starpu\_multiformat\_interface**

todo

**Data Fields**

<code>void *</code>	<code>cpu_ptr</code>	
<code>void *</code>	<code>cuda_ptr</code>	
<code>void *</code>	<code>opencl_ptr</code>	
<code>uint32_t</code>	<code>nx</code>	
<code>struct starpu_multiformat_data_interface_ops *</code>	<code>ops</code>	

**17.10.3 Macro Definition Documentation****#define STARPU\_MULTIFORMAT\_GET\_CPU\_PTR( *interface* )**

returns the local pointer to the data with CPU format.

**#define STARPU\_MULTIFORMAT\_GET\_CUDA\_PTR( *interface* )**

returns the local pointer to the data with CUDA format.

**#define STARPU\_MULTIFORMAT\_GET\_OPENCL\_PTR( *interface* )**

returns the local pointer to the data with OpenCL format.

**#define STARPU\_MULTIFORMAT\_GET\_NX( *interface* )**

returns the number of elements in the data.

**17.10.4 Function Documentation**

```
void starpu_multiformat_data_register ( starpu_data_handle_t * handle, unsigned home_node, void * ptr,  
uint32_t nobjects, struct starpu_multiformat_data_interface_ops * format_ops )
```

Register a piece of data that can be represented in different ways, depending upon the processing unit that manipulates it. It allows the programmer, for instance, to use an array of structures when working on a CPU, and a structure of arrays when working on a GPU. *nobjects* is the number of elements in the data. *format\_ops* describes the format.

## 17.11 Codelet And Tasks

This section describes the interface to manipulate codelets and tasks.

### Data Structures

- struct `starpu_codelet`
- struct `starpu_data_descr`
- struct `starpu_task`

### Macros

- #define `STARPU_CPU`
- #define `STARPU_CUDA`
- #define `STARPU_OPENCL`
- #define `STARPU_MULTIPLE_CPU_IMPLEMENTATIONS`
- #define `STARPU_MULTIPLE_CUDA_IMPLEMENTATIONS`
- #define `STARPU_MULTIPLE_OPENCL_IMPLEMENTATIONS`
- #define `STARPU_NMAXBUFS`
- #define `STARPU_TASK_INITIALIZER`
- #define `STARPU_TASK_GET_HANDLE(task, i)`
- #define `STARPU_TASK_SET_HANDLE(task, handle, i)`
- #define `STARPU_CODELET_GET_MODE(codelet, i)`
- #define `STARPU_CODELET_SET_MODE(codelet, mode, i)`
- #define `STARPU_TASK_INVALID`

### Typedefs

- typedef void(\* `starpu_cpu_func_t` )(void \*\*, void \*)
- typedef void(\* `starpu_cuda_func_t` )(void \*\*, void \*)
- typedef void(\* `starpu_opencl_func_t` )(void \*\*, void \*)

### Enumerations

- enum `starpu_codelet_type` { `STARPU_SEQ`, `STARPU_SPMD`, `STARPU_FORKJOIN` }
- enum `starpu_task_status` {
 `STARPU_TASK_INVALID`, `STARPU_TASK_INVALID`, `STARPU_TASK_BLOCKED`, `STARPU_TASK_READY`,
 `STARPU_TASK_RUNNING`, `STARPU_TASK_FINISHED`, `STARPU_TASK_BLOCKED_ON_TAG`, `STARPU_TASK_BLOCKED_ON_TASK`,
 `STARPU_TASK_BLOCKED_ON_DATA` }

### Functions

- void `starpu_codelet_init` (struct `starpu_codelet` \*cl)
- void `starpu_task_init` (struct `starpu_task` \*task)
- struct `starpu_task` \* `starpu_task_create` (void)
- struct `starpu_task` \* `starpu_task_dup` (struct `starpu_task` \*task)
- void `starpu_task_clean` (struct `starpu_task` \*task)
- void `starpu_task_destroy` (struct `starpu_task` \*task)
- int `starpu_task_wait` (struct `starpu_task` \*task) `STARPU_WARN_UNUSED_RESULT`
- int `starpu_task_submit` (struct `starpu_task` \*task) `STARPU_WARN_UNUSED_RESULT`
- int `starpu_task_submit_to_ctx` (struct `starpu_task` \*task, unsigned `sched_ctx_id`)
- int `starpu_task_wait_for_all` (void)
- int `starpu_task_wait_for_all_in_ctx` (unsigned `sched_ctx_id`)
- int `starpu_task_nready` (void)

- int `starpu_task_nsubmitted` (void)
- struct `starpu_task` \* `starpu_task_get_current` (void)
- void `starpu_codelet_display_stats` (struct `starpu_codelet` \*cl)
- int `starpu_task_wait_for_no_ready` (void)
- void `starpu_task_set_implementation` (struct `starpu_task` \*task, unsigned impl)
- unsigned `starpu_task_get_implementation` (struct `starpu_task` \*task)
- void `starpu_create_sync_task` (`starpu_tag_t` sync.tag, unsigned ndeps, `starpu_tag_t` \*deps, void(\*callback)(void \*), void \*callback\_arg)

### 17.11.1 Detailed Description

This section describes the interface to manipulate codelets and tasks.

### 17.11.2 Data Structure Documentation

#### `struct starpu_codelet`

The codelet structure describes a kernel that is possibly implemented on various targets. For compatibility, make sure to initialize the whole structure to zero, either by using explicit memset, or the function `starpu_codelet_init()`, or by letting the compiler implicitly do it in e.g. static storage case.

#### Data Fields

- uint32\_t `where`
- int(\*) `can_execute` )(unsigned workerid, struct `starpu_task` \*task, unsigned nimpl)
- enum `starpu_codelet_type` type
- int `max_parallelism`
- `starpu_cpu_func_t` cpu\_func
- `starpu_cuda_func_t` cuda\_func
- `starpu_opencl_func_t` opencl\_func
- `starpu_cpu_func_t` cpu\_funcs [STARPU\_MAXIMPLEMENTATIONS]
- `starpu_cuda_func_t` cuda\_funcs [STARPU\_MAXIMPLEMENTATIONS]
- `starpu_opencl_func_t` opencl\_funcs [STARPU\_MAXIMPLEMENTATIONS]
- unsigned `nbuffers`
- enum `starpu_data_access_mode` modes [STARPU\_NMAXBUFS]
- enum `starpu_data_access_mode` \* dyn\_modes
- struct `starpu_perfmodel` \* model
- struct `starpu_perfmodel` \* power\_model
- unsigned long `per_worker_stats` [STARPU\_NMAXWORKERS]
- const char \* name

#### Field Documentation

**`starpu_codelet::where`** Optional field to indicate which types of processing units are able to execute the codelet. The different values `STARPU_CPU`, `STARPU_CUDA`, `STARPU_OPENCL` can be combined to specify on which types of processing units the codelet can be executed. `STARPU_CPU|STARPU_CUDA` for instance indicates that the codelet is implemented for both CPU cores and CUDA devices while `STARPU_OPENCL` indicates that it is only available on OpenCL devices. If the field is unset, its value will be automatically set based on the availability of the XXX\_funcs fields defined below.

**`starpu_codelet::can_execute`** Define a function which should return 1 if the worker designated by workerid can execute the nimplth implementation of the given task, 0 otherwise.

**`starpu_codelet::type`** Optional field to specify the type of the codelet. The default is `STARPU_SEQ`, i.e. usual sequential implementation. Other values (`STARPU_SPMD` or `STARPU_FORKJOIN`) declare that a parallel implementation is also available. See [Parallel Tasks](#) for details.

**starpu\_codelet::max\_parallelism** Optional field. If a parallel implementation is available, this denotes the maximum combined worker size that StarPU will use to execute parallel tasks for this codelet.

#### **starpu\_codelet::cpu\_func**

**Deprecated** Optional field which has been made deprecated. One should use instead the field [starpu\\_codelet::cpu\\_funcs](#).

#### **starpu\_codelet::cuda\_func**

**Deprecated** Optional field which has been made deprecated. One should use instead the [starpu\\_codelet::cuda\\_funcs](#) field.

#### **starpu\_codelet::opencl\_func**

**Deprecated** Optional field which has been made deprecated. One should use instead the [starpu\\_codelet::opencl\\_funcs](#) field.

**starpu\_codelet::cpu\_funcs** Optional array of function pointers to the CPU implementations of the codelet. It must be terminated by a NULL value. The functions prototype must be:

```
void cpu_func(void *buffers[], void *cl_arg)
```

The first argument being the array of data managed by the data management library, and the second argument is a pointer to the argument passed from the field [starpu\\_task::cl\\_arg](#). If the field [starpu\\_codelet::where](#) is set, then the field [starpu\\_codelet::cpu\\_funcs](#) is ignored if [STARPU\\_CPU](#) does not appear in the field [starpu\\_codelet::where](#), it must be non-null otherwise.

**starpu\_codelet::cuda\_funcs** Optional array of function pointers to the CUDA implementations of the codelet. It must be terminated by a NULL value. The functions must be host-functions written in the CUDA runtime API. Their prototype must be:

```
void cuda_func(void *buffers[], void *cl_arg)
```

If the field [starpu\\_codelet::where](#) is set, then the field [starpu\\_codelet::cuda\\_funcs](#) is ignored if [STARPU\\_CUDA](#) does not appear in the field [starpu\\_codelet::where](#), it must be non-null otherwise.

**starpu\_codelet::opencl\_funcs** Optional array of function pointers to the OpenCL implementations of the codelet. It must be terminated by a NULL value. The functions prototype must be:

```
void opencl_func(void *buffers[], void *cl_arg)
```

If the field [starpu\\_codelet::where](#) field is set, then the field [starpu\\_codelet::opencl\\_funcs](#) is ignored if [STARPU\\_OPENCL](#) does not appear in the field [starpu\\_codelet::where](#), it must be non-null otherwise.

**starpu\_codelet::nbuffers** Specify the number of arguments taken by the codelet. These arguments are managed by the DSM and are accessed from the `void *buffers[]` array. The constant argument passed with the field [starpu\\_task::cl\\_arg](#) is not counted in this number. This value should not be above [STARPU\\_NMAXBUFS](#).

**starpu\_codelet::modes** Is an array of [starpu\\_data\\_access\\_mode](#). It describes the required access modes to the data needed by the codelet (e.g. [STARPU\\_RW](#)). The number of entries in this array must be specified in the field [starpu\\_codelet::nbuffers](#), and should not exceed [STARPU\\_NMAXBUFS](#). If unsufficient, this value can be set with the configure option `--enable-maxbuffers`.

**starpu\_codelet::dyn\_modes** Is an array of [starpu\\_data\\_access\\_mode](#). It describes the required access modes to the data needed by the codelet (e.g. [STARPU\\_RW](#)). The number of entries in this array must be specified in the field [starpu\\_codelet::nbuffers](#). This field should be used for codelets having a number of datas greater than [STARPU\\_NMAXBUFS](#) (see [Setting The Data Handles For A Task](#)). When defining a codelet, one should either define this field or the field [starpu\\_codelet::modes](#) defined above.

**starpu\_codelet::model** Optional pointer to the task duration performance model associated to this codelet. This optional field is ignored when set to `NULL` or when its field `starpu_perfmodel::symbol` is not set.

**starpu\_codelet::power\_model** Optional pointer to the task power consumption performance model associated to this codelet. This optional field is ignored when set to `NULL` or when its field `starpu_perfmodel::field` is not set. In the case of parallel codelets, this has to account for all processing units involved in the parallel execution.

**starpu\_codelet::per\_worker\_stats** Optional array for statistics collected at runtime: this is filled by StarPU and should not be accessed directly, but for example by calling the function `starpu_codelet_display_stats()` (See `starpu_codelet_display_stats()` for details).

**starpu\_codelet::name** Optional name of the codelet. This can be useful for debugging purposes.

### struct starpu\_data\_descr

This type is used to describe a data handle along with an access mode.

#### Data Fields

<code>starpu_data_handle_t</code>	handle	describes a data
enum <code>starpu_data_access_mode</code>	mode	describes its access mode

### struct starpu\_task

The structure describes a task that can be offloaded on the various processing units managed by StarPU. It instantiates a codelet. It can either be allocated dynamically with the function `starpu_task_create()`, or declared statically. In the latter case, the programmer has to zero the structure `starpu_task` and to fill the different fields properly. The indicated default values correspond to the configuration of a task allocated with `starpu_task_create()`.

#### Data Fields

- struct `starpu_codelet * cl`
- struct `starpu_data_descr buffers [STARPU_NMAXBUFS]`
- `starpu_data_handle_t handles [STARPU_NMAXBUFS]`
- `void * interfaces [STARPU_NMAXBUFS]`
- `starpu_data_handle_t * dyn_handles`
- `void ** dyn_interfaces`
- `void * cl_arg`
- `size_t cl_arg_size`
- `unsigned cl_arg_free`
- `void(* callback_func )(void *)`
- `void * callback_arg`
- `unsigned callback_arg_free`
- `void(* prologue_callback_func )(void *)`
- `void * prologue_callback_arg`
- `unsigned prologue_callback_arg_free`
- `unsigned use_tag`
- `starpu_tag_t tag_id`
- `unsigned sequential_consistency`
- `unsigned synchronous`
- `int priority`
- `unsigned execute_on_a_specific_worker`
- `unsigned workerid`
- `starpu_task_bundle_t bundle`

- int `detach`
- int `destroy`
- int `regenerate`
- enum `starpu_task_status` `status`
- struct `starpu_profiling_task_info *` `profiling_info`
- double `predicted`
- double `predicted_transfer`
- unsigned int `mf_skip`
- struct `starpu_task *` `prev`
- struct `starpu_task *` `next`
- void \* `starpu_private`
- int `magic`
- unsigned `sched_ctx`
- int `hypervisor_tag`
- double `flops`
- unsigned `scheduled`

## Field Documentation

**`starpu_task::cl`** Is a pointer to the corresponding structure `starpu_codelet`. This describes where the kernel should be executed, and supplies the appropriate implementations. When set to NULL, no code is executed during the tasks, such empty tasks can be useful for synchronization purposes.

### **`starpu_task::buffers`**

**Deprecated** This field has been made deprecated. One should use instead the field `starpu_task::handles` to specify the data handles accessed by the task. The access modes are now defined in the field `starpu_codelet::modes`.

**`starpu_task::handles`** Is an array of `starpu_data_handle_t`. It specifies the handles to the different pieces of data accessed by the task. The number of entries in this array must be specified in the field `starpu_codelet::nbuffers`, and should not exceed `STARPU_NMAXBUFS`. If unsufficient, this value can be set with the configure option `--enable-maxbuffers`.

**`starpu_task::interfaces`** The actual data pointers to the memory node where execution will happen, managed by the DSM.

**`starpu_task::dyn_handles`** Is an array of `starpu_data_handle_t`. It specifies the handles to the different pieces of data accessed by the task. The number of entries in this array must be specified in the field `starpu_codelet::nbuffers`. This field should be used for tasks having a number of datas greater than `STARPU_NMAXBUFS` (see [Setting The Data Handles For A Task](#)). When defining a task, one should either define this field or the field `starpu_task::handles` defined above.

**`starpu_task::dyn_interfaces`** The actual data pointers to the memory node where execution will happen, managed by the DSM. Is used when the field `starpu_task::dyn_handles` is defined.

**`starpu_task::cl_arg`** Optional pointer which is passed to the codelet through the second argument of the codelet implementation (e.g. `starpu_codelet::cpu_func` or `starpu_codelet::cuda_func`). The default value is NULL.

**`starpu_task::cl_arg_size`** Optional field. For some specific drivers, the pointer `starpu_task::cl_arg` cannot not be directly given to the driver function. A buffer of size `starpu_task::cl_arg_size` needs to be allocated on the driver. This buffer is then filled with the `starpu_task::cl_arg_size` bytes starting at address `starpu_task::cl_arg`. In this case, the argument given to the codelet is therefore not the `starpu_task::cl_arg` pointer, but the address of the buffer in local store (LS) instead. This field is ignored for CPU, CUDA and OpenCL codelets, where the `starpu_task::cl_arg` pointer is given as such.

**starpu\_task::cl\_arg\_free** Optional field. In case `starpu_task::cl_arg` was allocated by the application through `malloc()`, setting `starpu_task::cl_arg_free` to 1 makes StarPU automatically call `free(cl_arg)` when destroying the task. This saves the user from defining a callback just for that. This is mostly useful when targetting devices where the codelet does not execute in the same memory space as the main thread.

**starpu\_task::callback\_func** Optional field, the default value is `NULL`. This is a function pointer of prototype `void (*f)(void *)` which specifies a possible callback. If this pointer is non-null, the callback function is executed on the host after the execution of the task. Tasks which depend on it might already be executing. The callback is passed the value contained in the `starpu_task::callback_arg` field. No callback is executed if the field is set to `NULL`.

**starpu\_task::callback\_arg** Optional field, the default value is `NULL`. This is the pointer passed to the callback function. This field is ignored if the field `starpu_task::callback_func` is set to `NULL`.

**starpu\_task::callback\_arg\_free** Optional field. In case `starpu_task::callback_arg` was allocated by the application through `malloc()`, setting `starpu_task::callback_arg_free` to 1 makes StarPU automatically call `free(callback_arg)` when destroying the task.

**starpu\_task::prologue\_callback\_func** Optional field, the default value is `NULL`. This is a function pointer of prototype `void (*f)(void *)` which specifies a possible callback. If this pointer is non-null, the callback function is executed on the host when the task becomes ready for execution, before getting scheduled. The callback is passed the value contained in the `starpu_task::prologue_callback_arg` field. No callback is executed if the field is set to `NULL`.

**starpu\_task::prologue\_callback\_arg** Optional field, the default value is `NULL`. This is the pointer passed to the prologue callback function. This field is ignored if the field `starpu_task::prologue_callback_func` is set to `NULL`.

**starpu\_task::prologue\_callback\_arg\_free** Optional field. In case `starpu_task::prologue_callback_arg` was allocated by the application through `malloc()`, setting `starpu_task::prologue_callback_arg_free` to 1 makes StarPU automatically call `free(prologue_callback_arg)` when destroying the task.

**starpu\_task::use\_tag** Optional field, the default value is 0. If set, this flag indicates that the task should be associated with the tag contained in the `starpu_task::tag_id` field. Tag allow the application to synchronize with the task and to express task dependencies easily.

**starpu\_task::tag\_id** This optional field contains the tag associated to the task if the field `starpu_task::use_tag` is set, it is ignored otherwise.

**starpu\_task::sequential\_consistency** If this flag is set (which is the default), sequential consistency is enforced for the data parameters of this task for which sequential consistency is enabled. Clearing this flag permits to disable sequential consistency for this task, even if data have it enabled.

**starpu\_task::synchronous** If this flag is set, the function `starpu_task_submit()` is blocking and returns only when the task has been executed (or if no worker is able to process the task). Otherwise, `starpu_task_submit()` returns immediately.

**starpu\_task::priority** Optional field, the default value is `STARPU_DEFAULT_PRIO`. This field indicates a level of priority for the task. This is an integer value that must be set between the return values of the function `starpu_sched_get_min_priority()` for the least important tasks, and that of the function `starpu_sched_get_max_priority()` for the most important tasks (included). The `STARPU_MIN_PRIO` and `STARPU_MAX_PRIO` macros are provided for convenience and respectively returns the value of `starpu_sched_get_min_priority()` and `starpu_sched_get_max_priority()`. Default priority is `STARPU_DEFAULT_PRIO`, which is always defined as 0 in order to allow static task initialization. Scheduling strategies that take priorities into account can use this parameter to take better scheduling decisions, but the scheduling policy may also ignore it.

**starpu\_task::execute\_on\_a\_specific\_worker** Default value is 0. If this flag is set, StarPU will bypass the scheduler and directly affect this task to the worker specified by the field `starpu_task::workerid`.

**starpu\_task::workerid** Optional field. If the field `starpu_task::execute_on_a_specific_worker` is set, this field indicates the identifier of the worker that should process this task (as returned by `starpu_worker.get_id()`). This field is ignored if the field `starpu_task::execute_on_a_specific_worker` is set to 0.

**starpu\_task::bundle** Optional field. The bundle that includes this task. If no bundle is used, this should be NULL.

**starpu\_task::detach** Optional field, default value is 1. If this flag is set, it is not possible to synchronize with the task by the means of `starpu_task_wait()` later on. Internal data structures are only guaranteed to be freed once `starpu_task_wait()` is called if the flag is not set.

**starpu\_task::destroy** Optional value. Default value is 0 for `starpu_task_init()`, and 1 for `starpu_task_create()`. If this flag is set, the task structure will automatically be freed, either after the execution of the callback if the task is detached, or during `starpu_task_wait()` otherwise. If this flag is not set, dynamically allocated data structures will not be freed until `starpu_task_destroy()` is called explicitly. Setting this flag for a statically allocated task structure will result in undefined behaviour. The flag is set to 1 when the task is created by calling `starpu_task_create()`. Note that `starpu_task_wait_for_all()` will not free any task.

**starpu\_task::regenerate** Optional field. If this flag is set, the task will be re-submitted to StarPU once it has been executed. This flag must not be set if the flag `starpu_task::destroy` is set.

**starpu\_task::status** Optional field. Current state of the task.

**starpu\_task::profiling\_info** Optional field. Profiling information for the task.

**starpu\_task::predicted** Output field. Predicted duration of the task. This field is only set if the scheduling strategy uses performance models.

**starpu\_task::predicted\_transfer** Optional field. Predicted data transfer duration for the task in microseconds. This field is only valid if the scheduling strategy uses performance models.

**starpu\_task::mf\_skip** This is only used for tasks that use multiformat handle. This should only be used by StarPU.

**starpu\_task::prev** A pointer to the previous task. This should only be used by StarPU.

**starpu\_task::next** A pointer to the next task. This should only be used by StarPU.

**starpu\_task::starpu\_private** This is private to StarPU, do not modify. If the task is allocated by hand (without `starpu_task_create()`), this field should be set to NULL.

**starpu\_task::magic** This field is set when initializing a task. The function `starpu_task_submit()` will fail if the field does not have the right value. This will hence avoid submitting tasks which have not been properly initialised.

**starpu\_task::sched\_ctx** Scheduling context.

**starpu\_task::hypervisor\_tag** Helps the hypervisor monitor the execution of this task.

**starpu\_task::flops** This can be set to the number of floating points operations that the task will have to achieve. This is useful for easily getting GFlops curves from the tool `starpu_perfmmodel_plot`, and for the hypervisor load balancing.

**starpu\_task::scheduled** Whether the scheduler has pushed the task on some queue

### 17.11.3 Macro Definition Documentation

#### #define STARPU\_CPU

This macro is used when setting the field `starpu_codelet::where` to specify the codelet may be executed on a CPU processing unit.

#### #define STARPU\_CUDA

This macro is used when setting the field `starpu_codelet::where` to specify the codelet may be executed on a CUDA processing unit.

#### #define STARPU\_OPENCL

This macro is used when setting the field `starpu_codelet::where` to specify the codelet may be executed on a OpenCL processing unit.

#### #define STARPU\_MULTIPLE\_CPU\_IMPLEMENTATIONS

**Deprecated** Setting the field `starpu_codelet::cpu_func` with this macro indicates the codelet will have several implementations. The use of this macro is deprecated. One should always only define the field `starpu_codelet::cpu_funcs`.

#### #define STARPU\_MULTIPLE\_CUDA\_IMPLEMENTATIONS

**Deprecated** Setting the field `starpu_codelet::cuda_func` with this macro indicates the codelet will have several implementations. The use of this macro is deprecated. One should always only define the field `starpu_codelet::cuda_funcs`.

#### #define STARPU\_MULTIPLE\_OPENCL\_IMPLEMENTATIONS

**Deprecated** Setting the field `starpu_codelet::opencl_func` with this macro indicates the codelet will have several implementations. The use of this macro is deprecated. One should always only define the field `starpu_codelet::opencl_funcs`.

#### #define STARPU\_NMAXBUFS

Defines the maximum number of buffers that tasks will be able to take as parameters. The default value is 8, it can be changed by using the configure option `--enable-maxbuffers`.

#### #define STARPU\_TASK\_INITIALIZER

It is possible to initialize statically allocated tasks with this value. This is equivalent to initializing a structure `starpu_task` with the function `starpu_task.init()` function.

#### #define STARPU\_TASK\_GET\_HANDLE( *task*, *i* )

Return the *i* th data handle of the given task. If the task is defined with a static or dynamic number of handles, will either return the *i* th element of the field `starpu_task::handles` or the *i* th element of the field `starpu_task::dyn_handles` (see [Setting The Data Handles For A Task](#))

#### #define STARPU\_TASK\_SET\_HANDLE( *task*, *handle*, *i* )

Set the *i* th data handle of the given task with the given dat handle. If the task is defined with a static or dynamic number of handles, will either set the *i* th element of the field `starpu_task::handles` or the *i* th element of the field `starpu_task::dyn_handles` (see [Setting The Data Handles For A Task](#))

```
#define STARPU_CODELET_GET_MODE( codelet, i )
```

Return the access mode of the  $i$  th data handle of the given codelet. If the codelet is defined with a static or dynamic number of handles, will either return the  $i$  th element of the field `starpu_codelet::modes` or the  $i$  th element of the field `starpu_codelet::dyn_modes` (see [Setting The Data Handles For A Task](#))

```
#define STARPU_CODELET_SET_MODE( codelet, mode, i )
```

Set the access mode of the  $i$  th data handle of the given codelet. If the codelet is defined with a static or dynamic number of handles, will either set the  $i$  th element of the field `starpu_codelet::modes` or the  $i$  th element of the field `starpu_codelet::dyn_modes` (see [Setting The Data Handles For A Task](#))

#### `starpu_task_status::STARPU_TASK_INVALID`

The task has just been initialized.

### 17.11.4 Typedef Documentation

#### `starpu_cpu_func_t`

CPU implementation of a codelet.

#### `starpu_cuda_func_t`

CUDA implementation of a codelet.

#### `starpu_opencl_func_t`

OpenCL implementation of a codelet.

### 17.11.5 Enumeration Type Documentation

#### `enum starpu_codelet_type`

Describes the type of parallel task. See [Parallel Tasks](#) for details.

Enumerator

**`STARPU_SEQ`** (default) for classical sequential tasks.

**`STARPU_SPMD`** for a parallel task whose threads are handled by StarPU, the code has to use `starpu_combined_worker_get_size()` and `starpu_combined_worker_get_rank()` to distribute the work.

**`STARPU_FORKJOIN`** for a parallel task whose threads are started by the codelet function, which has to use `starpu_combined_worker_get_size()` to determine how many threads should be started.

#### `enum starpu_task_status`

Task status

Enumerator

**`STARPU_TASK_BLOCKED`** The task has just been submitted, and its dependencies has not been checked yet.

**`STARPU_TASK_READY`** The task is ready for execution.

**`STARPU_TASK_RUNNING`** The task is running on some worker.

**`STARPU_TASK_FINISHED`** The task is finished executing.

**`STARPU_TASK_BLOCKED_ON_TAG`** The task is waiting for a tag.

**`STARPU_TASK_BLOCKED_ON_TASK`** The task is waiting for a task.

**`STARPU_TASK_BLOCKED_ON_DATA`** The task is waiting for some data.

### 17.11.6 Function Documentation

**void starpu\_codelet\_init ( struct starpu\_codelet \* cl )**

Initialize `cl` with default values. Codelets should preferably be initialized statically as shown in [Defining A Codelet](#). However such a initialisation is not always possible, e.g. when using C++.

**void starpu\_task\_init ( struct starpu\_task \* task )**

Initialize task with default values. This function is implicitly called by [starpu\\_task\\_create\(\)](#). By default, tasks initialized with [starpu\\_task\\_init\(\)](#) must be deinitialized explicitly with [starpu\\_task\\_clean\(\)](#). Tasks can also be initialized statically, using [STARPU\\_TASK\\_INITIALIZER](#).

**struct starpu\_task \* starpu\_task\_create ( void )**

Allocate a task structure and initialize it with default values. Tasks allocated dynamically with [starpu\\_task\\_create\(\)](#) are automatically freed when the task is terminated. This means that the task pointer can not be used any more once the task is submitted, since it can be executed at any time (unless dependencies make it wait) and thus freed at any time. If the field [starpu\\_task::destroy](#) is explicitly unset, the resources used by the task have to be freed by calling [starpu\\_task\\_destroy\(\)](#).

**struct starpu\_task \* starpu\_task\_dup ( struct starpu\_task \* task )**

Allocate a task structure which is the exact duplicate of the given task.

**void starpu\_task\_clean ( struct starpu\_task \* task )**

Release all the structures automatically allocated to execute task, but not the task structure itself and values set by the user remain unchanged. It is thus useful for statically allocated tasks for instance. It is also useful when users want to execute the same operation several times with as least overhead as possible. It is called automatically by [starpu\\_task\\_destroy\(\)](#). It has to be called only after explicitly waiting for the task or after [starpu\\_shutdown\(\)](#) (waiting for the callback is not enough, since StarPU still manipulates the task after calling the callback).

**void starpu\_task\_destroy ( struct starpu\_task \* task )**

Free the resource allocated during [starpu\\_task\\_create\(\)](#) and associated with task. This function is already called automatically after the execution of a task when the field [starpu\\_task::destroy](#) is set, which is the default for tasks created by [starpu\\_task\\_create\(\)](#). Calling this function on a statically allocated task results in an undefined behaviour.

**int starpu\_task\_wait ( struct starpu\_task \* task )**

This function blocks until `task` has been executed. It is not possible to synchronize with a task more than once. It is not possible to wait for synchronous or detached tasks. Upon successful completion, this function returns 0. Otherwise, `-EINVAL` indicates that the specified task was either synchronous or detached.

**int starpu\_task\_submit ( struct starpu\_task \* task )**

This function submits task to StarPU. Calling this function does not mean that the task will be executed immediately as there can be data or task (tag) dependencies that are not fulfilled yet: StarPU will take care of scheduling this task with respect to such dependencies. This function returns immediately if the field [starpu\\_task::synchronous](#) is set to 0, and block until the termination of the task otherwise. It is also possible to synchronize the application with asynchronous tasks by the means of tags, using the function [starpu\\_tag\\_wait\(\)](#) function for instance. In case of success, this function returns 0, a return value of `-ENODEV` means that there is no worker able to process this task (e.g. there is no GPU available and this task is only implemented for CUDA devices). [starpu\\_task\\_submit\(\)](#) can be called from anywhere, including codelet functions and callbacks, provided that the field [starpu\\_task::synchronous](#) is set to 0.

```
int starpu_task_submit_to_ctx ( struct starpu_task * task, unsigned sched_ctx_id )
```

This function submits a task to StarPU to the context `sched_ctx_id`. By default `starpu_task_submit` submits the task to a global context that is created automatically by StarPU.

```
int starpu_task_wait_for_all ( void )
```

This function blocks until all the tasks that were submitted (to the current context or the global one if there aren't any) are terminated. It does not destroy these tasks.

```
int starpu_task_wait_for_all_in_ctx ( unsigned sched_ctx_id )
```

This function waits until all the tasks that were already submitted to the context `sched_ctx_id` have been executed

```
int starpu_task_nready ( void )
```

TODO

Return the number of submitted tasks which are ready for execution are already executing. It thus does not include tasks waiting for dependencies.

```
int starpu_task_nsubmitted ( void )
```

Return the number of submitted tasks which have not completed yet.

```
struct starpu_task * starpu_task_get_current ( void )
```

This function returns the task currently executed by the worker, or `NULL` if it is called either from a thread that is not a task or simply because there is no task being executed at the moment.

```
void starpu_codelet_display_stats ( struct starpu_codelet * cl )
```

Output on stderr some statistics on the codelet `cl`.

```
int starpu_task_wait_for_no_ready ( void )
```

This function waits until there is no more ready task.

```
void starpu_task_set_implementation ( struct starpu_task * task, unsigned impl )
```

This function should be called by schedulers to specify the codelet implementation to be executed when executing the task.

```
unsigned starpu_task_get_implementation ( struct starpu_task * task )
```

This function return the codelet implementation to be executed when executing the task.

```
void starpu_create_sync_task ( starpu_tag_t sync_tag, unsigned ndeps, starpu_tag_t * deps, void(*)(void *) callback, void * callback_arg )
```

This creates (and submits) an empty task that unlocks a tag once all its dependencies are fulfilled.

## 17.12 Insert\_Task

### Macros

- #define `STARPU_VALUE`
- #define `STARPU_CALLBACK`
- #define `STARPU_CALLBACK_WITH_ARG`

- #define STARPU\_CALLBACK\_ARG
- #define STARPU\_PRIORITY
- #define STARPU\_DATA\_ARRAY
- #define STARPU\_TAG
- #define STARPU\_FLOPS
- #define STARPU\_SCHED\_CTX

## Functions

- int `starpu_insert_task` (struct `starpu_codelet` \*cl,...)
- void `starpu_codelet_pack_args` (void \*\*arg\_buffer, size\_t \*arg\_buffer\_size,...)
- void `starpu_codelet_unpack_args` (void \*cl\_arg,...)
- struct `starpu_task` \* `starpu_task_build` (struct `starpu_codelet` \*cl,...)

### 17.12.1 Detailed Description

### 17.12.2 Macro Definition Documentation

#### #define STARPU\_VALUE

this macro is used when calling `starpu_insert_task()`, and must be followed by a pointer to a constant value and the size of the constant

#### #define STARPU\_CALLBACK

this macro is used when calling `starpu_insert_task()`, and must be followed by a pointer to a callback function

#### #define STARPU\_CALLBACK\_WITH\_ARG

this macro is used when calling `starpu_insert_task()`, and must be followed by two pointers: one to a callback function, and the other to be given as an argument to the callback function; this is equivalent to using both `STARPU_CALLBACK` and `STARPU_CALLBACK_WITH_ARG`.

#### #define STARPU\_CALLBACK\_ARG

this macro is used when calling `starpu_insert_task()`, and must be followed by a pointer to be given as an argument to the callback function

#### #define STARPU\_PRIORITY

this macro is used when calling `starpu_insert_task()`, and must be followed by a integer defining a priority level

#### #define STARPU\_DATA\_ARRAY

TODO

#### #define STARPU\_TAG

this macro is used when calling `starpu_insert_task()`, and must be followed by a tag.

#### #define STARPU\_FLOPS

this macro is used when calling `starpu_insert_task()`, and must be followed by an amount of floating point operations, as a double. Users **MUST** explicitly cast into double, otherwise parameter passing will not work.

#### #define STARPU\_SCHED\_CTX

this macro is used when calling `starpu_insert_task()`, and must be followed by the id of the scheduling context to which we want to submit the task.

### 17.12.3 Function Documentation

```
int starpu_insert_task ( struct starpu_codelet * cl, ... )
```

Create and submit a task corresponding to `cl` with the following arguments. The argument list must be zero-terminated.

The arguments following the codelet can be of the following types:

- `STARPU_R`, `STARPU_W`, `STARPU_RW`, `STARPU_SCRATCH`, `STARPU_REDUX` an access mode followed by a data handle;
- `STARPU_DATA_ARRAY` followed by an array of data handles and its number of elements;
- the specific values `STARPU_VALUE`, `STARPU_CALLBACK`, `STARPU_CALLBACK_ARG`, `STARPU_CALLBACK_WITH_ARG`, `STARPU_PRIORITY`, `STARPU_TAG`, `STARPU_FLOPS`, `STARPU_SCHED_CTX` followed by the appropriated objects as defined elsewhere.

When using `STARPU_DATA_ARRAY`, the access mode of the data handles is not defined.

Parameters to be passed to the codelet implementation are defined through the type `STARPU_VALUE`. The function `starpu_codelet_unpack_args()` must be called within the codelet implementation to retrieve them.

```
void starpu_codelet_pack_args ( void ** arg_buffer, size_t * arg_buffer_size, ... )
```

Pack arguments of type `STARPU_VALUE` into a buffer which can be given to a codelet and later unpacked with the function `starpu_codelet_unpack_args()`.

```
void starpu_codelet_unpack_args ( void * cl_arg, ... )
```

Retrieve the arguments of type `STARPU_VALUE` associated to a task automatically created using the function `starpu_insert_task()`.

```
struct starpu_task * starpu_task_build ( struct starpu_codelet * cl, ... )
```

Create a task corresponding to `cl` with the following arguments. The argument list must be zero-terminated. The arguments following the codelet are the same as the ones for the function `starpu_insert_task()`.

## 17.13 Explicit Dependencies

### Typedefs

- `typedef uint64_t starpu_tag_t`

### Functions

- `void starpu_task_declare_deps_array (struct starpu_task *task, unsigned ndeps, struct starpu_task *task_array[])`
- `void starpu_tag_declare_deps (starpu_tag_t id, unsigned ndeps,...)`
- `void starpu_tag_declare_deps_array (starpu_tag_t id, unsigned ndeps, starpu_tag_t *array)`
- `int starpu_tag_wait (starpu_tag_t id)`
- `int starpu_tag_wait_array (unsigned ntags, starpu_tag_t *id)`
- `void starpu_tag_restart (starpu_tag_t id)`
- `void starpu_tag_remove (starpu_tag_t id)`
- `void starpu_tag_notify_from_apps (starpu_tag_t id)`

### 17.13.1 Detailed Description

### 17.13.2 Typedef Documentation

#### `starpu_tag_t`

This type defines a task logical identifier. It is possible to associate a task with a unique `tag` chosen by the application, and to express dependencies between tasks by the means of those tags. To do so, fill the field `starpu_task::tag_id` with a tag number (can be arbitrary) and set the field `starpu_task::use_tag` to 1. If `starpu_tag_declare_deps()` is called with this tag number, the task will not be started until the tasks which holds the declared dependency tags are completed.

### 17.13.3 Function Documentation

```
void starpu_task_declare_deps_array ( struct starpu_task * task, unsigned ndeps, struct starpu_task * task_array[] )
```

Declare task dependencies between a `task` and an array of tasks of length `ndeps`. This function must be called prior to the submission of the task, but it may be called after the submission or the execution of the tasks in the array, provided the tasks are still valid (i.e. they were not automatically destroyed). Calling this function on a task that was already submitted or with an entry of `task_array` that is no longer a valid task results in an undefined behaviour. If `ndeps` is 0, no dependency is added. It is possible to call `starpu_task_declare_deps_array()` several times on the same task, in this case, the dependencies are added. It is possible to have redundancy in the task dependencies.

```
void starpu_tag_declare_deps ( starpu_tag_t id, unsigned ndeps, ... )
```

Specify the dependencies of the task identified by tag `id`. The first argument specifies the tag which is configured, the second argument gives the number of tag(s) on which `id` depends. The following arguments are the tags which have to be terminated to unlock the task. This function must be called before the associated task is submitted to StarPU with `starpu_task_submit()`.

**WARNING! Use with caution.** Because of the variable arity of `starpu_tag_declare_deps()`, note that the last arguments must be of type `starpu_tag_t`: constant values typically need to be explicitly casted. Otherwise, due to integer sizes and argument passing on the stack, the C compiler might consider the tag `0x200000003` instead of `0x2` and `0x3` when calling `starpu_tag_declare_deps(0x1, 2, 0x2, 0x3)`. Using the `starpu_tag_declare_deps_array()` function avoids this hazard.

```
/* Tag 0x1 depends on tags 0x32 and 0x52 */
starpu_tag_declare_deps((starpu_tag_t)0x1, 2,
                      starpu_tag_t)0x32, (starpu_tag_t)0x52);
```

```
void starpu_tag_declare_deps_array ( starpu_tag_t id, unsigned ndeps, starpu_tag_t * array )
```

This function is similar to `starpu_tag_declare_deps()`, except that it does not take a variable number of arguments but an array of tags of size `ndeps`.

```
/* Tag 0x1 depends on tags 0x32 and 0x52 */
starpu_tag_t tag_array[2] = {0x32, 0x52};
starpu_tag_declare_deps_array((starpu_tag_t)0x1, 2, tag_array);
```

```
int starpu_tag_wait ( starpu_tag_t id )
```

This function blocks until the task associated to tag `id` has been executed. This is a blocking call which must therefore not be called within tasks or callbacks, but only from the application directly. It is possible to synchronize with the same tag multiple times, as long as the `starpu_tag_remove()` function is not called. Note that it is still possible to synchronize with a tag associated to a task for which the strucuture `starpu_task` was freed (e.g. if the field `starpu_task::destroy` was enabled).

```
int starpu_tag_wait_array ( unsigned ntags, starpu_tag_t * id )
```

This function is similar to `starpu_tag_wait()` except that it blocks until all the `ntags` tags contained in the array `id` are terminated.

**void starpu.tag.restart( starpu.tag\_t id )**

This function can be used to clear the *already notified* status of a tag which is not associated with a task. Before that, calling [starpu.tag.notify\\_from\\_apps\(\)](#) again will not notify the successors. After that, the next call to [starpu.tag.notify\\_from\\_apps\(\)](#) will notify the successors.

**void starpu.tag.remove( starpu.tag\_t id )**

This function releases the resources associated to tag *id*. It can be called once the corresponding task has been executed and when there is no other tag that depend on this tag anymore.

**void starpu.tag.notify\_from\_apps( starpu.tag\_t id )**

This function explicitly unlocks tag *id*. It may be useful in the case of applications which execute part of their computation outside StarPU tasks (e.g. third-party libraries). It is also provided as a convenient tool for the programmer, for instance to entirely construct the task DAG before actually giving StarPU the opportunity to execute the tasks. When called several times on the same tag, notification will be done only on first call, thus implementing "OR" dependencies, until the tag is restarted using [starpu.tag.restart\(\)](#).

## 17.14 Implicit Data Dependencies

In this section, we describe how StarPU makes it possible to insert implicit task dependencies in order to enforce sequential data consistency. When this data consistency is enabled on a specific data handle, any data access will appear as sequentially consistent from the application. For instance, if the application submits two tasks that access the same piece of data in read-only mode, and then a third task that access it in write mode, dependencies will be added between the two first tasks and the third one. Implicit data dependencies are also inserted in the case of data accesses from the application.

### Functions

- void [starpu.data.set.default\\_sequential.consistency.flag](#) (unsigned flag)
- unsigned [starpu.data.get.default\\_sequential.consistency.flag](#) (void)
- void [starpu.data.set.sequential.consistency.flag](#) (starpu\_data\_handle\_t handle, unsigned flag)

#### 17.14.1 Detailed Description

In this section, we describe how StarPU makes it possible to insert implicit task dependencies in order to enforce sequential data consistency. When this data consistency is enabled on a specific data handle, any data access will appear as sequentially consistent from the application. For instance, if the application submits two tasks that access the same piece of data in read-only mode, and then a third task that access it in write mode, dependencies will be added between the two first tasks and the third one. Implicit data dependencies are also inserted in the case of data accesses from the application.

#### 17.14.2 Function Documentation

**[starpu.data.set.default\\_sequential.consistency.flag\( unsigned flag \)](#)**

Set the default sequential consistency flag. If a non-zero value is passed, a sequential data consistency will be enforced for all handles registered after this function call, otherwise it is disabled. By default, StarPU enables sequential data consistency. It is also possible to select the data consistency mode of a specific data handle with the function [starpu.data.set.sequential.consistency.flag\(\)](#).

**[unsigned starpu.data.get.default\\_sequential.consistency.flag\( void \)](#)**

Return the default sequential consistency flag

**void starpu\_data\_set\_sequential\_consistency\_flag ( starpu\_data\_handle\_t handle, unsigned flag )**

Set the data consistency mode associated to a data handle. The consistency mode set using this function has the priority over the default mode which can be set with [starpu\\_data\\_set\\_default\\_sequential\\_consistency\\_flag\(\)](#).

## 17.15 Performance Model

### Data Structures

- struct [starpu\\_perfmodel](#)
- struct [starpu\\_perfmodel\\_regression\\_model](#)
- struct [starpu\\_perfmodel\\_per\\_arch](#)
- struct [starpu\\_perfmodel\\_history\\_list](#)
- struct [starpu\\_perfmodel\\_history\\_entry](#)

### Enumerations

- enum [starpu\\_perfmodel\\_archtype](#) { [STARPU\\_CPU\\_DEFAULT](#), [STARPU\\_CUDA\\_DEFAULT](#), [STARPU\\_OPENCL\\_DEFAULT](#) }
- enum [starpu\\_perfmodel\\_type](#) {  
[STARPU\\_PER\\_ARCH](#), [STARPU\\_COMMON](#), [STARPU\\_HISTORY\\_BASED](#), [STARPU\\_REGRESSION\\_BASED](#),  
[STARPU\\_NL\\_REGRESSION\\_BASED](#) }

### Functions

- int [starpu\\_perfmodel\\_load\\_symbol](#) (const char \*symbol, struct [starpu\\_perfmodel](#) \*model)
- int [starpu\\_perfmodel\\_unload\\_model](#) (struct [starpu\\_perfmodel](#) \*model)
- void [starpu\\_perfmodel\\_debugfilepath](#) (struct [starpu\\_perfmodel](#) \*model, enum [starpu\\_perfmodel\\_archtype](#) arch, char \*path, size\_t maxlen, unsigned nimpl)
- void [starpu\\_perfmodel\\_get\\_arch\\_name](#) (enum [starpu\\_perfmodel\\_archtype](#) arch, char \*archname, size\_t maxlen, unsigned nimpl)
- enum [starpu\\_perfmodel\\_archtype](#) [starpu\\_worker\\_get\\_perf\\_archtype](#) (int workerid)
- int [starpu\\_perfmodel\\_list](#) (FILE \*output)
- void [starpu\\_perfmodel\\_directory](#) (FILE \*output)
- void [starpu\\_perfmodel\\_print](#) (struct [starpu\\_perfmodel](#) \*model, enum [starpu\\_perfmodel\\_archtype](#) arch, unsigned nimpl, char \*parameter, uint32\_t \*footprint, FILE \*output)
- int [starpu\\_perfmodel\\_print\\_all](#) (struct [starpu\\_perfmodel](#) \*model, char \*arch, char \*parameter, uint32\_t \*footprint, FILE \*output)
- void [starpu\\_bus\\_print\\_bandwidth](#) (FILE \*f)
- void [starpu\\_bus\\_print\\_affinity](#) (FILE \*f)
- void [starpu\\_perfmodel\\_update\\_history](#) (struct [starpu\\_perfmodel](#) \*model, struct [starpu\\_task](#) \*task, enum [starpu\\_perfmodel\\_archtype](#) arch, unsigned cpuid, unsigned nimpl, double measured)
- double [starpu\\_transfer\\_bandwidth](#) (unsigned src\_node, unsigned dst\_node)
- double [starpu\\_transfer\\_latency](#) (unsigned src\_node, unsigned dst\_node)
- double [starpu\\_transfer\\_predict](#) (unsigned src\_node, unsigned dst\_node, size\_t size)

#### 17.15.1 Detailed Description

#### 17.15.2 Data Structure Documentation

##### struct [starpu\\_perfmodel](#)

Contains all information about a performance model. At least the type and symbol fields have to be filled when defining a performance model for a codelet. For compatibility, make sure to initialize the whole structure to zero, either by using explicit memset, or by letting the compiler implicitly do it in e.g. static storage case. If not provided, other fields have to be zero.

## Data Fields

- enum `starpu_perfmodel_type` type
- double(\* `cost_model` )(struct `starpu_data_descr` \*)
- double(\* `cost_function` )(struct `starpu_task` \*, unsigned nimpl)
- size\_t(\* `size_base` )(struct `starpu_task` \*, unsigned nimpl)
- struct `starpu_perfmodel_per_arch_per_arch` [STARPU\_NARCH\_VARIATIONS][STARPU\_MAXIMPLEMENTATIONS]
- const char \* `symbol`
- unsigned `is_loaded`
- unsigned `benchmarking`
- `starpu_pthread_rwlock_t model_rwlock`

## Field Documentation

**`starpu_perfmodel::type`** is the type of performance model

- `STARPU_HISTORY_BASED`, `STARPU_REGRESSION_BASED`, `STARPU_NL_REGRESSION_BASED`: No other fields needs to be provided, this is purely history-based.
- `STARPU_PER_ARCH`: field `starpu_perfmodel::per_arch` has to be filled with functions which return the cost in micro-seconds.
- `STARPU_COMMON`: field `starpu_perfmodel::cost.function` has to be filled with a function that returns the cost in micro-seconds on a CPU, timing on other archs will be determined by multiplying by an arch-specific factor.

**`starpu_perfmodel::cost_model`**

**Deprecated** This field is deprecated. Use instead the field `starpu_perfmodel::cost.function` field.

**`starpu_perfmodel::cost.function`** Used by `STARPU_COMMON`: takes a task and implementation number, and must return a task duration estimation in micro-seconds.

**`starpu_perfmodel::size_base`** Used by `STARPU_HISTORY_BASED`, `STARPU_REGRESSION_BASED` and `STARPU_NL_REGRESSION_BASED`. If not NULL, takes a task and implementation number, and returns the size to be used as index for history and regression.

**`starpu_perfmodel::per_arch`** Used by `STARPU_PER_ARCH`: array of structures `starpu_per_arch_perfmodel`

**`starpu_perfmodel::symbol`** is the symbol name for the performance model, which will be used as file name to store the model. It must be set otherwise the model will be ignored.

**`starpu_perfmodel::is_loaded`** Whether the performance model is already loaded from the disk.

**`starpu_perfmodel::benchmarking`** Whether the performance model is still being calibrated.

**`starpu_perfmodel::model_rwlock`** Lock to protect concurrency between loading from disk (W), updating the values (W), and making a performance estimation (R).

**struct `starpu_perfmodel_regression_model`**

...

**Data Fields**

double	sumlny	sum of ln(measured)
double	sumlnx	sum of ln(size)
double	sumlnx2	sum of ln(size) <sup>2</sup>
unsigned long	minx	minimum size
unsigned long	maxx	maximum size
double	sumlnxlny	sum of ln(size)*ln(measured)
double	alpha	estimated = alpha * size ^ beta
double	beta	estimated = alpha * size ^ beta
unsigned	valid	whether the linear regression model is valid (i.e. enough measures)
double	a	estimated = a * size ^ b + c
double	b	estimated = a * size ^ b + c
double	c	estimated = a * size ^ b + c
unsigned	nl_valid	whether the non-linear regression model is valid (i.e. enough measures)
unsigned	nsample	number of sample values for non-linear regression

**struct starpu\_perfmodel\_per\_arch**

contains information about the performance model of a given arch.

**Data Fields**

- double(\* [cost\\_model](#) )(struct [starpu\\_data\\_descr](#) \*t)
- double(\* [cost\\_function](#) )(struct [starpu\\_task](#) \*task, enum [starpu\\_perfmodel\\_archtype](#) arch, unsigned nimpl)
- size\_t(\* [size\\_base](#) )(struct [starpu\\_task](#) \*, enum [starpu\\_perfmodel\\_archtype](#) arch, unsigned nimpl)
- struct  
    [starpu\\_perfmodel\\_history\\_table](#) \* [history](#)
- struct  
    [starpu\\_perfmodel\\_history\\_list](#) \* [list](#)
- struct  
    [starpu\\_perfmodel\\_regression\\_model](#) [regression](#)

**Field Documentation****[starpu\\_perfmodel\\_per\\_arch::cost\\_model](#)**

**Deprecated** This field is deprecated. Use instead the field [starpu\\_perfmodel\\_per\\_arch::cost\\_function](#).

**[starpu\\_perfmodel\\_per\\_arch::cost\\_function](#)** Used by [STARPU\\_PER\\_ARCH](#), must point to functions which take a task, the target arch and implementation number (as mere conveniency, since the array is already indexed by these), and must return a task duration estimation in micro-seconds.

**[starpu\\_perfmodel\\_per\\_arch::size\\_base](#)** Same as in structure [starpu\\_perfmodel](#), but per-arch, in case it depends on the architecture-specific implementation.

**[starpu\\_perfmodel\\_per\\_arch::history](#)** The history of performance measurements.

**[starpu\\_perfmodel\\_per\\_arch::list](#)** Used by [STARPU\\_HISTORY\\_BASED](#) and [STARPU\\_NL\\_REGRESSION\\_BASED](#), records all execution history measures.

**[starpu\\_perfmodel\\_per\\_arch::regression](#)** Used by [STARPU\\_REGRESSION\\_BASED](#) and [STARPU\\_NL\\_REGRESSION\\_BASED](#), contains the estimated factors of the regression.

**[struct starpu\\_perfmodel\\_history\\_list](#)**

todo

## Data Fields

<code>struct starpu_perfmodel_history_list *</code>	<code>next</code>	<code>todo</code>
<code>struct starpu_perfmodel_history_entry *</code>	<code>entry</code>	<code>todo</code>

**struct starpu\_perfmodel\_history\_entry**

todo  
Data Fields

<code>double</code>	<code>mean</code>	<code>mean_n = 1/n sum</code>
<code>double</code>	<code>deviation</code>	<code>n dev_n = sum2 - 1/n (sum)^2</code>
<code>double</code>	<code>sum</code>	<code>sum of samples (in <math>\mu</math>s)</code>
<code>double</code>	<code>sum2</code>	<code>sum of samples^2</code>
<code>unsigned</code>	<code>nsample</code>	<code>number of samples</code>
<code>uint32_t</code>	<code>footprint</code>	<code>data footprint</code>
<code>size_t</code>	<code>size</code>	<code>in bytes</code>
<code>double</code>	<code>flops</code>	<code>Provided by the application</code>

**17.15.3 Enumeration Type Documentation****enum starpu\_perfmodel\_archtype**

Enumerates the various types of architectures.

it is possible that we have multiple versions of the same kind of workers, for instance multiple GPUs or even different CPUs within the same machine so we do not use the archetype enum type directly for performance models.

- CPU types range within `STARPU_CPU_DEFAULT` (1 CPU), `STARPU_CPU_DEFAULT+1` (2 CPUs), ... `STARPU_CPU_DEFAULT + STARPU_MAXCPUS - 1` (`STARPU_MAXCPUS` CPUs).
- CUDA types range within `STARPU_CUDA_DEFAULT` (GPU number 0), `STARPU_CUDA_DEFAULT + 1` (GPU number 1), ..., `STARPU_CUDA_DEFAULT + STARPU_MAXCUDADEVS - 1` (GPU number `STARPU_MAXCUDADEVS - 1`).
- OpenCL types range within `STARPU_OPENCL_DEFAULT` (GPU number 0), `STARPU_OPENCL_DEFAULT + 1` (GPU number 1), ..., `STARPU_OPENCL_DEFAULT + STARPU_MAXOPENCLDEVS - 1` (GPU number `STARPU_MAXOPENCLDEVS - 1`).

Enumerator

**`STARPU_CPU_DEFAULT`** CPU combined workers between 0 and `STARPU_MAXCPUS-1`

**`STARPU_CUDA_DEFAULT`** CUDA workers

**`STARPU_OPENCL_DEFAULT`** OpenCL workers

**enum starpu\_perfmodel\_type**

TODO

Enumerator

**`STARPU_PER_ARCH`** Application-provided per-arch cost model function

**`STARPU_COMMON`** Application-provided common cost model function, with per-arch factor

**STARPU\_HISTORY\_BASED** Automatic history-based cost model

**STARPU\_REGRESSION\_BASED** Automatic linear regression-based cost model ( $\alpha * \text{size}^\beta$ )

**STARPU\_NL\_REGRESSION\_BASED** Automatic non-linear regression-based cost model ( $a * \text{size}^\beta + b + c$ )

#### 17.15.4 Function Documentation

**int starpu\_perfmodel\_load\_symbol ( const char \* symbol, struct starpu\_perfmodel \* model )**

loads a given performance model. The model structure has to be completely zero, and will be filled with the information saved in \$STARPU\_HOME/.starpu. The function is intended to be used by external tools that should read the performance model files.

**int starpu\_perfmodel\_unload\_model ( struct starpu\_perfmodel \* model )**

unloads the given model which has been previously loaded through the function [starpu\\_perfmodel.load\\_symbol\(\)](#)

**void starpu\_perfmodel\_debugfilepath ( struct starpu\_perfmodel \* model, enum starpu\_perfmodel\_archtype arch, char \* path, size\_t maxlen, unsigned nimpl )**

returns the path to the debugging information for the performance model.

**void starpu\_perfmodel\_get\_arch\_name ( enum starpu\_perfmodel\_archtype arch, char \* archname, size\_t maxlen, unsigned nimpl )**

returns the architecture name for `arch`

**enum starpu\_perfmodel\_archtype starpu\_worker\_get\_perf\_archtype ( int workerid )**

returns the architecture type of a given worker.

**int starpu\_perfmodel\_list ( FILE \* output )**

prints a list of all performance models on `output`

**int starpu\_perfmodel\_directory ( FILE \* output )**

prints the directory name storing performance models on `output`

**void starpu\_perfmodel\_print ( struct starpu\_perfmodel \* model, enum starpu\_perfmodel\_archtype arch, unsigned nimpl, char \* parameter, uint32\_t \* footprint, FILE \* output )**

todo

**int starpu\_perfmodel\_print\_all ( struct starpu\_perfmodel \* model, char \* arch, char \* parameter, uint32\_t \* footprint, FILE \* output )**

todo

**void starpu\_bus\_print\_bandwidth ( FILE \* f )**

prints a matrix of bus bandwidths on `f`.

**void starpu\_bus\_print\_affinity ( FILE \* f )**

prints the affinity devices on `f`.

```
void starpu_perfmodel_update_history ( struct starpu_perfmodel * model, struct starpu_task * task, enum starpu_perfmodel_archtype arch, unsigned cpuid, unsigned nimpl, double measured )
```

This feeds the performance model model with an explicit measurement measured (in  $\mu\text{s}$ ), in addition to measurements done by StarPU itself. This can be useful when the application already has an existing set of measurements done in good conditions, that StarPU could benefit from instead of doing on-line measurements. And example of use can be seen in [Performance Model Example](#).

```
double starpu_transfer_bandwidth ( unsigned src_node, unsigned dst_node )
```

Return the bandwidth of data transfer between two memory nodes

```
double starpu_transfer_latency ( unsigned src_node, unsigned dst_node )
```

Return the latency of data transfer between two memory nodes

```
double starpu_transfer_predict ( unsigned src_node, unsigned dst_node, size_t size )
```

Return the estimated time to transfer a given size between two memory nodes.

## 17.16 Profiling

### Data Structures

- struct [starpu\\_profiling\\_task\\_info](#)
- struct [starpu\\_profiling\\_worker\\_info](#)
- struct [starpu\\_profiling\\_bus\\_info](#)

### Macros

- `#define STARPU_PROFILING_DISABLE`
- `#define STARPU_PROFILING_ENABLE`

### Functions

- int [starpu\\_profiling\\_status\\_set](#) (int status)
- int [starpu\\_profiling\\_status\\_get](#) (void)
- void [starpu\\_profiling\\_init](#) ()
- void [starpu\\_profiling\\_set\\_id](#) (int new\_id)
- int [starpu\\_profiling\\_worker\\_get\\_info](#) (int workerid, struct [starpu\\_profiling\\_worker\\_info](#) \*worker\_info)
- int [starpu\\_bus\\_get\\_profiling\\_info](#) (int busid, struct [starpu\\_profiling\\_bus\\_info](#) \*bus\_info)
- int [starpu\\_bus\\_get\\_count](#) (void)
- int [starpu\\_bus\\_get\\_id](#) (int src, int dst)
- int [starpu\\_bus\\_get\\_src](#) (int busid)
- int [starpu\\_bus\\_get\\_dst](#) (int busid)
- double [starpu\\_timing\\_timespec\\_delay\\_us](#) (struct timespec \*start, struct timespec \*end)
- double [starpu\\_timing\\_timespec\\_to\\_us](#) (struct timespec \*ts)
- void [starpu\\_profiling\\_bus\\_helper\\_display\\_summary](#) (void)
- void [starpu\\_profiling\\_worker\\_helper\\_display\\_summary](#) (void)
- void [starpu\\_data\\_display\\_memory\\_stats](#) ()

#### 17.16.1 Detailed Description

#### 17.16.2 Data Structure Documentation

##### **struct starpu\_profiling\_task\_info**

This structure contains information about the execution of a task. It is accessible from the field [starpu\\_task::profiling\\_info](#) if profiling was enabled.

**Data Fields**

struct timespec	submit_time	Date of task submission (relative to the initialization of StarPU).
struct timespec	push_start_time	Time when the task was submitted to the scheduler.
struct timespec	push_end_time	Time when the scheduler finished with the task submission.
struct timespec	pop_start_time	Time when the scheduler started to be requested for a task, and eventually gave that task.
struct timespec	pop_end_time	Time when the scheduler finished providing the task for execution.
struct timespec	acquire_data_start_time	Time when the worker started fetching input data.
struct timespec	acquire_data_end_time	Time when the worker finished fetching input data.
struct timespec	start_time	Date of task execution beginning (relative to the initialization of StarPU).
struct timespec	end_time	Date of task execution termination (relative to the initialization of StarPU).
struct timespec	release_data_start_time	Time when the worker started releasing data.
struct timespec	release_data_end_time	Time when the worker finished releasing data.
struct timespec	callback_start_time	Time when the worker started the application callback for the task.
struct timespec	callback_end_time	Time when the worker finished the application callback for the task.
int	workerid	Identifier of the worker which has executed the task.
uint64_t	used_cycles	Number of cycles used by the task, only available in the MoviSim
uint64_t	stall_cycles	Number of cycles stalled within the task, only available in the MoviSim
double	power_consumed	Power consumed by the task, only available in the MoviSim

**struct starpu\_profiling\_worker\_info**

This structure contains the profiling information associated to a worker. The timing is provided since the previous call to [starpu\\_profiling\\_worker\\_get\\_info\(\)](#)

**Data Fields**

struct timespec	start_time	Starting date for the reported profiling measurements.
struct timespec	total_time	Duration of the profiling measurement interval.
struct timespec	executing_time	Time spent by the worker to execute tasks during the profiling measurement interval.
struct timespec	sleeping_time	Time spent idling by the worker during the profiling measurement interval.
int	executed_tasks	Number of tasks executed by the worker during the profiling measurement interval.
uint64_t	used_cycles	Number of cycles used by the worker, only available in the MoviSim
uint64_t	stall_cycles	Number of cycles stalled within the worker, only available in the MoviSim
double	power_consumed	Power consumed by the worker, only available in the MoviSim

**struct starpu\_profiling\_bus\_info**

todo

**Data Fields**

struct timespec	start_time	Time of bus profiling startup.
struct timespec	total_time	Total time of bus profiling.

int long long	transferred_bytes	Number of bytes transferred during profiling.
int	transfer_count	Number of transfers during profiling.

### 17.16.3 Macro Definition Documentation

#### **STARPU\_PROFILING\_DISABLE**

This value is used when calling the function [starpu\\_profiling\\_status\\_set\(\)](#) to disable profiling.

#### **STARPU\_PROFILING\_ENABLE**

This value is used when calling the function [starpu\\_profiling\\_status\\_set\(\)](#) to enable profiling.

### 17.16.4 Function Documentation

#### **int starpu\_profiling\_status\_set ( int status )**

This function sets the profiling status. Profiling is activated by passing [STARPU\\_PROFILING\\_ENABLE](#) in status. Passing [STARPU\\_PROFILING\\_DISABLE](#) disables profiling. Calling this function resets all profiling measurements. When profiling is enabled, the field [starpu\\_task::profiling\\_info](#) points to a valid structure [starpu\\_profiling\\_task\\_info](#) containing information about the execution of the task. Negative return values indicate an error, otherwise the previous status is returned.

#### **int starpu\_profiling\_status\_get ( void )**

Return the current profiling status or a negative value in case there was an error.

#### **int starpu\_profiling\_init ( void )**

This function resets performance counters and enable profiling if the environment variable [STARPU\\_PROFILING](#) is set to a positive value.

#### **void starpu\_profiling\_set\_id ( int new\_id )**

This function sets the ID used for profiling trace filename. It needs to be called before [starpu\\_init\(\)](#).

#### **int starpu\_profiling\_worker\_get\_info ( int workerid, struct starpu\_profiling\_worker\_info \* worker\_info )**

Get the profiling info associated to the worker identified by `workerid`, and reset the profiling measurements. If the argument `worker_info` is NULL, only reset the counters associated to worker `workerid`. Upon successful completion, this function returns 0. Otherwise, a negative value is returned.

#### **int starpu\_bus\_get\_profiling\_info ( int busid, struct starpu\_profiling\_bus\_info \* bus\_info )**

todo

#### **int starpu\_bus\_get\_count ( void )**

Return the number of buses in the machine

#### **int starpu\_bus\_get\_id ( int src, int dst )**

Return the identifier of the bus between `src` and `dst`

#### **int starpu\_bus\_get\_src ( int busid )**

Return the source point of bus `busid`

**int starpu\_bus\_get\_dst ( int busid )**

Return the destination point of bus busid

**double starpu\_timing\_timespec\_delay\_us ( struct timespec \* start, struct timespec \* end )**

Returns the time elapsed between start and end in microseconds.

**double starpu\_timing\_timespec\_to\_us ( struct timespec \* ts )**

Converts the given timespec ts into microseconds

**void starpu\_profiling\_bus\_helper\_display\_summary ( void )**

Displays statistics about the bus on stderr. if the environment variable `STARPU_BUS_STATS` is defined. The function is called automatically by `starpu_shutdown()`.

**void starpu\_profiling\_worker\_helper\_display\_summary ( void )**

Displays statistics about the workers on stderr if the environment variable `STARPU_WORKER_STATS` is defined. The function is called automatically by `starpu_shutdown()`.

**void starpu\_data\_display\_memory\_stats ( )**

Display statistics about the current data handles registered within StarPU. StarPU must have been configured with the configure option `--enable-memory-stats` (see [Memory Feedback](#)).

## 17.17 Theoretical Lower Bound on Execution Time

Compute theoretical upper computation efficiency bound corresponding to some actual execution.

### Functions

- `void starpu_bound_start (int deps, int prio)`
- `void starpu_bound_stop (void)`
- `void starpu_bound_print_dot (FILE *output)`
- `void starpu_bound_compute (double *res, double *integer_res, int integer)`
- `void starpu_bound_print_lp (FILE *output)`
- `void starpu_bound_print_mps (FILE *output)`
- `void starpu_bound_print (FILE *output, int integer)`

#### 17.17.1 Detailed Description

Compute theoretical upper computation efficiency bound corresponding to some actual execution.

#### 17.17.2 Function Documentation

**void starpu\_bound\_start ( int deps, int prio )**

Start recording tasks (resets stats). deps tells whether dependencies should be recorded too (this is quite expensive)

**void starpu\_bound\_stop ( void )**

Stop recording tasks

**void starpu\_bound\_print\_dot ( FILE \* output )**

Print the DAG that was recorded

```
void starpu_bound_compute( double * res, double * integer_res, int integer )
```

Get theoretical upper bound (in ms) (needs glpk support detected by configure script). It returns 0 if some performance models are not calibrated.

```
void starpu_bound_print_lp( FILE * output )
```

Emit the Linear Programming system on `output` for the recorded tasks, in the lp format

```
void starpu_bound_print_mps( FILE * output )
```

Emit the Linear Programming system on `output` for the recorded tasks, in the mps format

```
void starpu_bound_print( FILE * output, int integer )
```

Emit statistics of actual execution vs theoretical upper bound. `integer` permits to choose between integer solving (which takes a long time but is correct), and relaxed solving (which provides an approximate solution).

## 17.18 CUDA Extensions

### Macros

- `#define STARPU_USE_CUDA`
- `#define STARPU_MAXCUDADEVS`
- `#define STARPU_CUDA_REPORT_ERROR(status)`
- `#define STARPU_CUBLAS_REPORT_ERROR(status)`

### Functions

- `cudaStream_t starpu_cuda_get_local_stream(void)`
- `struct cudaDeviceProp * starpu_cuda_get_device_properties(unsigned workerid)`
- `void starpu_cuda_report_error(const char *func, const char *file, int line, cudaError_t status)`
- `int starpu_cuda_copy_async_sync(void *src_ptr, unsigned src_node, void *dst_ptr, unsigned dst_node, size_t ssize, cudaStream_t stream, enum cudaMemcpyKind kind)`
- `void starpu_cuda_set_device(unsigned devid)`
- `void starpu_cublas_init(void)`
- `void starpu_cublas_shutdown(void)`
- `void starpu_cublas_report_error(const char *func, const char *file, int line, cublasStatus status)`

#### 17.18.1 Detailed Description

#### 17.18.2 Macro Definition Documentation

##### `#define STARPU_USE_CUDA`

This macro is defined when StarPU has been installed with CUDA support. It should be used in your code to detect the availability of CUDA as shown in [Full source code for the 'Scaling a Vector' example](#).

##### `#define STARPU_MAXCUDADEVS`

This macro defines the maximum number of CUDA devices that are supported by StarPU.

##### `#define STARPU_CUDA_REPORT_ERROR( status )`

Calls `starpu_cuda_report_error()`, passing the current function, file and line position.

##### `#define STARPU_CUBLAS_REPORT_ERROR( status )`

Calls `starpu_cublas_report_error()`, passing the current function, file and line position.

### 17.18.3 Function Documentation

**cudaStream\_t starpu\_cuda\_get\_local\_stream ( void )**

This function gets the current worker's CUDA stream. StarPU provides a stream for every CUDA device controlled by StarPU. This function is only provided for convenience so that programmers can easily use asynchronous operations within codelets without having to create a stream by hand. Note that the application is not forced to use the stream provided by `starpu_cuda_get_local_stream()` and may also create its own streams. Synchronizing with `cudaThreadSynchronize()` is allowed, but will reduce the likelihood of having all transfers overlapped.

**const struct cudaDeviceProp \* starpu\_cuda\_get\_device\_properties ( unsigned workerid )**

This function returns a pointer to device properties for worker `workerid` (assumed to be a CUDA worker).

**void starpu\_cuda\_report\_error ( const char \* func, const char \* file, int line, cudaError\_t status )**

Report a CUDA error.

**int starpu\_cuda\_copy\_async\_sync ( void \* src\_ptr, unsigned src\_node, void \* dst\_ptr, unsigned dst\_node, size\_t ssize, cudaStream\_t stream, enum cudaMemcpyKind kind )**

Copy `ssize` bytes from the pointer `src_ptr` on `src_node` to the pointer `dst_ptr` on `dst_node`. The function first tries to copy the data asynchronous (unless `stream` is `NULL`). If the asynchronous copy fails or if `stream` is `NULL`, it copies the data synchronously. The function returns `-EAGAIN` if the asynchronous launch was successful. It returns `0` if the synchronous copy was successful, or fails otherwise.

**void starpu\_cuda\_set\_device ( unsigned devid )**

Calls `cudaSetDevice(devid)` or `cudaGLSetGLDevice(devid)`, according to whether `devid` is among the field `starpu_conf::cuda_opengl_interoperability`.

**void starpu\_cublas\_init ( void )**

This function initializes CUBLAS on every CUDA device. The CUBLAS library must be initialized prior to any CUBLAS call. Calling `starpu_cublas_init()` will initialize CUBLAS on every CUDA device controlled by StarPU. This call blocks until CUBLAS has been properly initialized on every device.

**void starpu\_cublas\_shutdown ( void )**

This function synchronously deinitializes the CUBLAS library on every CUDA device.

**void starpu\_cublas\_report\_error ( const char \* func, const char \* file, int line, cublasStatus status )**

Report a cublas error.

## 17.19 OpenCL Extensions

### Data Structures

- struct `starpu_opencl_program`

### Macros

- `#define STARPU_USE_OPENCL`
- `#define STARPU_MAXOPENCLDEVS`
- `#define STARPU_OPENCL_DATADIR`

## Writing OpenCL kernels

- void `starpu_opencl_get_context` (int devid, cl\_context \*context)
- void `starpu_opencl_get_device` (int devid, cl\_device\_id \*device)
- void `starpu_opencl_get_queue` (int devid, cl\_command\_queue \*queue)
- void `starpu_opencl_get_current_context` (cl\_context \*context)
- void `starpu_opencl_get_current_queue` (cl\_command\_queue \*queue)
- int `starpu_opencl_set_kernel_args` (cl\_int \*err, cl\_kernel \*kernel,...)

## Compiling OpenCL kernels

Source codes for OpenCL kernels can be stored in a file or in a string. StarPU provides functions to build the program executable for each available OpenCL device as a `cl_program` object. This program executable can then be loaded within a specific queue as explained in the next section. These are only helpers, Applications can also fill a `starpu_opencl_program` array by hand for more advanced use (e.g. different programs on the different OpenCL devices, for relocation purpose for instance).

- int `starpu_opencl_load_opencl_from_file` (const char \*source\_file\_name, struct `starpu_opencl_program` \*opencl\_programs, const char \*build\_options)
- int `starpu_opencl_load_opencl_from_string` (const char \*opencl\_program\_source, struct `starpu_opencl_program` \*opencl\_programs, const char \*build\_options)
- int `starpu_opencl_unload_opencl` (struct `starpu_opencl_program` \*opencl\_programs)
- void `starpu_opencl_load_program_source` (const char \*source\_file\_name, char \*located\_file\_name, char \*located\_dir\_name, char \*opencl\_program\_source)
- int `starpu_opencl_compile_opencl_from_file` (const char \*source\_file\_name, const char \*build\_options)
- int `starpu_opencl_compile_opencl_from_string` (const char \*opencl\_program\_source, const char \*file\_name, const char \*build\_options)
- int `starpu_opencl_load_binary_opencl` (const char \*kernel\_id, struct `starpu_opencl_program` \*opencl\_programs)

## Loading OpenCL kernels

- int `starpu_opencl_load_kernel` (cl\_kernel \*kernel, cl\_command\_queue \*queue, struct `starpu_opencl_program` \*opencl\_programs, const char \*kernel\_name, int devid)
- int `starpu_opencl_release_kernel` (cl\_kernel kernel)

## OpenCL statistics

- int `starpu_opencl_collect_stats` (cl\_event event)

## OpenCL utilities

- #define `STARPU_OPENCL_DISPLAY_ERROR`(status)
- #define `STARPU_OPENCL_REPORT_ERROR`(status)
- #define `STARPU_OPENCL_REPORT_ERROR_WITH_MSG`(msg, status)
- const char \* `starpu_opencl_error_string` (cl\_int status)
- void `starpu_opencl_display_error` (const char \*func, const char \*file, int line, const char \*msg, cl\_int status)
- static \_\_starpu\_inline void `starpu_opencl_report_error` (const char \*func, const char \*file, int line, const char \*msg, cl\_int status)
- cl\_int `starpu_opencl_allocate_memory` (cl\_mem \*addr, size\_t size, cl\_mem\_flags flags)
- cl\_int `starpu_opencl_copy_ram_to_opencl` (void \*ptr, unsigned src\_node, cl\_mem buffer, unsigned dst\_node, size\_t size, size\_t offset, cl\_event \*event, int \*ret)
- cl\_int `starpu_opencl_copy_opencl_to_ram` (cl\_mem buffer, unsigned src\_node, void \*ptr, unsigned dst\_node, size\_t size, size\_t offset, cl\_event \*event, int \*ret)
- cl\_int `starpu_opencl_copy_opencl_to_opencl` (cl\_mem src, unsigned src\_node, size\_t src\_offset, cl\_mem dst, unsigned dst\_node, size\_t dst\_offset, size\_t size, cl\_event \*event, int \*ret)
- cl\_int `starpu_opencl_copy_async_sync` (uintptr\_t src, size\_t src\_offset, unsigned src\_node, uintptr\_t dst, size\_t dst\_offset, unsigned dst\_node, size\_t size, cl\_event \*event)

**17.19.1 Detailed Description****17.19.2 Data Structure Documentation****struct starpu\_opencl\_program**

Stores the OpenCL programs as compiled for the different OpenCL devices.

**Data Fields**

cl_program	programs[STAR-PU_MAXOPENCLDEVS]	Stores each program for each OpenCL device.
------------	---------------------------------	---

**17.19.3 Macro Definition Documentation****#define STARPU\_USE\_OPENCL**

This macro is defined when StarPU has been installed with OpenCL support. It should be used in your code to detect the availability of OpenCL as shown in [Full source code for the 'Scaling a Vector' example](#).

**#define STARPU\_MAXOPENCLDEVS**

This macro defines the maximum number of OpenCL devices that are supported by StarPU.

**#define STARPU\_OPENCL\_DATADIR**

This macro defines the directory in which the OpenCL codelets of the applications provided with StarPU have been installed.

**#define STARPU\_OPENCL\_DISPLAY\_ERROR( *status* )**

Call the function [starpu\\_opencl\\_display\\_error\(\)](#) with the given error status, the current function name, current file and line number, and a empty message.

**#define STARPU\_OPENCL\_REPORT\_ERROR( *status* )**

Call the function [starpu\\_opencl\\_report\\_error\(\)](#) with the given error *status*, with the current function name, current file and line number, and a empty message.

**#define STARPU\_OPENCL\_REPORT\_ERROR\_WITH\_MSG( *msg*, *status* )**

Call the function [starpu\\_opencl\\_report\\_error\(\)](#) with the given *msg* and the given error *status*, with the current function name, current file and line number.

**17.19.4 Function Documentation****void starpu\_opencl\_get\_context ( int *devid*, cl\_context \* *context* )**

Places the OpenCL context of the device designated by *devid* into *context*.

**void starpu\_opencl\_get\_device ( int *devid*, cl\_device\_id \* *device* )**

Places the *cl\_device\_id* corresponding to *devid* in *device*.

**void starpu\_opencl\_get\_queue ( int *devid*, cl\_command\_queue \* *queue* )**

Places the command queue of the device designated by *devid* into *queue*.

**void starpu\_opencl\_get\_current\_context ( cl\_context \* *context* )**

Return the context of the current worker.

**void starpu\_opencl\_get\_current\_queue ( cl\_command\_queue \* *queue* )**

Return the computation kernel command queue of the current worker.

**int starpu\_opencl\_set\_kernel\_args ( cl\_int \* err, cl\_kernel \* kernel, ... )**

Sets the arguments of a given kernel. The list of arguments must be given as (size\_t size\_of\_the\_argument, cl\_mem \* pointer\_to\_the\_argument). The last argument must be 0. Returns the number of arguments that were successfully set. In case of failure, returns the id of the argument that could not be set and err is set to the error returned by OpenCL. Otherwise, returns the number of arguments that were set.

Here an example:

```
int n;
cl_int err;
cl_kernel kernel;
n = starpu_opencl_set_kernel_args(&err, 2, &kernel,
                                  sizeof(foo), &foo,
                                  sizeof(bar), &bar,
                                  0);
if (n != 2)
    fprintf(stderr, "Error : %d\n", err);
```

**int starpu\_opencl\_load\_opencl\_from\_file ( const char \* source\_file\_name, struct starpu\_opencl\_program \* opencl\_programs, const char \* build\_options )**

This function compiles an OpenCL source code stored in a file.

**int starpu\_opencl\_load\_opencl\_from\_string ( const char \* opencl\_program\_source, struct starpu\_opencl\_program \* opencl\_programs, const char \* build\_options )**

This function compiles an OpenCL source code stored in a string.

**int starpu\_opencl\_unload\_opencl ( struct starpu\_opencl\_program \* opencl\_programs )**

This function unloads an OpenCL compiled code.

**void starpu\_opencl\_load\_program\_source ( const char \* source\_file\_name, char \* located\_file\_name, char \* located\_dir\_name, char \* opencl\_program\_source )**

Store the contents of the file source\_file\_name in the buffer opencl\_program\_source. The file source\_file\_name can be located in the current directory, or in the directory specified by the environment variable **STARPU\_OPENCL\_PROGRAM\_DIR**, or in the directory share/starpu/opencl of the installation directory of StarPU, or in the source directory of StarPU. When the file is found, located\_file\_name is the full name of the file as it has been located on the system, located\_dir\_name the directory where it has been located. Otherwise, they are both set to the empty string.

**int starpu\_opencl\_compile\_opencl\_from\_file ( const char \* source\_file\_name, const char \* build\_options )**

Compile the OpenCL kernel stored in the file source\_file\_name with the given options build\_options and stores the result in the directory \$STARPU\_HOME/.starpu/opencl with the same filename as source\_file\_name. The compilation is done for every OpenCL device, and the filename is suffixed with the vendor id and the device id of the OpenCL device.

**int starpu\_opencl\_compile\_opencl\_from\_string ( const char \* opencl\_program\_source, const char \* file\_name, const char \* build\_options )**

Compile the OpenCL kernel in the string opencl\_program\_source with the given options build\_options and stores the result in the directory \$STARPU\_HOME/.starpu/opencl with the filename file\_name. The compilation is done for every OpenCL device, and the filename is suffixed with the vendor id and the device id of the OpenCL device.

**int starpu\_opencl\_load\_binary\_opencl ( const char \* kernel\_id, struct starpu\_opencl\_program \* opencl\_programs )**

Compile the binary OpenCL kernel identified with kernel\_id. For every OpenCL device, the binary OpenCL kernel will be loaded from the file \$STARPU\_HOME/.starpu/opencl/<kernel\_id>. <device\_type>.vendor\_id\_<vendor\_id>\_device\_id\_<device\_id>.

```
int starpu_opencl_load_kernel ( cl_kernel * kernel, cl_command_queue * queue, struct starpu_opencl_program * opencl_programs, const char * kernel_name, int devid )
```

Create a kernel `kernel` for device `devid`, on its computation command queue returned in `queue`, using program `opencl_programs` and name `kernel_name`.

```
int starpu_opencl_release_kernel ( cl_kernel kernel )
```

Release the given `kernel`, to be called after kernel execution.

```
int starpu_opencl_collect_stats ( cl_event event )
```

This function allows to collect statistics on a kernel execution. After termination of the kernels, the OpenCL codelet should call this function to pass it the even returned by `clEnqueueNDRangeKernel`, to let StarPU collect statistics about the kernel execution (used cycles, consumed power).

```
const char * starpu_opencl_error_string ( cl_int status )
```

Return the error message in English corresponding to `status`, an OpenCL error code.

```
void starpu_opencl_display_error ( const char * func, const char * file, int line, const char * msg, cl_int status )
```

Given a valid error status, prints the corresponding error message on `stdout`, along with the given function name `func`, the given filename `file`, the given line number `line` and the given message `msg`.

```
void starpu_opencl_report_error ( const char * func, const char * file, int line, const char * msg, cl_int status ) [static]
```

Call the function [starpu\\_opencl\\_display\\_error\(\)](#) and abort.

```
cl_int starpu_opencl_allocate_memory ( cl_mem * addr, size_t size, cl_mem_flags flags )
```

Allocate `size` bytes of memory, stored in `addr`. `flags` must be a valid combination of `cl_mem_flags` values.

```
cl_int starpu_opencl_copy_ram_to_opencl ( void * ptr, unsigned src_node, cl_mem buffer, unsigned dst_node, size_t size, size_t offset, cl_event * event, int * ret )
```

Copy `size` bytes from the given `ptr` on RAM `src_node` to the given `buffer` on OpenCL `dst_node`. `offset` is the offset, in bytes, in `buffer`. If `event` is `NULL`, the copy is synchronous, i.e the queue is synchronised before returning. If not `NULL`, `event` can be used after the call to wait for this particular copy to complete. This function returns `CL_SUCCESS` if the copy was successful, or a valid OpenCL error code otherwise. The integer pointed to by `ret` is set to `-EAGAIN` if the asynchronous launch was successful, or to 0 if `event` was `NULL`.

```
cl_int starpu_opencl_copy_opencl_to_ram ( cl_mem buffer, unsigned src_node, void * ptr, unsigned dst_node, size_t size, size_t offset, cl_event * event, int * ret )
```

Copy `size` bytes asynchronously from the given `buffer` on OpenCL `src_node` to the given `ptr` on RAM `dst_node`. `offset` is the offset, in bytes, in `buffer`. If `event` is `NULL`, the copy is synchronous, i.e the queue is synchronised before returning. If not `NULL`, `event` can be used after the call to wait for this particular copy to complete. This function returns `CL_SUCCESS` if the copy was successful, or a valid OpenCL error code otherwise. The integer pointed to by `ret` is set to `-EAGAIN` if the asynchronous launch was successful, or to 0 if `event` was `NULL`.

```
cl_int starpu_opencl_copy_opencl_to_opencl ( cl_mem src, unsigned src_node, size_t src_offset, cl_mem dst, unsigned dst_node, size_t dst_offset, size_t size, cl_event * event, int * ret )
```

Copy `size` bytes asynchronously from byte offset `src_offset` of `src` on OpenCL `src_node` to byte offset `dst_offset` of `dst` on OpenCL `dst_node`. If `event` is NULL, the copy is synchronous, i.e. the queue is synchronised before returning. If not NULL, `event` can be used after the call to wait for this particular copy to complete. This function returns `CL_SUCCESS` if the copy was successful, or a valid OpenCL error code otherwise. The integer pointed to by `ret` is set to `-EAGAIN` if the asynchronous launch was successful, or to 0 if `event` was NULL.

```
cl_int starpu_opencl_copy_async_sync ( uintptr_t src, size_t src_offset, unsigned src_node, uintptr_t dst, size_t dst_offset, unsigned dst_node, size_t size, cl_event * event )
```

Copy `size` bytes from byte offset `src_offset` of `src` on `src_node` to byte offset `dst_offset` of `dst` on `dst_node`. If `event` is NULL, the copy is synchronous, i.e. the queue is synchronised before returning. If not NULL, `event` can be used after the call to wait for this particular copy to complete. The function returns `-EAGAIN` if the asynchronous launch was successfull. It returns 0 if the synchronous copy was successful, or fails otherwise.

## 17.20 Miscellaneous Helpers

### Functions

- int `starpu_data_cpy` (`starpu_data_handle_t dst_handle, starpu_data_handle_t src_handle, int asynchronous, void(*callback_func)(void *), void *callback_arg)`)
- void `starpu_execute_on_each_worker` (`void(*func)(void *), void *arg, uint32_t where`)

#### 17.20.1 Detailed Description

#### 17.20.2 Function Documentation

```
int starpu_data_cpy ( starpu_data_handle_t dst_handle, starpu_data_handle_t src_handle, int asynchronous, void(*)(void *) callback_func, void * callback_arg )
```

Copy the content of `src_handle` into `dst_handle`. The parameter `asynchronous` indicates whether the function should block or not. In the case of an asynchronous call, it is possible to synchronize with the termination of this operation either by the means of implicit dependencies (if enabled) or by calling `starpu_task_wait_for_all()`. If `callback_func` is not NULL, this callback function is executed after the handle has been copied, and it is given the pointer `callback_arg` as argument.

```
void starpu_execute_on_each_worker ( void(*)(void *) func, void * arg, uint32_t where )
```

This function executes the given function on a subset of workers. When calling this method, the offloaded function `func` is executed by every StarPU worker that may execute the function. The argument `arg` is passed to the offloaded function. The argument `where` specifies on which types of processing units the function should be executed. Similarly to the field `starpu_codelet::where`, it is possible to specify that the function should be executed on every CUDA device and every CPU by passing `STARPU_CPU|STARPU_CUDA`. This function blocks until the function has been executed on every appropriate processing units, so that it may not be called from a callback function for instance.

## 17.21 FxT Support

### Data Structures

- struct `starpu_fxt_codelet_event`
- struct `starpu_fxt_options`

## Functions

- void `starpu_fxt_options_init` (struct `starpu_fxt_options` \*options)
- void `starpu_fxt_generate_trace` (struct `starpu_fxt_options` \*options)
- void `starpu_fxt_start_profiling` (void)
- void `starpu_fxt_stop_profiling` (void)

### 17.21.1 Detailed Description

### 17.21.2 Data Structure Documentation

#### `struct starpu_fxt_codelet_event`

todo  
Data Fields

char	symbol[256]	name of the codelet
int	workerid	
enum <code>starpu_perfmodel_archtype</code>	archtype	
uint32.t	hash	
size.t	size	
float	time	

#### `struct starpu_fxt_options`

todo  
Data Fields

unsigned	per_task.colour	
unsigned	no_counter	
unsigned	no_bus	
unsigned	ninputfiles	
char *	filenames[STAR-PU_FXT_MAX_FILES]	
char *	out_paje_path	
char *	distrib_time_path	
char *	activity_path	
char *	dag_path	
char *	file_prefix	In case we are going to gather multiple traces (e.g in the case of MPI processes), we may need to prefix the name of the containers.
uint64.t	file_offset	In case we are going to gather multiple traces (e.g in the case of MPI processes), we may need to prefix the name of the containers.
int	file_rank	In case we are going to gather multiple traces (e.g in the case of MPI processes), we may need to prefix the name of the containers.
char	worker_names[STAR-PU_NMAXWORKERS][256]	Output parameters

enum <code>starpu_perfmodel_archtype</code>	<code>worker_archtypes[S-TARPU_NMAX_WORKERS]</code>	Output parameters
int	<code>nworkers</code>	Output parameters
struct <code>starpu_fxt_codelet_event</code> **	<code>dumped_codelets</code>	In case we want to dump the list of codelets to an external tool
long	<code>dumped_codelets_count</code>	In case we want to dump the list of codelets to an external tool

### 17.21.3 Function Documentation

`void starpu_fxt_options_init( struct starpu_fxt_options * options )`

todo

`void starpu_fxt_generate_trace( struct starpu_fxt_options * options )`

todo

`void starpu_fxt_start_profiling( void )`

Start recording the trace. The trace is by default started from `starpu_init()` call, but can be paused by using `starpu_fxt_stop_profiling()`, in which case `starpu_fxt_start_profiling()` should be called to resume recording events.

`void starpu_fxt_stop_profiling( void )`

Stop recording the trace. The trace is by default stopped when calling `starpu_shutdown()`. `starpu_fxt_stop_profiling()` can however be used to stop it earlier. `starpu_fxt_start_profiling()` can then be called to start recording it again, etc.

## 17.22 FFT Support

### Functions

- `void * starpufft_malloc( size_t n )`
- `void starpufft_free( void *p )`
- `starpufft_plan starpufft_plan_dft_1d( int n, int sign, unsigned flags )`
- `starpufft_plan starpufft_plan_dft_2d( int n, int m, int sign, unsigned flags )`
- `struct starpu_task * starpufft_start( starpufft_plan p, void *in, void *out )`
- `struct starpu_task * starpufft_start_handle( starpufft_plan p, starpu_data_handle_t in, starpu_data_handle_t out )`
- `int starpufft_execute( starpufft_plan p, void *in, void *out )`
- `int starpufft_execute_handle( starpufft_plan p, starpu_data_handle_t in, starpu_data_handle_t out )`
- `void starpufft_cleanup( starpufft_plan p )`
- `void starpufft_destroy_plan( starpufft_plan p )`

### 17.22.1 Detailed Description

### 17.22.2 Function Documentation

`void * starpufft_malloc( size_t n )`

Allocates memory for `n` bytes. This is preferred over `malloc()`, since it allocates pinned memory, which allows overlapped transfers.

`void * starpufft_free( void * p )`

Release memory previously allocated.

**struct starpufft\_plan \* starpufft\_plan\_dft\_1d ( int n, int sign, unsigned flags )**

Initializes a plan for 1D FFT of size n. sign can be STARPUFFT\_FORWARD or STARPUFFT\_INVERSE. flags must be 0.

**struct starpufft\_plan \* starpufft\_plan\_dft\_2d ( int n, int m, int sign, unsigned flags )**

Initializes a plan for 2D FFT of size (n, m). sign can be STARPUFFT\_FORWARD or STARPUFFT\_INVERSE. flags must be 0.

**struct starpu\_task \* starpufft\_start ( starpufft\_plan p, void \* in, void \* out )**

Start an FFT previously planned as p, using in and out as input and output. This only submits the task and does not wait for it. The application should call [starpufft\\_cleanup\(\)](#) to unregister the

**struct starpu\_task \* starpufft\_start\_handle ( starpufft\_plan p, starpu\_data\_handle\_t in, starpu\_data\_handle\_t out )**

Start an FFT previously planned as p, using data handles in and out as input and output (assumed to be vectors of elements of the expected types). This only submits the task and does not wait for it.

**void starpufft\_execute ( starpufft\_plan p, void \* in, void \* out )**

Execute an FFT previously planned as p, using in and out as input and output. This submits and waits for the task.

**void starpufft\_execute\_handle ( starpufft\_plan p, starpu\_data\_handle\_t in, starpu\_data\_handle\_t out )**

Execute an FFT previously planned as p, using data handles in and out as input and output (assumed to be vectors of elements of the expected types). This submits and waits for the task.

**void starpufft\_cleanup ( starpufft\_plan p )**

Releases data for plan p, in the [starpufft\\_start\(\)](#) case.

**void starpufft\_destroy\_plan ( starpufft\_plan p )**

Destroys plan p, i.e. release all CPU (fftw) and GPU (cufft) resources.

## 17.23 MPI Support

### Initialisation

- #define STARPU\_USE\_MPI
- int [starpu\\_mpi\\_init](#) (int \*argc, char \*\*\*argv, int initialize\_mpi)
- int [starpu\\_mpi\\_initialize](#) (void)
- int [starpu\\_mpi\\_initialize\\_extended](#) (int \*rank, int \*world\_size)
- int [starpu\\_mpi\\_shutdown](#) (void)
- void [starpu\\_mpi\\_comm\\_amounts\\_retrieve](#) (size\_t \*comm\_amounts)

### Communication

- int [starpu\\_mpi\\_send](#) ([starpu\\_data\\_handle\\_t](#) data\_handle, int dest, int mpi\_tag, MPI\_Comm comm)
- int [starpu\\_mpi\\_recv](#) ([starpu\\_data\\_handle\\_t](#) data\_handle, int source, int mpi\_tag, MPI\_Comm comm, MPI\_Status \*status)
- int [starpu\\_mpi\\_isend](#) ([starpu\\_data\\_handle\\_t](#) data\_handle, [starpu\\_mpi\\_req](#) \*req, int dest, int mpi\_tag, MPI\_Comm comm)
- int [starpu\\_mpi irecv](#) ([starpu\\_data\\_handle\\_t](#) data\_handle, [starpu\\_mpi\\_req](#) \*req, int source, int mpi\_tag, MPI\_Comm comm)

- int `starpu_mpi_isend_detached` (`starpu_data_handle_t` data\_handle, int dest, int mpi\_tag, MPI\_Comm comm, void(\*callback)(void \*), void \*arg)
- int `starpu_mpi_irecv_detached` (`starpu_data_handle_t` data\_handle, int source, int mpi\_tag, MPI\_Comm comm, void(\*callback)(void \*), void \*arg)
- int `starpu_mpi_wait` (`starpu_mpi_req` \*req, MPI\_Status \*status)
- int `starpu_mpi_test` (`starpu_mpi_req` \*req, int \*flag, MPI\_Status \*status)
- int `starpu_mpi_barrier` (MPI\_Comm comm)
- int `starpu_mpi_isend_detached_unlock_tag` (`starpu_data_handle_t` data\_handle, int dest, int mpi\_tag, MPI\_Comm comm, `starpu_tag_t` tag)
- int `starpu_mpi_irecv_detached_unlock_tag` (`starpu_data_handle_t` data\_handle, int source, int mpi\_tag, MPI\_Comm comm, `starpu_tag_t` tag)
- int `starpu_mpi_isend_array_detached_unlock_tag` (unsigned array\_size, `starpu_data_handle_t` \*data\_handle, int \*dest, int \*mpi\_tag, MPI\_Comm \*comm, `starpu_tag_t` tag)
- int `starpu_mpi_irecv_array_detached_unlock_tag` (unsigned array\_size, `starpu_data_handle_t` \*data\_handle, int \*source, int \*mpi\_tag, MPI\_Comm \*comm, `starpu_tag_t` tag)

## Communication Cache

- void `starpu_mpi_cache_flush` (MPI\_Comm comm, `starpu_data_handle_t` data\_handle)
- void `starpu_mpi_cache_flush_all_data` (MPI\_Comm comm)

## MPI Insert Task

- #define `STARPU_EXECUTE_ON_NODE`
- #define `STARPU_EXECUTE_ON_DATA`
- int `starpu_data_set_tag` (`starpu_data_handle_t` handle, int tag)
- int `starpu_data_get_tag` (`starpu_data_handle_t` handle)
- int `starpu_data_set_rank` (`starpu_data_handle_t` handle, int rank)
- int `starpu_data_get_rank` (`starpu_data_handle_t` handle)
- int `starpu_mpi_insert_task` (MPI\_Comm comm, struct `starpu_codelet` \*codelet,...)
- void `starpu_mpi_get_data_on_node` (MPI\_Comm comm, `starpu_data_handle_t` data\_handle, int node)
- void `starpu_mpi_get_data_on_node_detached` (MPI\_Comm comm, `starpu_data_handle_t` data\_handle, int node, void(\*callback)(void \*), void \*arg)

## Collective Operations

- void `starpu_mpi_redux_data` (MPI\_Comm comm, `starpu_data_handle_t` data\_handle)
- int `starpu_mpi_scatter_detached` (`starpu_data_handle_t` \*data\_handles, int count, int root, MPI\_Comm comm, void(\*scallback)(void \*), void \*sarg, void(\*rcallback)(void \*), void \*rarg)
- int `starpu_mpi_gather_detached` (`starpu_data_handle_t` \*data\_handles, int count, int root, MPI\_Comm comm, void(\*scallback)(void \*), void \*sarg, void(\*rcallback)(void \*), void \*rarg)

### 17.23.1 Detailed Description

### 17.23.2 Macro Definition Documentation

#### #define `STARPU_USE_MPI`

This macro is defined when StarPU has been installed with MPI support. It should be used in your code to detect the availability of MPI.

#### #define `STARPU_EXECUTE_ON_NODE`

this macro is used when calling `starpu_mpi_insert_task()`, and must be followed by a integer value which specified the node on which to execute the codelet.

```
#define STARPU_EXECUTE_ON_DATA
```

this macro is used when calling `starpu_mpi_insert_task()`, and must be followed by a data handle to specify that the node owning the given data will execute the codelet.

### 17.23.3 Function Documentation

```
int starpu_mpi_init( int * argc, char *** argv, int initialize_mpi )
```

Initializes the starpumpi library. `initialize_mpi` indicates if MPI should be initialized or not by StarPU. If the value is not 0, MPI will be initialized by calling `MPI_Init_Thread(argc, argv, MPI_THREAD_SERIALIZED, ...)`.

```
int starpu_mpi_initialize( void )
```

**Deprecated** This function has been made deprecated. One should use instead the function `starpu_mpi_init()`. This function does not call `MPI_Init()`, it should be called beforehand.

```
int starpu_mpi_initialize_extended( int * rank, int * world_size )
```

**Deprecated** This function has been made deprecated. One should use instead the function `starpu_mpi_init()`. MPI will be initialized by starpumpi by calling `MPI_Init_Thread(argc, argv, MPI_THREAD_SERIALIZED, ...)`.

```
int starpu_mpi_shutdown( void )
```

Cleans the starpumpi library. This must be called between calling `starpu_mpi` functions and `starpu_shutdown()`. `MPI_Finalize()` will be called if StarPU-MPI has been initialized by `starpu_mpi_init()`.

```
void starpu_mpi_comm_amounts_retrieve( size_t * comm_amounts )
```

Retrieve the current amount of communications from the current node in the array `comm_amounts` which must have a size greater or equal to the world size. Communications statistics must be enabled (see `STARPU_COMM_STATS`).

```
int starpu_mpi_send( starpu_data_handle_t data_handle, int dest, int mpi_tag, MPI_Comm comm )
```

Performs a standard-mode, blocking send of `data_handle` to the node `dest` using the message tag `mpi_tag` within the communicator `comm`.

```
int starpu_mpi_recv( starpu_data_handle_t data_handle, int source, int mpi_tag, MPI_Comm comm, MPI_Status * status )
```

Performs a standard-mode, blocking receive in `data_handle` from the node `source` using the message tag `mpi_tag` within the communicator `comm`.

```
int starpu_mpi_isend( starpu_data_handle_t data_handle, starpu_mpi_req * req, int dest, int mpi_tag, MPI_Comm comm )
```

Posts a standard-mode, non blocking send of `data_handle` to the node `dest` using the message tag `mpi_tag` within the communicator `comm`. After the call, the pointer to the request `req` can be used to test or to wait for the completion of the communication.

```
int starpu_mpi_irecv( starpu_data_handle_t data_handle, starpu_mpi_req * req, int source, int mpi_tag, MPI_Comm comm )
```

Posts a nonblocking receive in `data_handle` from the node `source` using the message tag `mpi_tag` within the communicator `comm`. After the call, the pointer to the request `req` can be used to test or to wait for the completion of the communication.

```
int starpu_mpi_isend_detached ( starpu_data_handle_t data_handle, int dest, int mpi_tag, MPI_Comm comm, void(*)(void *) callback, void * arg )
```

Posts a standard-mode, non blocking send of *data\_handle* to the node *dest* using the message tag *mpi\_tag* within the communicator *comm*. On completion, the *callback* function is called with the argument *arg*. Similarly to the pthread detached functionality, when a detached communication completes, its resources are automatically released back to the system, there is no need to test or to wait for the completion of the request.

```
int starpu_mpi_irecv_detached ( starpu_data_handle_t data_handle, int source, int mpi_tag, MPI_Comm comm, void(*)(void *) callback, void * arg )
```

Posts a nonblocking receive in *data\_handle* from the node *source* using the message tag *mpi\_tag* within the communicator *comm*. On completion, the *callback* function is called with the argument *arg*. Similarly to the pthread detached functionality, when a detached communication completes, its resources are automatically released back to the system, there is no need to test or to wait for the completion of the request.

```
int starpu_mpi_wait ( starpu_mpi_req * req, MPI_Status * status )
```

Returns when the operation identified by request *req* is complete.

```
int starpu_mpi_test ( starpu_mpi_req * req, int * flag, MPI_Status * status )
```

If the operation identified by *req* is complete, set *flag* to 1. The *status* object is set to contain information on the completed operation.

```
int starpu_mpi_barrier ( MPI_Comm comm )
```

Blocks the caller until all group members of the communicator *comm* have called it.

```
int starpu_mpi_isend_detached_unlock_tag ( starpu_data_handle_t data_handle, int dest, int mpi_tag, MPI_Comm comm, starpu_tag_t tag )
```

Posts a standard-mode, non blocking send of *data\_handle* to the node *dest* using the message tag *mpi\_tag* within the communicator *comm*. On completion, *tag* is unlocked.

```
int starpu_mpi_irecv_detached_unlock_tag ( starpu_data_handle_t data_handle, int source, int mpi_tag, MPI_Comm comm, starpu_tag_t tag )
```

Posts a nonblocking receive in *data\_handle* from the node *source* using the message tag *mpi\_tag* within the communicator *comm*. On completion, *tag* is unlocked.

```
int starpu_mpi_isend_array_detached_unlock_tag ( unsigned array_size, starpu_data_handle_t * data_handle, int * dest, int * mpi_tag, MPI_Comm * comm, starpu_tag_t tag )
```

Posts *array\_size* standard-mode, non blocking send. Each post sends the n-th data of the array *data\_handle* to the n-th node of the array *dest* using the n-th message tag of the array *mpi\_tag* within the n-th communicator of the array *comm*. On completion of the all the requests, *tag* is unlocked.

```
int starpu_mpi_irecv_array_detached_unlock_tag ( unsigned array_size, starpu_data_handle_t * data_handle, int * source, int * mpi_tag, MPI_Comm * comm, starpu_tag_t tag )
```

Posts *array\_size* nonblocking receive. Each post receives in the n-th data of the array *data\_handle* from the n-th node of the array *source* using the n-th message tag of the array *mpi\_tag* within the n-th communicator of the array *comm*. On completion of the all the requests, *tag* is unlocked.

```
void starpu_mpi_cache_flush ( MPI_Comm comm, starpu_data_handle_t data_handle )
```

Clear the send and receive communication cache for the data *data\_handle*. The function has to be called synchronously by all the MPI nodes. The function does nothing if the cache mechanism is disabled (see [STARPU\\_MPI\\_CACHE](#)).

**void starpu\_mpi\_cache\_flush\_all\_data ( MPI\_Comm comm )**

Clear the send and receive communication cache for all data. The function has to be called synchronously by all the MPI nodes. The function does nothing if the cache mechanism is disabled (see [STARPU\\_MPI\\_CACHE](#)).

**int starpu\_data\_set\_tag ( starpu\_data\_handle\_t handle, int tag )**

Tell StarPU-MPI which MPI tag to use when exchanging the data.

**int starpu\_data\_get\_tag ( starpu\_data\_handle\_t handle )**

Returns the MPI tag to be used when exchanging the data.

**int starpu\_data\_set\_rank ( starpu\_data\_handle\_t handle, int rank )**

Tell StarPU-MPI which MPI node "owns" a given data, that is, the node which will always keep an up-to-date value, and will by default execute tasks which write to it.

**int starpu\_data\_get\_rank ( starpu\_data\_handle\_t handle )**

Returns the last value set by [starpu\\_data\\_set\\_rank\(\)](#).

**int starpu\_mpi\_insert\_task ( MPI\_Comm comm, struct starpu\_codelet \* codelet, ... )**

Create and submit a task corresponding to codelet with the following arguments. The argument list must be zero-terminated.

The arguments following the codelets are the same types as for the function [starpu.insert.task\(\)](#). The extra argument [STARPU\\_EXECUTE\\_ON\\_NODE](#) followed by an integer allows to specify the MPI node to execute the codelet. It is also possible to specify that the node owning a specific data will execute the codelet, by using [STARPU\\_EXECUTE\\_ON\\_DATA](#) followed by a data handle.

The internal algorithm is as follows:

1. Find out which MPI node is going to execute the codelet.
  - If there is only one node owning data in [STARPU\\_W](#) mode, it will be selected;
  - If there is several nodes owning data in [STARPU\\_W](#) mode, the one selected will be the one having the least data in R mode so as to minimize the amount of data to be transferred;
  - The argument [STARPU\\_EXECUTE\\_ON\\_NODE](#) followed by an integer can be used to specify the node;
  - The argument [STARPU\\_EXECUTE\\_ON\\_DATA](#) followed by a data handle can be used to specify that the node owing the given data will execute the codelet.
2. Send and receive data as requested. Nodes owning data which need to be read by the task are sending them to the MPI node which will execute it. The latter receives them.
3. Execute the codelet. This is done by the MPI node selected in the 1st step of the algorithm.
4. If several MPI nodes own data to be written to, send written data back to their owners.

The algorithm also includes a communication cache mechanism that allows not to send data twice to the same MPI node, unless the data has been modified. The cache can be disabled (see [STARPU\\_MPI\\_CACHE](#)).

**void starpu\_mpi\_get\_data\_on\_node ( MPI\_Comm comm, starpu\_data\_handle\_t data\_handle, int node )**

Transfer data `data_handle` to MPI node `node`, sending it from its owner if needed. At least the target node and the owner have to call the function.

**void starpu\_mpi\_get\_data\_on\_node\_detached ( MPI\_Comm comm, starpu\_data\_handle\_t data\_handle, int node, void(\*)(void \*) callback, void \* arg )**

Transfer data `data_handle` to MPI node `node`, sending it from its owner if needed. At least the target node and the owner have to call the function. On reception, the `callback` function is called with the argument `arg`.

```
void starpu_mpi_redux_data ( MPI_Comm comm, starpu_data_handle_t data_handle )
```

Perform a reduction on the given data. All nodes send the data to its owner node which will perform a reduction.

```
int starpu_mpi_scatter_detached ( starpu_data_handle_t * data_handles, int count, int root, MPI_Comm comm, void(*)(void *) callback, void * sarg, void(*)(void *) rcallback, void * rarg )
```

Scatter data among processes of the communicator based on the ownership of the data. For each data of the array `data_handles`, the process `root` sends the data to the process owning this data. Processes receiving data must have valid data handles to receive them. On completion of the collective communication, the `callback` function is called with the argument `sarg` on the process `root`, the `rcallback` function is called with the argument `rarg` on any other process.

```
int starpu_mpi_gather_detached ( starpu_data_handle_t * data_handles, int count, int root, MPI_Comm comm, void(*)(void *) callback, void * sarg, void(*)(void *) rcallback, void * rarg )
```

Gather data from the different processes of the communicator onto the process `root`. Each process owning data handle in the array `data_handles` will send them to the process `root`. The process `root` must have valid data handles to receive the data. On completion of the collective communication, the `rcallback` function is called with the argument `rarg` on the process `root`, the `callback` function is called with the argument `sarg` on any other process.

## 17.24 Task Bundles

### Typedefs

- `typedef struct _starpu_task_bundle * starpu_task_bundle_t`

### Functions

- `void starpu_task_bundle_create (starpu_task_bundle_t *bundle)`
- `int starpu_task_bundle_insert (starpu_task_bundle_t bundle, struct starpu_task *task)`
- `int starpu_task_bundle_remove (starpu_task_bundle_t bundle, struct starpu_task *task)`
- `void starpu_task_bundle_close (starpu_task_bundle_t bundle)`
- `double starpu_task_bundle_expected_length (starpu_task_bundle_t bundle, enum starpu_perfmodel_archtype arch, unsigned nimpl)`
- `double starpu_task_bundle_expected_power (starpu_task_bundle_t bundle, enum starpu_perfmodel_archtype arch, unsigned nimpl)`
- `double starpu_task_bundle_expected_data_transfer_time (starpu_task_bundle_t bundle, unsigned memory_node)`

#### 17.24.1 Detailed Description

#### 17.24.2 Typedef Documentation

##### `starpu_task_bundle_t`

Opaque structure describing a list of tasks that should be scheduled on the same worker whenever it's possible. It must be considered as a hint given to the scheduler as there is no guarantee that they will be executed on the same worker.

#### 17.24.3 Function Documentation

```
void starpu_task_bundle_create ( starpu_task_bundle_t * bundle )
```

Factory function creating and initializing `bundle`, when the call returns, memory needed is allocated and `bundle` is ready to use.

```
int starpu_task_bundle_insert ( starpu_task_bundle_t bundle, struct starpu_task * task )
```

Insert task in bundle. Until task is removed from bundle its expected length and data transfer time will be considered along those of the other tasks of bundle. This function must not be called if bundle is already closed and/or task is already submitted. On success, it returns 0. There are two cases of error : if bundle is already closed it returns -EPERM, if task was already submitted it returns -EINVAL.

```
int starpu_task_bundle_remove ( starpu_task_bundle_t bundle, struct starpu_task * task )
```

Remove task from bundle. Of course task must have been previously inserted in bundle. This function must not be called if bundle is already closed and/or task is already submitted. Doing so would result in undefined behaviour. On success, it returns 0. If bundle is already closed it returns -ENOENT.

```
void starpu_task_bundle_close ( starpu_task_bundle_t bundle )
```

Inform the runtime that the user will not modify bundle anymore, it means no more inserting or removing task. Thus the runtime can destroy it when possible.

```
double starpu_task_bundle_expected_length ( starpu_task_bundle_t bundle, enum starpu_perfmodel_archtype arch, unsigned nimpl )
```

Return the expected duration of bundle in micro-seconds.

```
double starpu_task_bundle_expected_power ( starpu_task_bundle_t bundle, enum starpu_perfmodel_archtype arch, unsigned nimpl )
```

Return the expected power consumption of bundle in J.

```
double starpu_task_bundle_expected_data_transfer_time ( starpu_task_bundle_t bundle, unsigned memory_node )
```

Return the time (in micro-seconds) expected to transfer all data used within bundle.

## 17.25 Task Lists

### Data Structures

- struct [starpu\\_task\\_list](#)

### Functions

- static STARPU\_INLINE void [starpu\\_task\\_list\\_init](#) (struct [starpu\\_task\\_list](#) \*list)
- static STARPU\_INLINE void [starpu\\_task\\_list\\_push\\_front](#) (struct [starpu\\_task\\_list](#) \*list, struct [starpu\\_task](#) \*task)
- static STARPU\_INLINE void [starpu\\_task\\_list\\_push\\_back](#) (struct [starpu\\_task\\_list](#) \*list, struct [starpu\\_task](#) \*task)
- static STARPU\_INLINE struct  
    [starpu\\_task](#) \* [starpu\\_task\\_list\\_front](#) (struct [starpu\\_task\\_list](#) \*list)
- static STARPU\_INLINE struct  
    [starpu\\_task](#) \* [starpu\\_task\\_list\\_back](#) (struct [starpu\\_task\\_list](#) \*list)
- static STARPU\_INLINE int [starpu\\_task\\_list\\_empty](#) (struct [starpu\\_task\\_list](#) \*list)
- static STARPU\_INLINE void [starpu\\_task\\_list\\_erase](#) (struct [starpu\\_task\\_list](#) \*list, struct [starpu\\_task](#) \*task)
- static STARPU\_INLINE struct  
    [starpu\\_task](#) \* [starpu\\_task\\_list\\_pop\\_front](#) (struct [starpu\\_task\\_list](#) \*list)
- static STARPU\_INLINE struct  
    [starpu\\_task](#) \* [starpu\\_task\\_list\\_pop\\_back](#) (struct [starpu\\_task\\_list](#) \*list)
- static STARPU\_INLINE struct  
    [starpu\\_task](#) \* [starpu\\_task\\_list\\_begin](#) (struct [starpu\\_task\\_list](#) \*list)
- static STARPU\_INLINE struct  
    [starpu\\_task](#) \* [starpu\\_task\\_list\\_next](#) (struct [starpu\\_task](#) \*task)

### 17.25.1 Detailed Description

### 17.25.2 Data Structure Documentation

**struct starpu\_task\_list**

Stores a double-chained list of tasks

## Data Fields

struct starpu_task *	head	head of the list
struct starpu_task *	tail	tail of the list

## 17.25.3 Function Documentation

**void starpu\_task\_list\_init ( struct starpu\_task\_list \* *list* ) [static]**

Initialize a list structure

**void starpu\_task\_list\_push\_front ( struct starpu\_task\_list \* *list*, struct starpu\_task \* *task* ) [static]**

Push task at the front of list

**void starpu\_task\_list\_push\_back ( struct starpu\_task\_list \* *list*, struct starpu\_task \* *task* ) [static]**

Push task at the back of list

**struct starpu\_task \* starpu\_task\_list\_front ( struct starpu\_task\_list \* *list* ) [static]**

Get the front of list (without removing it)

**struct starpu\_task \* starpu\_task\_list\_back ( struct starpu\_task\_list \* *list* ) [static]**

Get the back of list (without removing it)

**int starpu\_task\_list\_empty ( struct starpu\_task\_list \* *list* ) [static]**

Test if list is empty

**void starpu\_task\_list\_erase ( struct starpu\_task\_list \* *list*, struct starpu\_task \* *task* ) [static]**

Remove task from list

**struct starpu\_task \* starpu\_task\_list\_pop\_front ( struct starpu\_task\_list \* *list* ) [static]**

Remove the element at the front of list

**struct starpu\_task \* starpu\_task\_list\_pop\_back ( struct starpu\_task\_list \* *list* ) [static]**

Remove the element at the back of list

**struct starpu\_task \* starpu\_task\_list\_begin ( struct starpu\_task\_list \* *list* ) [static]**

Get the first task of list.

**struct starpu\_task \* starpu\_task\_list\_next ( struct starpu\_task \* *task* ) [static]**

Get the next task of list. This is not erase-safe.

## 17.26 Parallel Tasks

### Functions

- int `starpu_combined_worker_get_size` (void)
- int `starpu_combined_worker_get_rank` (void)
- unsigned `starpu_combined_worker_get_count` (void)
- int `starpu_combined_worker_get_id` (void)
- int `starpu_combined_worker_assign_workerid` (int nworkers, int workerid\_array[])
- int `starpu_combined_worker_get_description` (int workerid, int \*worker\_size, int \*\*combined\_workerid)
- int `starpu_combined_worker_can_execute_task` (unsigned workerid, struct `starpu_task` \*task, unsigned nimpl)
- void `starpu_parallel_task_barrier_init` (struct `starpu_task` \*task, int workerid)

#### 17.26.1 Detailed Description

#### 17.26.2 Function Documentation

**int starpu\_combined\_worker\_get\_size ( void )**

Return the size of the current combined worker, i.e. the total number of cpus running the same task in the case of `STARPU_SPMD` parallel tasks, or the total number of threads that the task is allowed to start in the case of `STARPU_FORKJOIN` parallel tasks.

**int starpu\_combined\_worker\_get\_rank ( void )**

Return the rank of the current thread within the combined worker. Can only be used in `STARPU_FORKJOIN` parallel tasks, to know which part of the task to work on.

**unsigned starpu\_combined\_worker\_get\_count ( void )**

Return the number of different combined workers.

**int starpu\_combined\_worker\_get\_id ( void )**

Return the identifier of the current combined worker.

**int starpu\_combined\_worker\_assign\_workerid ( int nworkers, int workerid\_array[] )**

Register a new combined worker and get its identifier

**int starpu\_combined\_worker\_get\_description ( int workerid, int \* worker\_size, int \*\* combined\_workerid )**

Get the description of a combined worker

**int starpu\_combined\_worker\_can\_execute\_task ( unsigned workerid, struct starpu\_task \* task, unsigned nimpl )**

Variant of `starpu_worker_can_execute_task()` compatible with combined workers

**void starpu\_parallel\_task\_barrier\_init ( struct starpu\_task \* task, int workerid )**

Initialise the barrier for the parallel task, and dispatch the task between the different combined workers.

## 17.27 Running Drivers

### Functions

- int `starpu_driver_run` (struct `starpu_driver` \**d*)
- int `starpu_driver_init` (struct `starpu_driver` \**d*)
- int `starpu_driver_run_once` (struct `starpu_driver` \**d*)
- int `starpu_driver_deinit` (struct `starpu_driver` \**d*)
- void `starpu_drivers_request_termination` (void)

### 17.27.1 Detailed Description

### 17.27.2 Function Documentation

**int starpu\_driver\_run ( struct starpu\_driver \* *d* )**

Initialize the given driver, run it until it receives a request to terminate, deinitialize it and return 0 on success. It returns -EINVAL if *d*->type is not a valid StarPU device type (`STARPU_CPU_WORKER`, `STARPU_CUDA_WORKER` or `STARPU_OPENCL_WORKER`). This is the same as using the following functions: calling `starpu_driver_init()`, then calling `starpu_driver_run_once()` in a loop, and eventually `starpu_driver_deinit()`.

**int starpu\_driver\_init ( struct starpu\_driver \* *d* )**

Initialize the given driver. Returns 0 on success, -EINVAL if *d*->type is not a valid `starpu_worker_archtype`.

**int starpu\_driver\_run\_once ( struct starpu\_driver \* *d* )**

Run the driver once, then returns 0 on success, -EINVAL if *d*->type is not a valid `starpu_worker_archtype`.

**int starpu\_driver\_deinit ( struct starpu\_driver \* *d* )**

Deinitialize the given driver. Returns 0 on success, -EINVAL if *d*->type is not a valid `starpu_worker_archtype`.

**void starpu\_drivers\_request\_termination ( void )**

Notify all running drivers they should terminate.

## 17.28 Expert Mode

### Functions

- void `starpu_wake_all_blocked_workers` (void)
- int `starpu_progression_hook_register` (unsigned(\*func)(void \*arg), void \*arg)
- void `starpu_progression_hook_deregister` (int hook\_id)

### 17.28.1 Detailed Description

### 17.28.2 Function Documentation

**void starpu\_wake\_all\_blocked\_workers ( void )**

Wake all the workers, so they can inspect data requests and task submissions again.

**int starpu\_progression\_hook\_register ( unsigned(\*)(void \*arg) *func*, void \* *arg* )**

Register a progression hook, to be called when workers are idle.

**void starpu\_progression\_hook\_deregister ( int *hook\_id* )**

Unregister a given progression hook.

## 17.29 StarPU-Top Interface

### Data Structures

- struct `starpu_top_data`
- struct `starpu_top_param`

### Enumerations

- enum `starpu_top_data_type` { `STARPU_TOP_DATA_BOOLEAN`, `STARPU_TOP_DATA_INTEGER`, `STARPU_TOP_DATA_FLOAT` }
- enum `starpu_top_param_type` { `STARPU_TOP_PARAM_BOOLEAN`, `STARPU_TOP_PARAM_INTEGER`, `STARPU_TOP_PARAM_FLOAT`, `STARPU_TOP_PARAM_ENUM` }
- enum `starpu_top_message_type` {  
`TOP_TYPE_GO`, `TOP_TYPE_SET`, `TOP_TYPE_CONTINUE`, `TOP_TYPE_ENABLE`,  
`TOP_TYPE_DISABLE`, `TOP_TYPE_DEBUG`, `TOP_TYPE_UNKNOW` }

### Functions to call before the initialisation

- struct `starpu_top_data` \* `starpu_top_add_data_boolean` (const char \*data\_name, int active)
- struct `starpu_top_data` \* `starpu_top_add_data_integer` (const char \*data\_name, int minimum\_value, int maximum\_value, int active)
- struct `starpu_top_data` \* `starpu_top_add_data_float` (const char \*data\_name, double minimum\_value, double maximum\_value, int active)
- struct `starpu_top_param` \* `starpu_top_register_parameter_boolean` (const char \*param\_name, int \*parameter\_field, void(\*callback)(struct `starpu_top_param` \*))
- struct `starpu_top_param` \* `starpu_top_register_parameter_float` (const char \*param\_name, double \*parameter\_field, double minimum\_value, double maximum\_value, void(\*callback)(struct `starpu_top_param` \*))
- struct `starpu_top_param` \* `starpu_top_register_parameter_integer` (const char \*param\_name, int \*parameter\_field, int minimum\_value, int maximum\_value, void(\*callback)(struct `starpu_top_param` \*))
- struct `starpu_top_param` \* `starpu_top_register_parameter_enum` (const char \*param\_name, int \*parameter\_field, char \*\*values, int nb\_values, void(\*callback)(struct `starpu_top_param` \*))

### Initialisation

- void `starpu_top_init_and_wait` (const char \*server\_name)

### To call after initialisation

- void `starpu_top_update_parameter` (const struct `starpu_top_param` \*param)
- void `starpu_top_update_data_boolean` (const struct `starpu_top_data` \*data, int value)
- void `starpu_top_update_data_integer` (const struct `starpu_top_data` \*data, int value)
- void `starpu_top_update_data_float` (const struct `starpu_top_data` \*data, double value)
- void `starpu_top_task_prevision` (struct `starpu_task` \*task, int devid, unsigned long long start, unsigned long long end)
- void `starpu_top_debug_log` (const char \*message)
- void `starpu_top_debug_lock` (const char \*message)

#### 17.29.1 Detailed Description

#### 17.29.2 Data Structure Documentation

##### `struct starpu_top_data`

todo

**Data Fields**

unsigned int	<b>id</b>	todo
const char *	<b>name</b>	todo
int	<b>int_min_value</b>	todo
int	<b>int_max_value</b>	todo
double	<b>double_min_value</b>	todo
double	<b>double_max_value</b>	todo
int	<b>active</b>	todo
enum <b>starpu_top_data_type</b>	<b>type</b>	todo
struct <b>starpu_top_data</b> *	<b>next</b>	todo

**struct starpu\_top\_param**

todo

**Data Fields**

- unsigned int **id**
- const char \* **name**
- enum **starpu\_top\_param\_type** **type**
- void \* **value**
- char \*\* **enum\_values**
- int **nb\_values**
- void(\* **callback** )(struct **starpu\_top\_param** \*)
- int **int\_min\_value**
- int **int\_max\_value**
- double **double\_min\_value**
- double **double\_max\_value**
- struct **starpu\_top\_param** \* **next**

**Field Documentation****starpu\_top\_param::id** todo**starpu\_top\_param::name** todo**starpu\_top\_param::type** todo**starpu\_top\_param::value** todo**starpu\_top\_param::enum\_values** only for enum type can be NULL**starpu\_top\_param::nb\_values** todo**starpu\_top\_param::callback** todo**starpu\_top\_param::int\_min\_value** only for integer type**starpu\_top\_param::int\_max\_value** todo

`starpu_top_param::double_min_value` only for double type

`starpu_top_param::double_max_value` todo

`starpu_top_param::next` todo

### 17.29.3 Enumeration Type Documentation

`enum starpu_top_data_type`

StarPU-Top Data type

Enumerator

`STARPU_TOP_DATA_BOOLEAN` todo

`STARPU_TOP_DATA_INTEGER` todo

`STARPU_TOP_DATA_FLOAT` todo

`enum starpu_top_param_type`

StarPU-Top Parameter type

Enumerator

`STARPU_TOP_PARAM_BOOLEAN` todo

`STARPU_TOP_PARAM_INTEGER` todo

`STARPU_TOP_PARAM_FLOAT` todo

`STARPU_TOP_PARAM_ENUM` todo

`enum starpu_top_message_type`

StarPU-Top Message type

Enumerator

`TOP_TYPE_GO` todo

`TOP_TYPE_SET` todo

`TOP_TYPE_CONTINUE` todo

`TOP_TYPE_ENABLE` todo

`TOP_TYPE_DISABLE` todo

`TOP_TYPE_DEBUG` todo

`TOP_TYPE_UNKNOW` todo

### 17.29.4 Function Documentation

`struct starpu_top_data * starpu_top_add_data_boolean ( const char * data_name, int active )`

This fonction register a data named `data_name` of type boolean. If `active=0`, the value will NOT be displayed to user by default. Any other value will make the value displayed by default.

`struct starpu_top_data * starpu_top_add_data_integer ( const char * data_name, int minimum_value, int maximum_value, int active )`

This fonction register a data named `data_name` of type integer. The minimum and maximum value will be usefull to define the scale in UI. If `active=0`, the value will NOT be displayed to user by default. Any other value will make the value displayed by default.

```
struct starpu_top_data * starpu_top_add_data_float ( const char * data_name, double minimum_value,
double maximum_value, int active )
```

This fonction register a data named `data_name` of type float. The minimum and maximum value will be usefull to define the scale in UI. If `active=0`, the value will NOT be displayed to user by default. Any other value will make the value displayed by default.

```
struct starpu_top_param * starpu_top_register_parameter_boolean ( const char * param_name, int *
parameter_field, void(*)(struct starpu_top_param *) callback )
```

This fonction register a parameter named `param_name`, of type boolean. The `callback` fonction will be called when the parameter is modified by UI, and can be null.

```
struct starpu_top_param * starpu_top_register_parameter_float ( const char * param_name, double *
parameter_field, double minimum_value, double maximum_value, void(*)(struct starpu_top_param *) callback )
```

This fonction register a parameter named `param_name`, of type integer. Minimum and maximum value will be used to prevent user seting incorrect value. The `callback` fonction will be called when the parameter is modified by UI, and can be null.

```
struct starpu_top_param * starpu_top_register_parameter_integer ( const char * param_name, int *
parameter_field, int minimum_value, int maximum_value, void(*)(struct starpu_top_param *) callback )
```

This fonction register a parameter named `param_name`, of type float. Minimum and maximum value will be used to prevent user seting incorrect value. The `callback` fonction will be called when the parameter is modified by UI, and can be null.

```
struct starpu_top_param * starpu_top_register_parameter_enum ( const char * param_name, int *
parameter_field, char ** values, int nb_values, void(*)(struct starpu_top_param *) callback )
```

This fonction register a parameter named `param_name`, of type enum. Minimum and maximum value will be used to prevent user seting incorrect value. The `callback` fonction will be called when the parameter is modified by UI, and can be null.

```
void starpu_top_init_and_wait ( const char * server_name )
```

This function must be called when all parameters and data have been registered AND initialised (for parameters). This function will wait for a TOP to connect, send initialisation sentences, and wait for the GO message.

```
void starpu_top_update_parameter ( const struct starpu_top_param * param )
```

This function should be called after every modification of a parameter from something other than `starpu_top`. This fonction notice UI that the configuration changed.

```
void starpu_top_update_data_boolean ( const struct starpu_top_data * data, int value )
```

This function updates the value of the `starpu_top_data` on UI.

```
void starpu_top_update_data_integer ( const struct starpu_top_data * data, int value )
```

This function updates the value of the `starpu_top_data` on UI.

```
void starpu_top_update_data_float ( const struct starpu_top_data * data, double value )
```

This function updates the value of the `starpu_top_data` on UI.

```
void starpu_top_task_prevision ( struct starpu_task * task, int devid, unsigned long long start, unsigned
long long end )
```

This function notifies UI than the task have been planed to run from start to end, on computation-core.

```
void starpu_top_debug_log ( const char * message )
```

This function is useful in debug mode. The starpu developper doesn't need to check if the debug mode is active. This is checked by starpu\_top itself. It just send a message to display by UI.

```
void starpu_top_debug_lock ( const char * message )
```

This function is useful in debug mode. The starpu developper doesn't need to check if the debug mode is active. This is checked by starpu\_top itself. It send a message and wait for a continue message from UI to return. The lock (wich create a stop-point) should be called only by the main thread. Calling it from more than one thread is not supported.

## 17.30 Scheduling Contexts

StarPU permits on one hand grouping workers in combined workers in order to execute a parallel task and on the other hand grouping tasks in bundles that will be executed by a single specified worker. In contrast when we group workers in scheduling contexts we submit starpu tasks to them and we schedule them with the policy assigned to the context. Scheduling contexts can be created, deleted and modified dynamically.

### Data Structures

- struct `starpu_sched_ctx_performance_counters`

### Scheduling Contexts Basic API

- #define `STARPU_SCHED_CTX_POLICY_NAME`
- #define `STARPU_SCHED_CTX_POLICY_STRUCT`
- #define `STARPU_SCHED_CTX_POLICY_MIN_PRIO`
- #define `STARPU_SCHED_CTX_POLICY_MAX_PRIO`
- unsigned `starpu_sched_ctx_create` (int \*workerids\_ctx, int nworkers\_ctx, const char \*sched\_ctx\_name,...)
- unsigned `starpu_sched_ctx_create_inside_interval` (const char \*policy\_name, const char \*sched\_name, int min\_ncpus, int max\_ncpus, int min\_ngpus, int max\_ngpus, unsigned allow\_overlap)
- void `starpu_sched_ctx_register_close_callback` (unsigned sched\_ctx\_id, void(\*close\_callback)(unsigned sched\_ctx\_id, void \*args), void \*args)
- void `starpu_sched_ctx_add_workers` (int \*workerids\_ctx, int nworkers\_ctx, unsigned sched\_ctx\_id)
- void `starpu_sched_ctx_remove_workers` (int \*workerids\_ctx, int nworkers\_ctx, unsigned sched\_ctx\_id)
- void `starpu_sched_ctx_delete` (unsigned sched\_ctx\_id)
- void `starpu_sched_ctx_set_inheritor` (unsigned sched\_ctx\_id, unsigned inheritor)
- void `starpu_sched_ctx_set_context` (unsigned \*sched\_ctx\_id)
- unsigned `starpu_sched_ctx_get_context` (void)
- void `starpu_sched_ctx_stop_task_submission` (void)
- void `starpu_sched_ctx_finished_submit` (unsigned sched\_ctx\_id)
- unsigned `starpu_sched_ctx_get_workers_list` (unsigned sched\_ctx\_id, int \*\*workerids)
- unsigned `starpu_sched_ctx_get_nworkers` (unsigned sched\_ctx\_id)
- unsigned `starpu_sched_ctx_get_nshared_workers` (unsigned sched\_ctx\_id, unsigned sched\_ctx\_id2)
- unsigned `starpu_sched_ctx_contains_worker` (int workerid, unsigned sched\_ctx\_id)
- unsigned `starpu_sched_ctx_worker_get_id` (unsigned sched\_ctx\_id)
- unsigned `starpu_sched_ctx_overlapping_ctxs_on_worker` (int workerid)

### Scheduling Context Priorities

- #define `STARPU_MIN_PRIO`
- #define `STARPU_MAX_PRIO`
- #define `STARPU_DEFAULT_PRIO`
- int `starpu_sched_ctx_set_min_priority` (unsigned sched\_ctx\_id, int min\_prio)
- int `starpu_sched_ctx_set_max_priority` (unsigned sched\_ctx\_id, int max\_prio)
- int `starpu_sched_ctx_get_min_priority` (unsigned sched\_ctx\_id)
- int `starpu_sched_ctx_get_max_priority` (unsigned sched\_ctx\_id)

## Scheduling Context Worker Collection

- struct `starpu_worker_collection` \* `starpu_sched_ctx_create_worker_collection` (unsigned `sched_ctx_id`, enum `starpu_worker_collection_type` `type`)
- void `starpu_sched_ctx_delete_worker_collection` (unsigned `sched_ctx_id`)
- struct `starpu_worker_collection` \* `starpu_sched_ctx_get_worker_collection` (unsigned `sched_ctx_id`)

## Scheduling Context Link with Hypervisor

- void `starpu_sched_ctx_set_perf_counters` (unsigned `sched_ctx_id`, void \*`perf_counters`)
- void `starpu_sched_ctx_call_pushed_task_cb` (int `workerid`, unsigned `sched_ctx_id`)
- void `starpu_sched_ctx_notify_hypervisor_exists` (void)
- unsigned `starpu_sched_ctx_check_if_hypervisor_exists` (void)
- void `starpu_sched_ctx_set_policy_data` (unsigned `sched_ctx_id`, void \*`policy_data`)
- void \* `starpu_sched_ctx_get_policy_data` (unsigned `sched_ctx_id`)

### 17.30.1 Detailed Description

StarPU permits on one hand grouping workers in combined workers in order to execute a parallel task and on the other hand grouping tasks in bundles that will be executed by a single specified worker. In contrast when we group workers in scheduling contexts we submit starpu tasks to them and we schedule them with the policy assigned to the context. Scheduling contexts can be created, deleted and modified dynamically.

### 17.30.2 Data Structure Documentation

#### `struct starpu_sched_ctx_performance_counters`

Performance counters used by the starpu to indicate the hypervisor how the application and the resources are executing.

#### Data Fields

- void(\* `notify_idle_cycle` )(unsigned `sched_ctx_id`, int `worker`, double `idle_time`)
- void(\* `notify_poped_task` )(unsigned `sched_ctx_id`, int `worker`)
- void(\* `notify_pushed_task` )(unsigned `sched_ctx_id`, int `worker`)
- void(\* `notify_post_exec_task` )(struct `starpu_task` \*`task`, size\_t `data_size`, uint32\_t `footprint`, int `hypervisor_tag`, double `flops`)
- void(\* `notify_submitted_job` )(struct `starpu_task` \*`task`, uint32\_t `footprint`, size\_t `data_size`)
- void(\* `notify_empty_ctx` )(unsigned `sched_ctx_id`, struct `starpu_task` \*`task`)
- void(\* `notify_delete_context` )(unsigned `sched_ctx`)

#### Field Documentation

**`starpu_sched_ctx_performance_counters::notify_idle_cycle`** Informs the hypervisor for how long a worker has been idle in the specified context

**`starpu_sched_ctx_performance_counters::notify_poped_task`** Informs the hypervisor that a task executing a specified number of instructions has been popped from the worker

**`starpu_sched_ctx_performance_counters::notify_pushed_task`** Notifies the hypervisor that a task has been scheduled on the queue of the worker corresponding to the specified context

**`starpu_sched_ctx_performance_counters::notify_submitted_job`** Notifies the hypervisor that a task has just been submitted

**`starpu_sched_ctx_performance_counters::notify_delete_context`** Notifies the hypervisor that the context was deleted

### 17.30.3 Macro Definition Documentation

**#define STARPU\_SCHED\_CTX\_POLICY\_NAME**

This macro is used when calling [starpu\\_sched\\_ctx\\_create\(\)](#) to specify a name for a scheduling policy

**#define STARPU\_SCHED\_CTX\_POLICY\_STRUCT**

This macro is used when calling [starpu\\_sched\\_ctx\\_create\(\)](#) to specify a pointer to a scheduling policy

**#define STARPU\_SCHED\_CTX\_POLICY\_MIN\_PRIO**

This macro is used when calling [starpu\\_sched\\_ctx\\_create\(\)](#) to specify a minimum scheduler priority value.

**#define STARPU\_SCHED\_CTX\_POLICY\_MAX\_PRIO**

This macro is used when calling [starpu\\_sched\\_ctx\\_create\(\)](#) to specify a maximum scheduler priority value.

**#define STARPU\_MIN\_PRIO**

Provided for legacy reasons.

**#define STARPU\_MAX\_PRIO**

Provided for legacy reasons.

**#define STARPU\_DEFAULT\_PRIO**

By convention, the default priority level should be 0 so that we can statically allocate tasks with a default priority.

### 17.30.4 Function Documentation

**unsigned starpu\_sched\_ctx\_create ( int \* workerids\_ctx, int nworkers\_ctx, const char \* sched\_ctx\_name, ... )**

This function creates a scheduling context with the given parameters (see below) and assigns the workers in `workerids_ctx` to execute the tasks submitted to it. The return value represents the identifier of the context that has just been created. It will be further used to indicate the context the tasks will be submitted to. The return value should be at most `STARPU_NMAX_SCHED_CTXS`.

The arguments following the name of the scheduling context can be of the following types:

- `STARPU_SCHED_CTX_POLICY_NAME`, followed by the name of a predefined scheduling policy
- `STARPU_SCHED_CTX_POLICY_STRUCT`, followed by a pointer to a custom scheduling policy (struct `starpu_sched_policy *`)
- `STARPU_SCHED_CTX_POLICY_MIN_PRIO`, followed by a integer representing the minimum priority value to be defined for the scheduling policy.
- `STARPU_SCHED_CTX_POLICY_MAX_PRIO`, followed by a integer representing the maximum priority value to be defined for the scheduling policy.

**unsigned starpu\_sched\_ctx\_create\_inside\_interval ( const char \* policy\_name, const char \* sched\_ctx\_name, int min\_ncpus, int max\_ncpus, int min\_ngpus, int max\_ngpus, unsigned allow\_overlap )**

Create a context indicating an approximate interval of resources

**void starpu\_sched\_ctx\_register\_close\_callback ( unsigned sched\_ctx\_id, void(\*)(unsigned sched\_ctx\_id, void \*args) close\_callback, void \* args )**

Execute the callback whenever the last task of the context finished executing, it is called with the prameters: `sched_ctx` and any other paramter needed by the application (packed in a `void*`)

**void starpu\_sched\_ctx\_add\_workers ( int \* workerids\_ctx, int nworkers\_ctx, unsigned sched\_ctx\_id )**

This function adds dynamically the workers in `workerids_ctx` to the context `sched_ctx_id`. The last argument cannot be greater than `STARPU_NMAX_SCHED_CTXS`.

**void starpu\_sched\_ctx\_remove\_workers ( int \* workerids\_ctx, int nworkers\_ctx, unsigned sched\_ctx\_id )**

This function removes the workers in `workerids_ctx` from the context `sched_ctx_id`. The last argument cannot be greater than `STARPU_NMAX_SCHED_CTXS`.

**void starpu\_sched\_ctx\_delete ( unsigned sched\_ctx\_id )**

Delete scheduling context `sched_ctx_id` and transfer remaining workers to the inheritor scheduling context.

**void starpu\_sched\_ctx\_set\_inheritor ( unsigned sched\_ctx\_id, unsigned inheritor )**

Indicate which context will inherit the resources of this context when it will be deleted.

**void starpu\_sched\_ctx\_set\_context ( unsigned \* sched\_ctx\_id )**

Set the scheduling context the subsequent tasks will be submitted to

**unsigned starpu\_sched\_ctx\_get\_context ( void )**

Return the scheduling context the tasks are currently submitted to, or `::STARPU_NMAX_SCHED_CTXS` if no default context has been defined by calling the function `starpu_sched_ctx_set_context()`.

**void starpu\_sched\_ctx\_stop\_task\_submission ( void )**

Stop submitting tasks from the empty context list until the next time the context has time to check the empty context list

**void starpu\_sched\_ctx\_finished\_submit ( unsigned sched\_ctx\_id )**

Indicate starpu that the application finished submitting to this context in order to move the workers to the inheritor as soon as possible.

**unsigned starpu\_sched\_ctx\_get\_workers\_list ( unsigned sched\_ctx\_id, int \*\* workerids )**

Returns the list of workers in the array `workerids`, the returned value is the number of workers. The user should free the `workerids` table after finishing using it (it is allocated inside the function with the proper size)

**unsigned starpu\_sched\_ctx\_get\_nworkers ( unsigned sched\_ctx\_id )**

Return the number of workers managed by the specified contexts (Usually needed to verify if it manages any workers or if it should be blocked)

**unsigned starpu\_sched\_ctx\_get\_nshared\_workers ( unsigned sched\_ctx\_id, unsigned sched\_ctx\_id2 )**

Return the number of workers shared by two contexts.

**unsigned starpu\_sched\_ctx\_contains\_worker ( int workerid, unsigned sched\_ctx\_id )**

Return 1 if the worker belongs to the context and 0 otherwise

**unsigned starpu\_sched\_ctx\_worker\_get\_id ( unsigned sched\_ctx\_id )**

Return the workerid if the worker belongs to the context and -1 otherwise. If the thread calling this function is not a worker the function returns -1 as it calls the function `starpu_worker_get_id()`

**unsigned starpu\_sched\_ctx\_overlapping\_ctxs\_on\_worker ( int workerid )**

Check if a worker is shared between several contexts

**int starpu\_sched\_ctx\_set\_min\_priority ( unsigned sched\_ctx\_id, int min\_prio )**

Defines the minimum task priority level supported by the scheduling policy of the given scheduler context. The default minimum priority level is the same as the default priority level which is 0 by convention. The application may access that value by calling the function `starpu_sched_ctx_get_min_priority()`. This function should only be called from the initialization method of the scheduling policy, and should not be used directly from the application.

**int starpu\_sched\_ctx\_set\_max\_priority ( unsigned sched\_ctx\_id, int max\_prio )**

Defines the maximum priority level supported by the scheduling policy of the given scheduler context. The default maximum priority level is 1. The application may access that value by calling the `starpu_sched_ctx_get_max_priority` function. This function should only be called from the initialization method of the scheduling policy, and should not be used directly from the application.

**int starpu\_sched\_ctx\_get\_min\_priority ( unsigned sched\_ctx\_id )**

Returns the current minimum priority level supported by the scheduling policy of the given scheduler context.

**int starpu\_sched\_ctx\_get\_max\_priority ( unsigned sched\_ctx\_id )**

Returns the current maximum priority level supported by the scheduling policy of the given scheduler context.

**struct starpu\_worker\_collection \* starpu\_sched\_ctx\_create\_worker\_collection ( unsigned sched\_ctx\_id, enum starpu\_worker\_collection\_type type )**

Create a worker collection of the type indicated by the last parameter for the context specified through the first parameter.

**void starpu\_sched\_ctx\_delete\_worker\_collection ( unsigned sched\_ctx\_id )**

Delete the worker collection of the specified scheduling context

**struct starpu\_worker\_collection \* starpu\_sched\_ctx\_get\_worker\_collection ( unsigned sched\_ctx\_id )**

Return the worker collection managed by the indicated context

**void starpu\_sched\_ctx\_set\_perf\_counters ( unsigned sched\_ctx\_id, void \* perf\_counters )**

Indicates to starpu the pointer to the performance counter

**void starpu\_sched\_ctx\_call\_pushed\_task\_cb ( int workerid, unsigned sched\_ctx\_id )**

Callback that lets the scheduling policy tell the hypervisor that a task was pushed on a worker

**void starpu\_sched\_ctx\_notify\_hypervisor\_exists ( void )**

Allow the hypervisor to let starpu know he's initialised

**unsigned starpu\_sched\_ctx\_check\_if\_hypervisor\_exists ( void )**

Ask starpu if he is informed if the hypervisor is initialised

**void starpu\_sched\_ctx\_set\_policy\_data ( unsigned sched\_ctx\_id, void \* policy\_data )**

Allocate the scheduling policy data (private information of the scheduler like queues, variables, additional condition variables) the context

```
void * starpu_sched_ctx_get_policy_data ( unsigned sched_ctx_id )
```

Return the scheduling policy data (private information of the scheduler) of the contexts previously assigned to.

## 17.31 Scheduling Policy

TODO. While StarPU comes with a variety of scheduling policies (see [Task Scheduling Policy](#)), it may sometimes be desirable to implement custom policies to address specific problems. The API described below allows users to write their own scheduling policy.

### Data Structures

- struct [starpu\\_sched\\_policy](#)

### Functions

- struct [starpu\\_sched\\_policy](#) \*\* [starpu\\_sched\\_get\\_predefined\\_policies](#) ()
- void [starpu\\_worker\\_get\\_sched\\_condition](#) (int workerid, [starpu\(pthread\\_mutex\\_t](#) \*\*sched\_mutex, [starpu\(pthread\\_cond\\_t](#) \*\*sched\_cond)
- int [starpu\\_sched\\_set\\_min\\_priority](#) (int min\_prio)
- int [starpu\\_sched\\_set\\_max\\_priority](#) (int max\_prio)
- int [starpu\\_sched\\_get\\_min\\_priority](#) (void)
- int [starpu\\_sched\\_get\\_max\\_priority](#) (void)
- int [starpu\\_push\\_local\\_task](#) (int workerid, struct [starpu\\_task](#) \*task, int back)
- int [starpu\\_push\\_task\\_end](#) (struct [starpu\\_task](#) \*task)
- int [starpu\\_worker\\_can\\_execute\\_task](#) (unsigned workerid, struct [starpu\\_task](#) \*task, unsigned nimpl)
- double [starpu\\_timing\\_now](#) (void)
- uint32\_t [starpu\\_task\\_footprint](#) (struct [starpu\\_perfmodel](#) \*model, struct [starpu\\_task](#) \*task, enum [starpu\\_perfmodel\\_archtype](#) arch, unsigned nimpl)
- double [starpu\\_task\\_expected\\_length](#) (struct [starpu\\_task](#) \*task, enum [starpu\\_perfmodel\\_archtype](#) arch, unsigned nimpl)
- double [starpu\\_worker\\_get\\_relative\\_speedup](#) (enum [starpu\\_perfmodel\\_archtype](#) perf\_archtype)
- double [starpu\\_task\\_expected\\_data\\_transfer\\_time](#) (unsigned memory\_node, struct [starpu\\_task](#) \*task)
- double [starpu\\_data\\_expected\\_transfer\\_time](#) ([starpu\\_data\\_handle\\_t](#) handle, unsigned memory\_node, enum [starpu\\_data\\_access\\_mode](#) mode)
- double [starpu\\_task\\_expected\\_power](#) (struct [starpu\\_task](#) \*task, enum [starpu\\_perfmodel\\_archtype](#) arch, unsigned nimpl)
- double [starpu\\_task\\_expected\\_conversion\\_time](#) (struct [starpu\\_task](#) \*task, enum [starpu\\_perfmodel\\_archtype](#) arch, unsigned nimpl)
- int [starpu\\_get\\_prefetch\\_flag](#) (void)
- int [starpu\\_prefetch\\_task\\_input\\_on\\_node](#) (struct [starpu\\_task](#) \*task, unsigned node)
- void [starpu\\_sched\\_ctx\\_worker\\_shares\\_tasks\\_lists](#) (int workerid, int sched\_ctx\_id)

#### 17.31.1 Detailed Description

TODO. While StarPU comes with a variety of scheduling policies (see [Task Scheduling Policy](#)), it may sometimes be desirable to implement custom policies to address specific problems. The API described below allows users to write their own scheduling policy.

#### 17.31.2 Data Structure Documentation

##### **struct starpu\_sched\_policy**

This structure contains all the methods that implement a scheduling policy. An application may specify which scheduling strategy in the field [starpu.conf::sched\\_policy](#) passed to the function [starpu\\_init\(\)](#).

## Data Fields

- void(\* `init_sched` )(unsigned sched\_ctx\_id)
- void(\* `deinit_sched` )(unsigned sched\_ctx\_id)
- int(\* `push_task` )(struct `starpu_task` \*)
- void(\* `push_task_notify` )(struct `starpu_task` \*, int workerid, unsigned sched\_ctx\_id)
- struct `starpu_task` \*(\* `pop_task` )(unsigned sched\_ctx\_id)
- struct `starpu_task` \*(\* `pop_every_task` )(unsigned sched\_ctx\_id)
- void(\* `pre_exec_hook` )(struct `starpu_task` \*)
- void(\* `post_exec_hook` )(struct `starpu_task` \*)
- void(\* `add_workers` )(unsigned sched\_ctx\_id, int \*workerids, unsigned nworkers)
- void(\* `remove_workers` )(unsigned sched\_ctx\_id, int \*workerids, unsigned nworkers)
- const char \* `policy_name`
- const char \* `policy_description`

## Field Documentation

**`starpu_sched_policy::init_sched`** Initialize the scheduling policy.

**`starpu_sched_policy::deinit_sched`** Cleanup the scheduling policy.

**`starpu_sched_policy::push_task`** Insert a task into the scheduler.

**`starpu_sched_policy::push_task_notify`** Notify the scheduler that a task was pushed on a given worker. This method is called when a task that was explicitly assigned to a worker becomes ready and is about to be executed by the worker. This method therefore permits to keep the state of the scheduler coherent even when StarPU bypasses the scheduling strategy.

**`starpu_sched_policy::pop_task`** Get a task from the scheduler. The mutex associated to the worker is already taken when this method is called. If this method is defined as NULL, the worker will only execute tasks from its local queue. In this case, the push\_task method should use the `starpu_push_local_task` method to assign tasks to the different workers.

**`starpu_sched_policy::pop_every_task`** Remove all available tasks from the scheduler (tasks are chained by the means of the field `starpu_task::prev` and `starpu_task::next`). The mutex associated to the worker is already taken when this method is called. This is currently not used.

**`starpu_sched_policy::pre_exec_hook`** Optional field. This method is called every time a task is starting.

**`starpu_sched_policy::post_exec_hook`** Optional field. This method is called every time a task has been executed.

**`starpu_sched_policy::add_workers`** Initialize scheduling structures corresponding to each worker used by the policy.

**`starpu_sched_policy::remove_workers`** Deinitialize scheduling structures corresponding to each worker used by the policy.

**`starpu_sched_policy::policy_name`** Optional field. Name of the policy.

**`starpu_sched_policy::policy_description`** Optional field. Human readable description of the policy.

### 17.31.3 Function Documentation

**struct starpu\_sched\_policy \*\* starpu\_sched\_get\_predefined\_policies ( )**

Return an NULL-terminated array of all the predefined scheduling policies.

**void starpu\_worker\_get\_sched\_condition ( int *workerid*, starpu\_pthread\_mutex\_t \*\* *sched\_mutex*, starpu\_pthread\_cond\_t \*\* *sched\_cond* )**

When there is no available task for a worker, StarPU blocks this worker on a condition variable. This function specifies which condition variable (and the associated mutex) should be used to block (and to wake up) a worker. Note that multiple workers may use the same condition variable. For instance, in the case of a scheduling strategy with a single task queue, the same condition variable would be used to block and wake up all workers.

**int starpu\_sched\_set\_min\_priority ( int *min\_prio* )**

TODO: check if this is correct Defines the minimum task priority level supported by the scheduling policy. The default minimum priority level is the same as the default priority level which is 0 by convention. The application may access that value by calling the function [starpu\\_sched\\_get\\_min\\_priority\(\)](#). This function should only be called from the initialization method of the scheduling policy, and should not be used directly from the application.

**int starpu\_sched\_set\_max\_priority ( int *max\_prio* )**

TODO: check if this is correct Defines the maximum priority level supported by the scheduling policy. The default maximum priority level is 1. The application may access that value by calling the function [starpu\\_sched\\_get\\_max\\_priority\(\)](#). This function should only be called from the initialization method of the scheduling policy, and should not be used directly from the application.

**int starpu\_sched\_get\_min\_priority ( void )**

TODO: check if this is correct Returns the current minimum priority level supported by the scheduling policy

**int starpu\_sched\_get\_max\_priority ( void )**

TODO: check if this is correct Returns the current maximum priority level supported by the scheduling policy

**int starpu\_push\_local\_task ( int *workerid*, struct starpu\_task \* *task*, int *back* )**

The scheduling policy may put tasks directly into a worker's local queue so that it is not always necessary to create its own queue when the local queue is sufficient. If *back* is not 0, *task* is put at the back of the queue where the worker will pop tasks first. Setting *back* to 0 therefore ensures a FIFO ordering.

**int starpu\_push\_task\_end ( struct starpu\_task \* *task* )**

This function must be called by a scheduler to notify that the given task has just been pushed.

**int starpu\_worker\_can\_execute\_task ( unsigned *workerid*, struct starpu\_task \* *task*, unsigned *nimpl* )**

Check if the worker specified by *workerid* can execute the codelet. Schedulers need to call it before assigning a task to a worker, otherwise the task may fail to execute.

**double starpu\_timing\_now ( void )**

Return the current date in micro-seconds.

**uint32\_t starpu\_task\_footprint ( struct starpu\_perfmodel \* *model*, struct starpu\_task \* *task*, enum starpu\_perfmodel\_archtype *arch*, unsigned *nimpl* )**

Returns the footprint for a given task

**double starpu\_task\_expected\_length ( struct starpu\_task \* task, enum starpu\_perfmodel\_archtype arch, unsigned nimpl )**

Returns expected task duration in micro-seconds.

**double starpu\_worker\_get\_relative\_speedup ( enum starpu\_perfmodel\_archtype perf\_archtype )**

Returns an estimated speedup factor relative to CPU speed

**double starpu\_task\_expected\_data\_transfer\_time ( unsigned memory\_node, struct starpu\_task \* task )**

Returns expected data transfer time in micro-seconds.

**double starpu\_data\_expected\_transfer\_time ( starpu\_data\_handle\_t handle, unsigned memory\_node, enum starpu\_data\_access\_mode mode )**

Predict the transfer time (in micro-seconds) to move handle to a memory node

**double starpu\_task\_expected\_power ( struct starpu\_task \* task, enum starpu\_perfmodel\_archtype arch, unsigned nimpl )**

Returns expected power consumption in J

**double starpu\_task\_expected\_conversion\_time ( struct starpu\_task \* task, enum starpu\_perfmodel\_archtype arch, unsigned nimpl )**

Returns expected conversion time in ms (multiformat interface only)

**int starpu\_get\_prefetch\_flag ( void )**

Whether **STARPU\_PREFETCH** was set

**int starpu\_prefetch\_task\_input\_on\_node ( struct starpu\_task \* task, unsigned node )**

Prefetch data for a given task on a given node

**void starpu\_sched\_ctx\_worker\_shares\_tasks\_lists ( int workerid, int sched\_ctx\_id )**

The scheduling policies indicates if the worker may pop tasks from the list of other workers or if there is a central list with task for all the workers

## 17.32 Scheduling Context Hypervisor - Regular usage

### Macros

- #define SC\_HYPERVERSOR\_MAX\_IDLE
- #define SC\_HYPERVERSOR\_PRIORITY
- #define SC\_HYPERVERSOR\_MIN\_WORKERS
- #define SC\_HYPERVERSOR\_MAX\_WORKERS
- #define SC\_HYPERVERSOR\_GRANULARITY
- #define SC\_HYPERVERSOR\_FIXED\_WORKERS
- #define SC\_HYPERVERSOR\_MIN\_TASKS
- #define SC\_HYPERVERSOR\_NEW\_WORKERS\_MAX\_IDLE
- #define SC\_HYPERVERSOR\_TIME\_TO\_APPLY
- #define SC\_HYPERVERSOR\_ISPEED\_W\_SAMPLE
- #define SC\_HYPERVERSOR\_ISPEED\_CTX\_SAMPLE
- #define SC\_HYPERVERSOR\_NULL

## Functions

- void \* `sc_hypervisor_init` (struct `sc_hypervisor_policy` \*policy)
- void `sc_hypervisor_shutdown` (void)
- void `sc_hypervisor_register_ctx` (unsigned sched\_ctx, double total\_flops)
- void `sc_hypervisor_unregister_ctx` (unsigned sched\_ctx)
- void `sc_hypervisor_stop_resize` (unsigned sched\_ctx)
- void `sc_hypervisor_start_resize` (unsigned sched\_ctx)
- const char \* `sc_hypervisor_get_policy` ()
- void `sc_hypervisor_add_workers_to_sched_ctx` (int \*workers\_to\_add, unsigned nworkers\_to\_add, unsigned sched\_ctx)
- void `sc_hypervisor_remove_workers_from_sched_ctx` (int \*workers\_to\_remove, unsigned nworkers\_to\_remove, unsigned sched\_ctx, unsigned now)
- void `sc_hypervisor_move_workers` (unsigned sender\_sched\_ctx, unsigned receiver\_sched\_ctx, int \*workers\_to\_move, unsigned nworkers\_to\_move, unsigned now)
- void `sc_hypervisor_size_ctxs` (unsigned \*sched\_ctxs, int nsched\_ctxs, int \*workers, int nworkers)
- void `sc_hypervisor_set_type_of_task` (struct `starpu_codelet` \*cl, unsigned sched\_ctx, uint32\_t footprint, size\_t data\_size)
- void `sc_hypervisor_update_diff_total_flops` (unsigned sched\_ctx, double diff\_total\_flops)
- void `sc_hypervisor_update_diff_elapsed_flops` (unsigned sched\_ctx, double diff\_task\_flops)
- void `sc_hypervisor_ctl` (unsigned sched\_ctx,...)

### 17.32.1 Detailed Description

### 17.32.2 Macro Definition Documentation

#### `#define SC_HYPERVISOR_MAX_IDLE`

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 3 arguments: an array of int for the workerids to apply the condition, an int to indicate the size of the array, and a double value indicating the maximum idle time allowed for a worker before the resizing process should be triggered

#### `#define SC_HYPERVISOR_PRIORITY`

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 3 arguments: an array of int for the workerids to apply the condition, an int to indicate the size of the array, and an int value indicating the priority of the workers previously mentioned. The workers with the smallest priority are moved the first.

#### `#define SC_HYPERVISOR_MIN_WORKERS`

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 1 argument(int) indicating the minimum number of workers a context should have, underneath this limit the context cannot execute.

#### `#define SC_HYPERVISOR_MAX_WORKERS`

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 1 argument(int) indicating the maximum number of workers a context should have, above this limit the context would not be able to scale

#### `#define SC_HYPERVISOR_GRANULARITY`

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 1 argument(int) indicating the granularity of the resizing process (the number of workers should be moved from the context once it is resized) This parameter is ignore for the Gflops rate based strategy (see [Resizing Strategies](#)), the number of workers that have to be moved is calculated by the strategy.

#### `#define SC_HYPERVISOR_FIXED_WORKERS`

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 2 arguments: an array of int for the workerids to apply the condition and an int to indicate the size of the array. These workers are not allowed to be moved from the context.

**#define SC\_HYPERVISOR\_MIN\_TASKS**

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 1 argument (int) that indicated the minimum number of tasks that have to be executed before the context could be resized. This parameter is ignored for the Application Driven strategy (see [Resizing Strategies](#)) where the user indicates exactly when the resize should be done.

**#define SC\_HYPERVISOR\_NEW\_WORKERS\_MAX\_IDLE**

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 1 argument, a double value indicating the maximum idle time allowed for workers that have just been moved from other contexts in the current context.

**#define SC\_HYPERVISOR\_TIME\_TO\_APPLY**

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 1 argument (int) indicating the tag an executed task should have such that this configuration should be taken into account.

**#define SC\_HYPERVISOR\_ISPEED\_W\_SAMPLE**

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 1 argument, a double, that indicates the number of flops needed to be executed before computing the speed of a worker

**#define SC\_HYPERVISOR\_ISPEED\_CTX\_SAMPLE**

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 1 argument, a double, that indicates the number of flops needed to be executed before computing the speed of a context

**#define SC\_HYPERVISOR\_NULL**

This macro is used when calling `sc_hypervisor_ctl()` and must be followed by 1 arguments

### 17.32.3 Function Documentation

**void \* sc\_hypervisor\_init ( struct sc\_hypervisor\_policy \* policy )**

There is a single hypervisor that is in charge of resizing contexts and the resizing strategy is chosen at the initialization of the hypervisor. A single resize can be done at a time.

The Scheduling Context Hypervisor Plugin provides a series of performance counters to StarPU. By incrementing them, StarPU can help the hypervisor in the resizing decision making process.

This function initializes the hypervisor to use the strategy provided as parameter and creates the performance counters (see `starpu_sched_ctx_performance_counters`). These performance counters represent actually some callbacks that will be used by the contexts to notify the information needed by the hypervisor.

Note: The Hypervisor is actually a worker that takes this role once certain conditions trigger the resizing process (there is no additional thread assigned to the hypervisor).

**void sc\_hypervisor\_shutdown ( void )**

The hypervisor and all information concerning it is cleaned. There is no synchronization between this function and `starpu_shutdown()`. Thus, this should be called after `starpu_shutdown()`, because the performance counters will still need allocated callback functions.

**void sc\_hypervisor\_register\_ctx ( unsigned sched\_ctx, double total\_flops )**

Scheduling Contexts that have to be resized by the hypervisor must be first registered to the hypervisor. This function registers the context to the hypervisor, and indicate the number of flops the context will execute (used for Gflops rate based strategy or any other custom strategy needing it, for the others we can pass 0.0)

**void sc\_hypervisor\_unregister\_ctx ( unsigned sched\_ctx )**

Whenever we want to exclude contexts from the resizing process we have to unregister them from the hypervisor.

**void sc\_hypervisor\_stop\_resize ( unsigned sched\_ctx )**

The user can totally forbid the resizing of a certain context or can then change his mind and allow it (in this case the resizing is managed by the hypervisor, that can forbid it or allow it)

**void sc\_hypervisor\_start\_resize ( unsigned sched\_ctx )**

Allow resizing of a context. The user can then provide information to the hypervisor concerning the conditions of resizing.

**char \* sc\_hypervisor\_get\_policy ( )**

Returns the name of the resizing policy the hypervisor uses

**void sc\_hypervisor\_add\_workers\_to\_sched\_ctx ( int \* workers\_to\_add, unsigned nworkers\_to\_add, unsigned sched\_ctx )**

Ask the hypervisor to add workers to a sched\_ctx

**void sc\_hypervisor\_remove\_workers\_from\_sched\_ctx ( int \* workers\_to\_remove, unsigned nworkers\_to\_remove, unsigned sched\_ctx, unsigned now )**

Ask the hypervisor to remove workers from a sched\_ctx

**void sc\_hypervisor\_move\_workers ( unsigned sender\_sched\_ctx, unsigned receiver\_sched\_ctx, int \* workers\_to\_move, unsigned nworkers\_to\_move, unsigned now )**

Moves workers from one context to another

**void sc\_hypervisor\_size\_ctxs ( unsigned \* sched\_ctxs, int nsched\_ctxs, int \* workers, int nworkers )**

Ask the hypervisor to chose a distribution of workers in the required contexts

**void sc\_hypervisor\_set\_type\_of\_task ( struct starpu\_codelet \* cl, unsigned sched\_ctx, uint32\_t footprint, size\_t data\_size )**

Indicate the types of tasks a context will execute in order to better decide the sizing of ctxs

**void sc\_hypervisor\_update\_diff\_total\_flops ( unsigned sched\_ctx, double diff\_total\_flops )**

Change dynamically the total number of flops of a context, move the deadline of the finishing time of the context

**void sc\_hypervisor\_update\_diff\_elapsed\_flops ( unsigned sched\_ctx, double diff\_task\_flops )**

Change dynamically the number of the elapsed flops in a context, modify the past in order to better compute the speed

**void sc\_hypervisor\_ctl ( unsigned sched\_ctx, ... )**

Inputs conditions to the context sched\_ctx with the following arguments. The argument list must be zero-terminated.

## 17.33 Scheduling Context Hypervisor - Building a new resizing policy

### Data Structures

- struct [sc\\_hypervisor\\_policy](#)
- struct [sc\\_hypervisor\\_policy\\_config](#)
- struct [sc\\_hypervisor\\_wrapper](#)

- struct `sc_hypervisor_resize_ack`
- struct `sc_hypervisor_policy_task_pool`

## Functions

- void `sc_hypervisor_post_resize_request` (unsigned sched\_ctx, int task\_tag)
- unsigned `sc_hypervisor_get_size_req` (unsigned \*\*sched\_ctxs, int \*nsched\_ctxs, int \*\*workers, int \*nworkers)
- void `sc_hypervisor_save_size_req` (unsigned \*sched\_ctxs, int nsched\_ctxs, int \*workers, int nworkers)
- void `sc_hypervisor_free_size_req` (void)
- unsigned `sc_hypervisor_can_resize` (unsigned sched\_ctx)
- struct  
  `sc_hypervisor_policy_config` \* `sc_hypervisor_get_config` (unsigned sched\_ctx)
- void `sc_hypervisor_set_config` (unsigned sched\_ctx, void \*config)
- unsigned \* `sc_hypervisor_get_sched_ctxs` ()
- int `sc_hypervisor_get_nsched_ctxs` ()
- struct `sc_hypervisor_wrapper` \* `sc_hypervisor_get_wrapper` (unsigned sched\_ctx)
- double `sc_hypervisor_get_elapsed_flops_per_sched_ctx` (struct `sc_hypervisor_wrapper` \*sc\_w)

### 17.33.1 Detailed Description

### 17.33.2 Data Structure Documentation

#### `struct sc_hypervisor_policy`

This structure contains all the methods that implement a hypervisor resizing policy.

#### Data Fields

- const char \* `name`
- unsigned `custom`
- void(\* `size_ctxs`) (unsigned \*sched\_ctxs, int nsched\_ctxs, int \*workers, int nworkers)
- void(\* `resize_ctxs`) (unsigned \*sched\_ctxs, int nsched\_ctxs, int \*workers, int nworkers)
- void(\* `handle_idle_cycle`) (unsigned sched\_ctx, int worker)
- void(\* `handle_pushed_task`) (unsigned sched\_ctx, int worker)
- void(\* `handle_popped_task`) (unsigned sched\_ctx, int worker, struct `starpu_task` \*task, uint32\_t footprint)
- void(\* `handle_idle_end`) (unsigned sched\_ctx, int worker)
- void(\* `handle_post_exec_hook`) (unsigned sched\_ctx, int task\_tag)
- void(\* `handle_submitted_job`) (struct `starpu_codelet` \*cl, unsigned sched\_ctx, uint32\_t footprint, size\_t data\_size)
- void(\* `end_ctx`) (unsigned sched\_ctx)

#### Field Documentation

**`sc_hypervisor_policy::name`** Indicates the name of the policy, if there is not a custom policy, the policy corresponding to this name will be used by the hypervisor

**`sc_hypervisor_policy::custom`** Indicates whether the policy is custom or not

**`sc_hypervisor_policy::size_ctxs`** Distribute workers to contexts even at the beginning of the program

**`sc_hypervisor_policy::resize_ctxs`** Require explicit resizing

**`sc_hypervisor_policy::handle_idle_cycle`** It is called whenever the indicated worker executes another idle cycle in sched\_ctx

**sc\_hypervisor\_policy::handle\_pushed\_task** It is called whenever a task is pushed on the worker's queue corresponding to the context sched\_ctx

**sc\_hypervisor\_policy::handle\_poped\_task** It is called whenever a task is popped from the worker's queue corresponding to the context sched\_ctx

The hypervisor takes a decision when another task was popped from this worker in this ctx

**sc\_hypervisor\_policy::handle\_idle\_end** It is called whenever a task is executed on the indicated worker and context after a long period of idle time

**sc\_hypervisor\_policy::handle\_post\_exec\_hook** It is called whenever a tag task has just been executed. The table of resize requests is provided as well as the tag

**sc\_hypervisor\_policy::handle\_submitted\_job** The hypervisor takes a decision when a job was submitted in this ctx

**sc\_hypervisor\_policy::end\_ctx** The hypervisor takes a decision when a certain ctx was deleted

### struct sc\_hypervisor\_policy\_config

This structure contains all configuration information of a context. It contains configuration information for each context, which can be used to construct new resize strategies.

#### Data Fields

int	min_nworkers	Indicates the minimum number of workers needed by the context
int	max_nworkers	Indicates the maximum number of workers needed by the context
int	granularity	Indicates the workers granularity of the context
int	priority[STARPU_NMAXWORKERS]	Indicates the priority of each worker in the context
double	max_idle[STARPU_NMAXWORKERS]	Indicates the maximum idle time accepted before a resize is triggered
double	min_working[STARPU_NMAXWORKERS]	Indicates that underneath this limit the priority of the worker is reduced
int	fixed_workers[STARPU_NMAXWORKERS]	Indicates which workers can be moved and which ones are fixed
double	new_workers_-max_idle	Indicates the maximum idle time accepted before a resize is triggered for the workers that just arrived in the new context
double	ispeed_w_-sample[STARPU_NMAXWORKERS]	Indicates the sample used to compute the instant speed per worker
double	ispeed_ctx_-sample	Indicates the sample used to compute the instant speed per ctxs
double	time_sample	todo

### struct sc\_hypervisor\_wrapper

This structure is a wrapper of the contexts available in StarPU and contains all information about a context obtained by incrementing the performance counters.

## Data Fields

unsigned	sched.ctx	The context wrapped
struct sc_hypervisor_- policy_config *	config	The corresponding resize configuration
double	start_time_w[ST- ARPU_NMAXW- ORKERS]	
double	current_idle_- time[STARPU_- NMAXWORKE- RS]	The idle time counter of each worker of the context
double	idle_time[STAR- PU_NMAXWOR- KERS]	The time the workers were idle from the last resize
double	idle_start_time[STARPU_NMAXWORKERS]	The moment when the workers started being idle
double	exec_time[STAR- PU_NMAXWOR- KERS]	
double	exec_start_- time[STARPU_- NMAXWORKE- RS]	
int	worker_to_be_- removed[STAR- PU_NMAXWOR- KERS]	The list of workers that will leave this contexts (lazy resizing process)
int	pushed_tasks[STARPU_NMAXWORKERS]	The number of pushed tasks of each worker of the context
int	poped_tasks[STARPU_NMAXW- ORKERS]	The number of poped tasks of each worker of the context
double	total_flops	The total number of flops to execute by the context
double	total_elapsed_- flops[STARPU_- NMAXWORKE- RS]	The number of flops executed by each workers of the context
double	elapsed_flops[STARPU_NMAX- WORKERS]	The number of flops executed by each worker of the context from last resize
size_t	elapsed_data[STARPU_NMAX- WORKERS]	The quantity of data (in bytes) used to execute tasks on each worker in this ctx

int	elapsed_tasks[S-TARPU_NMAX-WORKERS]	The nr of tasks executed on each worker in this ctx
double	ref_speed[2]	The average speed of the workers (type of workers) when they belonged to this context 0 - cuda 1 - cpu
double	submitted_flops	The number of flops submitted to this ctx
double	remaining_flops	The number of flops that still have to be executed by the workers in the context
double	start_time	The time when he started executed
double	real_start_time	The first time a task was pushed to this context
double	hyp_react_start_time	
struct sc_hypervisor_resize_ack	resize_ack	The structure confirming the last resize finished and a new one can be done
starpu_pthread_mutex_t	mutex	The mutex needed to synchronize the acknowledgment of the workers into the receiver context
unsigned	total_flops_available	A boolean indicating if the hypervisor can use the flops corresponding to the entire execution of the context
unsigned	to_be_sized	
unsigned	compute_idle[S-TARPU_NMAX-WORKERS]	

### struct sc\_hypervisor\_resize\_ack

This structures checks if the workers moved to another context are actually taken into account in that context.

#### Data Fields

int	receiver_sched_ctx	The context receiving the new workers
int *	moved_workers	The workers moved to the receiver context
int	nmoved_workers	The number of workers moved
int *	acked_workers	If the value corresponding to a worker is 1, this one is taken into account in the new context if 0 not yet

### struct sc\_hypervisor\_policy\_task\_pool

task wrapper linked list

#### Data Fields

struct starpu_codelet *	cl	Which codelet has been executed
uint32_t	footprint	Task footprint key
unsigned	sched_ctx_id	Context the task belongs to
unsigned long	n	Number of tasks of this kind
size_t	data_size	The quantity of data(in bytes) needed by the task to execute
struct sc_hypervisor_policy_task_pool *	next	Other task kinds

### 17.33.3 Function Documentation

`void sc_hypervisor_post_resize_request ( unsigned sched_ctx, int task_tag )`

Requires resizing the context `sched_ctx` whenever a task tagged with the id `task_tag` finished executing

```
unsigned sc_hypervisor_get_size_req ( unsigned ** sched_ctxs, int * nsched_ctxs, int ** workers, int * nworkers )
```

Check if there are pending demands of resizing

```
void sc_hypervisor_save_size_req ( unsigned * sched_ctxs, int nsched_ctxs, int * workers, int nworkers )
```

Save a demand of resizing

```
void sc_hypervisor_free_size_req ( void )
```

Clear the list of pending demands of resizing

```
unsigned sc_hypervisor_can_resize ( unsigned sched_ctx )
```

Check out if a context can be resized

```
struct sc_hypervisor_policy_config * sc_hypervisor_get_config ( unsigned sched_ctx )
```

Returns the configuration structure of a context

```
void sc_hypervisor_set_config ( unsigned sched_ctx, void * config )
```

Set a certain configuration to a contexts

```
unsigned * sc_hypervisor_get_sched_ctxs ( )
```

Gets the contexts managed by the hypervisor

```
int sc_hypervisor_get_nsched_ctxs ( )
```

Gets the number of contexts managed by the hypervisor

```
struct sc_hypervisor_wrapper * sc_hypervisor_get_wrapper ( unsigned sched_ctx )
```

Returns the wrapper corresponding the context `sched_ctx`

```
double sc_hypervisor_get_elapsed_flops_per_sched_ctx ( struct sc_hypervisor_wrapper * sc_w )
```

Returns the flops of a context elapsed from the last resize

# Chapter 18

## File Index

### 18.1 File List

Here is a list of all documented files with brief descriptions:

sc_hypervisor.h . . . . .	214
sc_hypervisor_config.h . . . . .	215
sc_hypervisor_lp.h . . . . .	215
sc_hypervisor_monitoring.h . . . . .	216
sc_hypervisor_policy.h . . . . .	216
starpu.h . . . . .	189
starpu_bound.h . . . . .	190
starpu_config.h . . . . .	190
starpu_cublas.h . . . . .	191
starpu_cuda.h . . . . .	191
starpu_data.h . . . . .	192
starpu_data_filters.h . . . . .	193
starpu_data_interfaces.h . . . . .	195
starpu_DEPRECATED_api.h . . . . .	199
starpu_driver.h . . . . .	199
starpu_expert.h . . . . .	199
starpu_fxt.h . . . . .	199
starpu_hash.h . . . . .	200
starpu_mpi.h . . . . .	213
starpu_opencl.h . . . . .	200
starpu_perfmodel.h . . . . .	201
starpu_profiling.h . . . . .	202
starpu_rand.h . . . . .	203
starpu_sched_ctx.h . . . . .	203
starpu_sched_ctx_hypervisor.h . . . . .	204
starpu_scheduler.h . . . . .	204
starpu_stdlib.h . . . . .	205
starpu_task.h . . . . .	206
starpu_task_bundle.h . . . . .	207
starpu_task_list.h . . . . .	207
starpu_task_util.h . . . . .	208
starpu_thread.h . . . . .	208
starpu_thread_util.h . . . . .	209
starpu_top.h . . . . .	210
starpu_util.h . . . . .	211
starpu_worker.h . . . . .	212
starpufft.h . . . . .	??



# Chapter 19

## File Documentation

### 19.1 starpu.h File Reference

```
#include <stdlib.h>
#include <stdint.h>
#include <starpu_config.h>
#include <starpu_opencl.h>
#include <starpu_thread.h>
#include <starpu_thread_util.h>
#include <starpu_util.h>
#include <starpu_data.h>
#include <starpu_data_interfaces.h>
#include <starpu_data_filters.h>
#include <starpu_stdlb.h>
#include <starpu_perfmodel.h>
#include <starpu_worker.h>
#include <starpu_task.h>
#include <starpu_task_list.h>
#include <starpu_task_util.h>
#include <starpu_sched_ctx.h>
#include <starpu_expert.h>
#include <starpu_rand.h>
#include <starpu_cuda.h>
#include <starpu_cublas.h>
#include <starpu_bound.h>
#include <starpu_hash.h>
#include <starpu_profiling.h>
#include <starpu_top.h>
#include <starpu_fxt.h>
#include <starpu_driver.h>
#include "starpu_deprecated_api.h"
```

#### Data Structures

- struct `starpu_conf`

#### Macros

- `#define main`

#### Functions

- int `starpu_conf_init` (struct `starpu_conf` \*`conf`)

- int `starpu_init` (struct `starpu_conf` \*conf) `STARPU_WARN_UNUSED_RESULT`
- void `starpu_shutdown` (void)
- void `starpu_topology_print` (FILE \*f)
- int `starpu_asynchronous_copy_disabled` (void)
- int `starpu_asynchronous_cuda_copy_disabled` (void)
- int `starpu_asynchronous_opencl_copy_disabled` (void)
- void `starpu_display_stats` ()
- void `starpu_get_version` (int \*major, int \*minor, int \*release)

## 19.2 starpu\_bound.h File Reference

```
#include <stdio.h>
```

### Functions

- void `starpu_bound_start` (int deps, int prio)
- void `starpu_bound_stop` (void)
- void `starpu_bound_print_dot` (FILE \*output)
- void `starpu_bound_compute` (double \*res, double \*integer\_res, int integer)
- void `starpu_bound_print_lp` (FILE \*output)
- void `starpu_bound_print_mps` (FILE \*output)
- void `starpu_bound_print` (FILE \*output, int integer)

## 19.3 starpu\_config.h File Reference

```
#include <sys/types.h>
```

### Macros

- #define `STARPU_MAJOR_VERSION`
- #define `STARPU_MINOR_VERSION`
- #define `STARPU_RELEASE_VERSION`
- #define `STARPU_USE_CPU`
- #define `STARPU_USE_CUDA`
- #define `STARPU_USE_OPENCL`
- #define `STARPU_SIMGRID`
- #define `STARPU_HAVE_ICC`
- #define `STARPU_ATLAS`
- #define `STARPU_GOTO`
- #define `STARPU_MKL`
- #define `STARPU_SYSTEM_BLAS`
- #define `STARPU_BUILD_DIR`
- #define `STARPU_OPENCL_DATADIR`
- #define `STARPU_HAVE_MAGMA`
- #define `STARPU_OPENGL_RENDER`
- #define `STARPU_USE_GTK`
- #define `STARPU_HAVE_X11`
- #define `STARPU_HAVE_POSIX_MEMALIGN`
- #define `STARPU_HAVE_MEMALIGN`
- #define `STARPU_HAVE_MALLOC_H`
- #define `STARPU_HAVE_SYNC_BOOL_COMPARE_AND_SWAP`
- #define `STARPU_HAVE_SYNC_FETCH_AND_ADD`

- #define **STARPU\_HAVE\_SYNC\_FETCH\_AND\_OR**
- #define **STARPU\_HAVE\_SYNC\_LOCK\_TEST\_AND\_SET**
- #define **STARPU\_MODEL\_DEBUG**
- #define **STARPU\_NO\_ASSERT**
- #define **STARPU\_HAVE\_FFTW**
- #define **STARPU\_HAVE\_FFTWF**
- #define **STARPU\_HAVE\_FFTWL**
- #define **STARPU\_HAVE\_CURAND**
- #define **STARPU\_MAXNODES**
- #define **STARPU\_NMAXBUFS**
- #define **STARPU\_MAXCPUS**
- #define **STARPU\_MAXCUDADEVS**
- #define **STARPU\_MAXOPENCLDEVS**
- #define **STARPU\_NMAXWORKERS**
- #define **STARPU\_NMAX\_SCHED\_CTXS**
- #define **STARPU\_MAXIMPLEMENTATIONS**
- #define **STARPU\_USE\_SC\_HYPERVISOR**
- #define **STARPU\_HAVE\_GLPK\_H**
- #define **STARPU\_HAVE\_LIBNUMA**
- #define **STARPU\_HAVE\_WINDOWS**
- #define **STARPU\_HAVE\_UNSETENV**
- #define **\_\_starpu\_func\_\_**
- #define **\_\_starpu\_inline**
- #define **STARPU\_QUICK\_CHECK**
- #define **STARPU\_USE\_DRAND48**
- #define **STARPU\_USE\_ERAND48\_R**
- #define **STARPU\_HAVE\_NEARBYINTF**
- #define **STARPU\_HAVE\_RINTF**
- #define **STARPU\_USE\_TOP**
- #define **STARPU\_HAVE\_HWLOC**

#### Initialisation

- #define **STARPU\_USE\_MPI**

#### Typedefs

- typedef **ssize\_t starpu\_size\_t**

## 19.4 starpu\_cublas.h File Reference

#### Functions

- void **starpu\_cublas\_init** (void)
- void **starpu\_cublas\_shutdown** (void)

## 19.5 starpu\_cuda.h File Reference

```
#include <starpu_config.h>
#include <cuda.h>
#include <cuda_runtime.h>
#include <cuda_runtime_api.h>
#include <cublas.h>
```

## Macros

- `#define STARPU_CUBLAS_REPORT_ERROR(status)`
- `#define STARPU_CUDA_REPORT_ERROR(status)`

## Functions

- `void starpu_cublas_report_error (const char *func, const char *file, int line, cublasStatus status)`
- `void starpu_cuda_report_error (const char *func, const char *file, int line, cudaError_t status)`
- `cudaStream_t starpu_cuda_get_local_stream (void)`
- `struct cudaDeviceProp * starpu_cuda_get_device_properties (unsigned workerid)`
- `int starpu_cuda_copy_async_sync (void *src_ptr, unsigned src_node, void *dst_ptr, unsigned dst_node, size_t ssize, cudaStream_t stream, enum cudaMemcpyKind kind)`
- `void starpu_cuda_set_device (unsigned devid)`

## 19.6 starpu\_data.h File Reference

```
#include <starpu.h>
```

## Data Structures

- `struct starpu_data_descr`

## Macros

- `#define starpu_data_malloc_pinned_if_possible`
- `#define starpu_data_free_pinned_if_possible`

## Typedefs

- `typedef struct _starpu_data_state * starpu_data_handle_t`

## Enumerations

- `enum starpu_data_access_mode {  
 STARPU_NONE, STARPU_R, STARPU_W, STARPU_RW,  
 STARPU_SCRATCH, STARPU_REDUX }`
- `enum starpu_node_kind { STARPU_UNUSED, STARPU_CPU_RAM, STARPU_CUDA_RAM, STARPU_OPE-NCL_RAM }`

## Functions

- `void starpu_data_display_memory_stats ()`
- `unsigned starpu_worker_get_memory_node (unsigned workerid)`
- `unsigned starpu_memory_nodes_get_count (void)`
- `enum starpu_node_kind starpu_node_get_kind (unsigned node)`
- `void starpu_data_set_sequential_consistency_flag (starpu_data_handle_t handle, unsigned flag)`
- `unsigned starpu_data_get_default_sequential_consistency_flag (void)`
- `void starpu_data_set_default_sequential_consistency_flag (unsigned flag)`
- `unsigned starpu_data_test_if_allocated_on_node (starpu_data_handle_t handle, unsigned memory_node)`

### Basic Data Management API

*Data management is done at a high-level in StarPU: rather than accessing a mere list of contiguous buffers, the tasks may manipulate data that are described by a high-level construct which we call data interface.*

*An example of data interface is the "vector" interface which describes a contiguous data array on a specific memory node. This interface is a simple structure containing the number of elements in the array, the size of the*

*elements, and the address of the array in the appropriate address space (this address may be invalid if there is no valid copy of the array in the memory node). More informations on the data interfaces provided by StarPU are given in [Data Interfaces](#).*

*When a piece of data managed by StarPU is used by a task, the task implementation is given a pointer to an interface describing a valid copy of the data that is accessible from the current processing unit.*

*Every worker is associated to a memory node which is a logical abstraction of the address space from which the processing unit gets its data. For instance, the memory node associated to the different CPU workers represents main memory (RAM), the memory node associated to a GPU is DRAM embedded on the device. Every memory node is identified by a logical index which is accessible from the function `starpu_worker_get_memory_node()`. When registering a piece of data to StarPU, the specified memory node indicates where the piece of data initially resides (we also call this memory node the home node of a piece of data).*

- `void starpu_data_unregister (starpu_data_handle_t handle)`
- `void starpu_data_unregister_no_coherency (starpu_data_handle_t handle)`
- `void starpu_data_unregister_submit (starpu_data_handle_t handle)`
- `void starpu_data_invalidate (starpu_data_handle_t handle)`
- `void starpu_data_invalidate_submit (starpu_data_handle_t handle)`
- `void starpu_data_advise_as_important (starpu_data_handle_t handle, unsigned is_important)`
- `int starpu_data_request_allocation (starpu_data_handle_t handle, unsigned node)`
- `int starpu_data_prefetch_on_node (starpu_data_handle_t handle, unsigned node, unsigned async)`
- `void starpu_data_set_wt_mask (starpu_data_handle_t handle, uint32_t wt_mask)`
- `void starpu_data_query_status (starpu_data_handle_t handle, int memory_node, int *is_allocated, int *is_valid, int *is_requested)`
- `void starpu_data_set_reduction_methods (starpu_data_handle_t handle, struct starpu_codelet *redux_cl, struct starpu_codelet *init_cl)`

### MPI Insert Task

- `int starpu_data_set_rank (starpu_data_handle_t handle, int rank)`
- `int starpu_data_get_rank (starpu_data_handle_t handle)`
- `int starpu_data_set_tag (starpu_data_handle_t handle, int tag)`
- `int starpu_data_get_tag (starpu_data_handle_t handle)`

### Access registered data from the application

- `#define STARPU_DATA_ACQUIRE_CB(handle, mode, code)`
- `int starpu_data_acquire (starpu_data_handle_t handle, enum starpu_data_access_mode mode)`
- `int starpu_data_acquire_on_node (starpu_data_handle_t handle, unsigned node, enum starpu_data_access_mode mode)`
- `int starpu_data_acquire_cb (starpu_data_handle_t handle, enum starpu_data_access_mode mode, void(*callback)(void *), void *arg)`
- `int starpu_data_acquire_on_node_cb (starpu_data_handle_t handle, unsigned node, enum starpu_data_access_mode mode, void(*callback)(void *), void *arg)`
- `void starpu_data_release (starpu_data_handle_t handle)`
- `void starpu_data_release_on_node (starpu_data_handle_t handle, unsigned node)`

## 19.7 starpu\_data\_filters.h File Reference

```
#include <starpu.h>
#include <stdarg.h>
```

### Data Structures

- `struct starpu_data_filter`

## Functions

### Basic API

- void `starpu_data_partition` (`starpu_data_handle_t` initial\_handle, `struct starpu_data_filter *f`)
- void `starpu_data_unpartition` (`starpu_data_handle_t` root\_data, unsigned gathering\_node)
- int `starpu_data_get_nb_children` (`starpu_data_handle_t` handle)
- `starpu_data_handle_t starpu_data_get_child` (`starpu_data_handle_t` handle, unsigned i)
- `starpu_data_handle_t starpu_data_get_sub_data` (`starpu_data_handle_t` root\_data, unsigned depth,...)
- `starpu_data_handle_t starpu_data_vget_sub_data` (`starpu_data_handle_t` root\_data, unsigned depth, va\_list pa)
- void `starpu_data_map_filters` (`starpu_data_handle_t` root\_data, unsigned nfilters,...)
- void `starpu_data_vmap_filters` (`starpu_data_handle_t` root\_data, unsigned nfilters, va\_list pa)

### Predefined BCSR Filter Functions

This section gives a partial list of the predefined partitioning functions for BCSR data. Examples on how to use them are shown in [Partitioning Data](#). The complete list can be found in the file `starpu_data_filters.h`.

- void `starpu_bcsr_filter_canonical_block` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_csr_filter_vertical_block` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)

### Predefined Matrix Filter Functions

This section gives a partial list of the predefined partitioning functions for matrix data. Examples on how to use them are shown in [Partitioning Data](#). The complete list can be found in the file `starpu_data_filters.h`.

- void `starpu_matrix_filter_block` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_matrix_filter_block_shadow` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_matrix_filter_vertical_block` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_matrix_filter_vertical_block_shadow` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)

### Predefined Vector Filter Functions

This section gives a partial list of the predefined partitioning functions for vector data. Examples on how to use them are shown in [Partitioning Data](#). The complete list can be found in the file `starpu_data_filters.h`.

- void `starpu_vector_filter_block` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_vector_filter_block_shadow` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_vector_filter_list` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_vector_filter_divide_in_2` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)

### Predefined Block Filter Functions

This section gives a partial list of the predefined partitioning functions for block data. Examples on how to use them are shown in [Partitioning Data](#). The complete list can be found in the file `starpu_data_filters.h`. A usage example is available in `examples/filters/shadow3d.c`

- void `starpu_block_filter_block` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_block_filter_block_shadow` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_block_filter_vertical_block` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)
- void `starpu_block_filter_vertical_block_shadow` (void \*father\_interface, void \*child\_interface, `struct starpu_data_filter *f`, unsigned id, unsigned nparts)

- void `starpu_block_filter_depth_block` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)
- void `starpu_block_filter_depth_block_shadow` (void \*father\_interface, void \*child\_interface, struct `starpu_data_filter` \*f, unsigned id, unsigned nparts)

## 19.8 starpu\_data\_interfaces.h File Reference

```
#include <starpu.h>
#include <cuda_runtime.h>
```

### Data Structures

- struct `starpu_data_copy_methods`
- struct `starpu_data_interface_ops`
- struct `starpu_matrix_interface`
- struct `starpu_coo_interface`
- struct `starpu_block_interface`
- struct `starpu_vector_interface`
- struct `starpu_variable_interface`
- struct `starpu_csr_interface`
- struct `starpu_bcsr_interface`
- struct `starpu_multiformat_data_interface_ops`
- struct `starpu_multiformat_interface`

### Macros

- #define `STARPU_MULTIFORMAT_GET_CPU_PTR`(interface)
- #define `STARPU_MULTIFORMAT_GET_CUDA_PTR`(interface)
- #define `STARPU_MULTIFORMAT_GET_OPENCL_PTR`(interface)
- #define `STARPU_MULTIFORMAT_GET_NX`(interface)

### Accessing COO Data Interfaces

- #define `STARPU_COO_GET_COLUMNS`(interface)
- #define `STARPU_COO_GET_COLUMNS_DEV_HANDLE`(interface)
- #define `STARPU_COO_GET_ROWS`(interface)
- #define `STARPU_COO_GET_ROWS_DEV_HANDLE`(interface)
- #define `STARPU_COO_GET_VALUES`(interface)
- #define `STARPU_COO_GET_VALUES_DEV_HANDLE`(interface)
- #define `STARPU_COO_GET_OFFSET`
- #define `STARPU_COO_GET_NX`(interface)
- #define `STARPU_COO_GET_NY`(interface)
- #define `STARPU_COO_GET_NVALUES`(interface)
- #define `STARPU_COO_GET_ELEMSIZE`(interface)

### Enumerations

- enum `starpu_data_interface_id` {
 `STARPU_UNKNOWN_INTERFACE_ID`, `STARPU_MATRIX_INTERFACE_ID`, `STARPU_BLOCK_INTERFACE_ID`, `STARPU_VECTOR_INTERFACE_ID`,  
`STARPU_CSR_INTERFACE_ID`, `STARPU_BCSR_INTERFACE_ID`, `STARPU_VARIABLE_INTERFACE_ID`, `STARPU_VOID_INTERFACE_ID`,  
`STARPU_MULTIFORMAT_INTERFACE_ID`, `STARPU_COO_INTERFACE_ID`, `STARPU_MAX_INTERFACE_ID`
}

## Functions

- void `starpu_multiformat_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, void \*ptr, uint32\_t nobjects, struct `starpu_multiformat_data_interface_ops` \*format\_ops)

### Defining Interface

*Applications can provide their own interface as shown in [Defining A New Data Interface](#).*

- int `starpu_interface_copy` (uintptr\_t src, size\_t src\_offset, unsigned src\_node, uintptr\_t dst, size\_t dst\_offset, unsigned dst\_node, size\_t size, void \*async\_data)
- uintptr\_t `starpu_malloc_on_node` (unsigned dst\_node, size\_t size)
- void `starpu_free_on_node` (unsigned dst\_node, uintptr\_t addr, size\_t size)
- int `starpu_data_interface_get_next_id` (void)

### Basic Data Management API

*Data management is done at a high-level in StarPU: rather than accessing a mere list of contiguous buffers, the tasks may manipulate data that are described by a high-level construct which we call data interface.*

*An example of data interface is the "vector" interface which describes a contiguous data array on a specific memory node. This interface is a simple structure containing the number of elements in the array, the size of the elements, and the address of the array in the appropriate address space (this address may be invalid if there is no valid copy of the array in the memory node). More informations on the data interfaces provided by StarPU are given in [Data Interfaces](#).*

*When a piece of data managed by StarPU is used by a task, the task implementation is given a pointer to an interface describing a valid copy of the data that is accessible from the current processing unit.*

*Every worker is associated to a memory node which is a logical abstraction of the address space from which the processing unit gets its data. For instance, the memory node associated to the different CPU workers represents main memory (RAM), the memory node associated to a GPU is DRAM embedded on the device. Every memory node is identified by a logical index which is accessible from the function `starpu_worker_get_memory_node()`. When registering a piece of data to StarPU, the specified memory node indicates where the piece of data initially resides (we also call this memory node the home node of a piece of data).*

- void `starpu_data_register` (`starpu_data_handle_t` \*handleptr, unsigned home\_node, void \*data\_interface, struct `starpu_data_interface_ops` \*ops)
- void `starpu_data_register_same` (`starpu_data_handle_t` \*handledst, `starpu_data_handle_t` handlesrc)
- `starpu_data_handle_t starpu_data_lookup` (const void \*ptr)

### Accessing Data Interfaces

*Each data interface is provided with a set of field access functions. The ones using a void \* parameter aimed to be used in codelet implementations (see for example the code in [Vector Scaling Using StarPU's API](#)).*

- void \* `starpu_data_handle_to_pointer` (`starpu_data_handle_t` handle, unsigned node)
- void \* `starpu_data_get_local_ptr` (`starpu_data_handle_t` handle)
- enum `starpu_data_interface_id` `starpu_data_get_interface_id` (`starpu_data_handle_t` handle)
- int `starpu_data_pack` (`starpu_data_handle_t` handle, void \*\*ptr, starpu\_ssize\_t \*count)
- int `starpu_data_unpack` (`starpu_data_handle_t` handle, void \*ptr, size\_t count)
- size\_t `starpu_data_get_size` (`starpu_data_handle_t` handle)

### Registering Data

*There are several ways to register a memory region so that it can be managed by StarPU. The functions below allow the registration of vectors, 2D matrices, 3D matrices as well as BCSR and CSR sparse matrices.*

- void \* `starpu_data_get_interface_on_node` (`starpu_data_handle_t` handle, unsigned memory\_node)
- void `starpu_matrix_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, uintptr\_t ptr, uint32\_t ld, uint32\_t nx, uint32\_t ny, size\_t elemsize)
- void `starpu_coo_data_register` (`starpu_data_handle_t` \*handleptr, unsigned home\_node, uint32\_t nx, uint32\_t ny, uint32\_t n\_values, uint32\_t \*columns, uint32\_t \*rows, uintptr\_t values, size\_t elemsize)
- void `starpu_block_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, uintptr\_t ptr, uint32\_t ld\_y, uint32\_t ld\_z, uint32\_t nx, uint32\_t ny, uint32\_t nz, size\_t elemsize)
- void `starpu_vector_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, uintptr\_t ptr, uint32\_t nx, size\_t elemsize)

- void `starpu_variable_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, `uintptr_t` ptr, `size_t` size)
- void `starpu_void_data_register` (`starpu_data_handle_t` \*handle)
- void `starpu_csr_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, `uint32_t` nnz, `uint32_t` nrow, `uintptr_t` nzval, `uint32_t` \*colind, `uint32_t` \*rowptr, `uint32_t` firstentry, `size_t` elemsize)
- void `starpu_bcsr_data_register` (`starpu_data_handle_t` \*handle, unsigned home\_node, `uint32_t` nnz, `uint32_t` nrow, `uintptr_t` nzval, `uint32_t` \*colind, `uint32_t` \*rowptr, `uint32_t` firstentry, `uint32_t` r, `uint32_t` c, `size_t` elemsize)

## Variables

- struct `starpu_data_interface_ops` `starpu_interface_matrix.ops`

## Accessing Variable Data Interfaces

- #define `STARPU_VARIABLE_GET_PTR`(interface)
- #define `STARPU_VARIABLE_GET_ELEMSIZE`(interface)
- #define `STARPU_VARIABLE_GET_DEV_HANDLE`(interface)
- #define `STARPU_VARIABLE_GET_OFFSET`
- `size_t` `starpu_variable_get_elemsize` (`starpu_data_handle_t` handle)
- `uintptr_t` `starpu_variable_get_local_ptr` (`starpu_data_handle_t` handle)

## Accessing Vector Data Interfaces

- #define `STARPU_VECTOR_GET_PTR`(interface)
- #define `STARPU_VECTOR_GET_DEV_HANDLE`(interface)
- #define `STARPU_VECTOR_GET_OFFSET`(interface)
- #define `STARPU_VECTOR_GET_NX`(interface)
- #define `STARPU_VECTOR_GET_ELEMSIZE`(interface)
- `uint32_t` `starpu_vector_get_nx` (`starpu_data_handle_t` handle)
- `size_t` `starpu_vector_get_elemsize` (`starpu_data_handle_t` handle)
- `uintptr_t` `starpu_vector_get_local_ptr` (`starpu_data_handle_t` handle)

## Accessing Matrix Data Interfaces

- #define `STARPU_MATRIX_GET_PTR`(interface)
- #define `STARPU_MATRIX_GET_DEV_HANDLE`(interface)
- #define `STARPU_MATRIX_GET_OFFSET`(interface)
- #define `STARPU_MATRIX_GET_NX`(interface)
- #define `STARPU_MATRIX_GET_NY`(interface)
- #define `STARPU_MATRIX_GET_LD`(interface)
- #define `STARPU_MATRIX_GET_ELEMSIZE`(interface)
- `uint32_t` `starpu_matrix_get_nx` (`starpu_data_handle_t` handle)
- `uint32_t` `starpu_matrix_get_ny` (`starpu_data_handle_t` handle)
- `uint32_t` `starpu_matrix_get_local_ld` (`starpu_data_handle_t` handle)
- `uintptr_t` `starpu_matrix_get_local_ptr` (`starpu_data_handle_t` handle)
- `size_t` `starpu_matrix_get_elemsize` (`starpu_data_handle_t` handle)

## Accessing Block Data Interfaces

- #define `STARPU_BLOCK_GET_PTR`(interface)
- #define `STARPU_BLOCK_GET_DEV_HANDLE`(interface)
- #define `STARPU_BLOCK_GET_OFFSET`(interface)
- #define `STARPU_BLOCK_GET_NX`(interface)
- #define `STARPU_BLOCK_GET_NY`(interface)
- #define `STARPU_BLOCK_GET_NZ`(interface)

- `#define STARPU_BLOCK_GET_LDY(interface)`
- `#define STARPU_BLOCK_GET_LDZ(interface)`
- `#define STARPU_BLOCK_GET_ELEMSIZE(interface)`
- `uint32_t starpu_block_get_nx (starpu_data_handle_t handle)`
- `uint32_t starpu_block_get_ny (starpu_data_handle_t handle)`
- `uint32_t starpu_block_get_nz (starpu_data_handle_t handle)`
- `uint32_t starpu_block_get_local_ldy (starpu_data_handle_t handle)`
- `uint32_t starpu_block_get_local_ldz (starpu_data_handle_t handle)`
- `uintptr_t starpu_block_get_local_ptr (starpu_data_handle_t handle)`
- `size_t starpu_block_get_elemsize (starpu_data_handle_t handle)`

## Accessing BCSR Data Interfaces

- `#define STARPU_BCSR_GET_NNZ(interface)`
- `#define STARPU_BCSR_GET_NZVAL(interface)`
- `#define STARPU_BCSR_GET_NZVAL_DEV_HANDLE(interface)`
- `#define STARPU_BCSR_GET_COLIND(interface)`
- `#define STARPU_BCSR_GET_COLIND_DEV_HANDLE(interface)`
- `#define STARPU_BCSR_GET_ROWPTR(interface)`
- `#define STARPU_BCSR_GET_ROWPTR_DEV_HANDLE(interface)`
- `#define STARPU_BCSR_GET_OFFSET`
- `uint32_t starpu_bcsr_get_nnz (starpu_data_handle_t handle)`
- `uint32_t starpu_bcsr_get_nrow (starpu_data_handle_t handle)`
- `uint32_t starpu_bcsr_get_firstentry (starpu_data_handle_t handle)`
- `uintptr_t starpu_bcsr_get_local_nzval (starpu_data_handle_t handle)`
- `uint32_t * starpu_bcsr_get_local_colind (starpu_data_handle_t handle)`
- `uint32_t * starpu_bcsr_get_local_rowptr (starpu_data_handle_t handle)`
- `uint32_t starpu_bcsr_get_r (starpu_data_handle_t handle)`
- `uint32_t starpu_bcsr_get_c (starpu_data_handle_t handle)`
- `size_t starpu_bcsr_get_elemsize (starpu_data_handle_t handle)`

## Accessing CSR Data Interfaces

- `#define STARPU_CSR_GET_NNZ(interface)`
- `#define STARPU_CSR_GET_NROW(interface)`
- `#define STARPU_CSR_GET_NZVAL(interface)`
- `#define STARPU_CSR_GET_NZVAL_DEV_HANDLE(interface)`
- `#define STARPU_CSR_GET_COLIND(interface)`
- `#define STARPU_CSR_GET_COLIND_DEV_HANDLE(interface)`
- `#define STARPU_CSR_GET_ROWPTR(interface)`
- `#define STARPU_CSR_GET_ROWPTR_DEV_HANDLE(interface)`
- `#define STARPU_CSR_GET_OFFSET`
- `#define STARPU_CSR_GET_FIRSTENTRY(interface)`
- `#define STARPU_CSR_GET_ELEMSIZE(interface)`
- `uint32_t starpu_csr_get_nnz (starpu_data_handle_t handle)`
- `uint32_t starpu_csr_get_nrow (starpu_data_handle_t handle)`
- `uint32_t starpu_csr_get_firstentry (starpu_data_handle_t handle)`
- `uintptr_t starpu_csr_get_local_nzval (starpu_data_handle_t handle)`
- `uint32_t * starpu_csr_get_local_colind (starpu_data_handle_t handle)`
- `uint32_t * starpu_csr_get_local_rowptr (starpu_data_handle_t handle)`
- `size_t starpu_csr_get_elemsize (starpu_data_handle_t handle)`

## 19.9 starpu\_deprecated\_api.h File Reference

### 19.10 starpu\_driver.h File Reference

```
#include <starpu_config.h>
#include <starpu_opencl.h>
```

#### Data Structures

- struct `starpu_driver`
- union `starpu_driver.id`

#### Functions

- int `starpu_driver_run` (struct `starpu_driver` \*d)
- void `starpu_drivers_request_termination` (void)
- int `starpu_driver_init` (struct `starpu_driver` \*d)
- int `starpu_driver_run_once` (struct `starpu_driver` \*d)
- int `starpu_driver_deinit` (struct `starpu_driver` \*d)

## 19.11 starpu\_expert.h File Reference

```
#include <starpu.h>
```

#### Functions

- void `starpu_wake_all_blocked_workers` (void)
- int `starpu_progression_hook_register` (unsigned(\*func)(void \*arg), void \*arg)
- void `starpu_progression_hook_deregister` (int hook\_id)

## 19.12 starpu\_fxt.h File Reference

```
#include <starpu_perfmodel.h>
```

#### Data Structures

- struct `starpu_fxt_codelet_event`
- struct `starpu_fxt_options`

#### Macros

- #define `STARPU_FXT_MAX_FILES`

#### Functions

- void `starpu_fxt_options_init` (struct `starpu_fxt_options` \*options)
- void `starpu_fxt_generate_trace` (struct `starpu_fxt_options` \*options)
- void `starpu_fxt_start_profiling` (void)
- void `starpu_fxt_stop_profiling` (void)

## 19.13 starpu\_hash.h File Reference

```
#include <stdint.h>
#include <stddef.h>
```

### Functions

#### Defining Interface

*Applications can provide their own interface as shown in [Defining A New Data Interface](#).*

- uint32\_t `starpu_hash_crc32c_be_n` (const void \*input, size\_t n, uint32\_t inputcrc)
- uint32\_t `starpu_hash_crc32c_be` (uint32\_t input, uint32\_t inputcrc)
- uint32\_t `starpu_hash_crc32c_string` (const char \*str, uint32\_t inputcrc)

## 19.14 starpu\_opencl.h File Reference

```
#include <starpu_config.h>
#include <CL/cl.h>
#include <assert.h>
```

### Data Structures

- struct `starpu_opencl_program`

### Functions

#### Writing OpenCL kernels

- void `starpu_opencl_get_context` (int devid, cl\_context \*context)
- void `starpu_opencl_get_device` (int devid, cl\_device\_id \*device)
- void `starpu_opencl_get_queue` (int devid, cl\_command\_queue \*queue)
- void `starpu_opencl_get_current_context` (cl\_context \*context)
- void `starpu_opencl_get_current_queue` (cl\_command\_queue \*queue)
- int `starpu_opencl_set_kernel_args` (cl\_int \*err, cl\_kernel \*kernel,...)

#### Compiling OpenCL kernels

*Source codes for OpenCL kernels can be stored in a file or in a string. StarPU provides functions to build the program executable for each available OpenCL device as a `cl_program` object. This program executable can then be loaded within a specific queue as explained in the next section. These are only helpers, Applications can also fill a `starpu_opencl_program` array by hand for more advanced use (e.g. different programs on the different OpenCL devices, for relocation purpose for instance).*

- void `starpu_opencl_load_program_source` (const char \*source\_file\_name, char \*located\_file\_name, char \*located\_dir\_name, char \*opencl\_program\_source)
- int `starpu_opencl_compile_opencl_from_file` (const char \*source\_file\_name, const char \*build\_options)
- int `starpu_opencl_compile_opencl_from_string` (const char \*opencl\_program\_source, const char \*file\_name, const char \*build\_options)
- int `starpu_opencl_load_binary_opencl` (const char \*kernel\_id, struct `starpu_opencl_program` \*opencl\_programs)
- int `starpu_opencl_load_opencl_from_file` (const char \*source\_file\_name, struct `starpu_opencl_program` \*opencl\_programs, const char \*build\_options)
- int `starpu_opencl_load_opencl_from_string` (const char \*opencl\_program\_source, struct `starpu_opencl_program` \*opencl\_programs, const char \*build\_options)
- int `starpu_opencl_unload_opencl` (struct `starpu_opencl_program` \*opencl\_programs)

#### Loading OpenCL kernels

- int `starpu_opencl_load_kernel` (cl\_kernel \*kernel, cl\_command\_queue \*queue, struct `starpu_opencl_program` \*opencl\_programs, const char \*kernel\_name, int devid)
- int `starpu_opencl_release_kernel` (cl\_kernel kernel)

### OpenCL statistics

- int `starpu_opencl_collect_stats` (cl\_event event)

### OpenCL utilities

- #define `STARPU_OPENCL_DISPLAY_ERROR`(status)
- #define `STARPU_OPENCL_REPORT_ERROR`(status)
- #define `STARPU_OPENCL_REPORT_ERROR_WITH_MSG`(msg, status)
- const char \* `starpu_opencl_error_string` (cl\_int status)
- void `starpu_opencl_display_error` (const char \*func, const char \*file, int line, const char \*msg, cl\_int status)
- static \_\_starpu\_inline void `starpu_opencl_report_error` (const char \*func, const char \*file, int line, const char \*msg, cl\_int status)
- cl\_int `starpu_opencl_allocate_memory` (cl\_mem \*addr, size\_t size, cl\_mem\_flags flags)
- cl\_int `starpu_opencl_copy_ram_to_opencl` (void \*ptr, unsigned src\_node, cl\_mem buffer, unsigned dst\_node, size\_t size, size\_t offset, cl\_event \*event, int \*ret)
- cl\_int `starpu_opencl_copy_opencl_to_ram` (cl\_mem buffer, unsigned src\_node, void \*ptr, unsigned dst\_node, size\_t size, size\_t offset, cl\_event \*event, int \*ret)
- cl\_int `starpu_opencl_copy_opencl_to_opencl` (cl\_mem src, unsigned src\_node, size\_t src\_offset, cl\_mem dst, unsigned dst\_node, size\_t dst\_offset, size\_t size, cl\_event \*event, int \*ret)
- cl\_int `starpu_opencl_copy_async_sync` (uintptr\_t src, size\_t src\_offset, unsigned src\_node, uintptr\_t dst, size\_t dst\_offset, unsigned dst\_node, size\_t size, cl\_event \*event)

## 19.15 starpu\_perfmodel.h File Reference

```
#include <starpu.h>
#include <stdio.h>
#include <starpu_util.h>
```

### Data Structures

- struct `starpu_perfmodel_history_entry`
- struct `starpu_perfmodel_history_list`
- struct `starpu_perfmodel_regression_model`
- struct `starpu_perfmodel_per_arch`
- struct `starpu_perfmodel`

### Macros

- #define `STARPU_NARCH_VARIATIONS`
- #define `starpu_per_arch_perfmodel`

### Enumerations

- enum `starpu_perfmodel_archtype` { `STARPU_CPU_DEFAULT`, `STARPU_CUDA_DEFAULT`, `STARPU_OPENCL_DEFAULT` }
- enum `starpu_perfmodel_type` {
`STARPU_PER_ARCH`, `STARPU_COMMON`, `STARPU_HISTORY_BASED`, `STARPU_REGRESSION_BASED`,
`STARPU_NL_REGRESSION_BASED` }

## Functions

- enum `starpu_perfmodel_archtype` `starpu_worker_get_perf_archtype` (int workerid)
- int `starpu_perfmodel_load_symbol` (const char \*symbol, struct `starpu_perfmodel` \*model)
- int `starpu_perfmodel_unload_model` (struct `starpu_perfmodel` \*model)
- void `starpu_perfmodel_debugfilepath` (struct `starpu_perfmodel` \*model, enum `starpu_perfmodel_archtype` arch, char \*path, size\_t maxlen, unsigned nimpl)
- void `starpu_perfmodel_get_arch_name` (enum `starpu_perfmodel_archtype` arch, char \*archname, size\_t maxlen, unsigned nimpl)
- double `starpu_perfmodel_history_based_expected_perf` (struct `starpu_perfmodel` \*model, enum `starpu_perfmodel_archtype` arch, uint32\_t footprint)
- int `starpu_perfmodel_list` (FILE \*output)
- void `starpu_perfmodel_print` (struct `starpu_perfmodel` \*model, enum `starpu_perfmodel_archtype` arch, unsigned nimpl, char \*parameter, uint32\_t \*footprint, FILE \*output)
- int `starpu_perfmodel_print_all` (struct `starpu_perfmodel` \*model, char \*arch, char \*parameter, uint32\_t \*footprint, FILE \*output)
- void `starpu_perfmodel_directory` (FILE \*output)
- void `starpu_perfmodel_update_history` (struct `starpu_perfmodel` \*model, struct `starpu_task` \*task, enum `starpu_perfmodel_archtype` arch, unsigned cpuid, unsigned nimpl, double measured)
- void `starpu_bus_print_bandwidth` (FILE \*f)
- void `starpu_bus_print_affinity` (FILE \*f)
- double `starpu_transfer_bandwidth` (unsigned src\_node, unsigned dst\_node)
- double `starpu_transfer_latency` (unsigned src\_node, unsigned dst\_node)
- double `starpu_transfer_predict` (unsigned src\_node, unsigned dst\_node, size\_t size)

## 19.16 starpu\_profiling.h File Reference

```
#include <starpu.h>
#include <errno.h>
#include <time.h>
```

## Data Structures

- struct `starpu_profiling_task_info`
- struct `starpu_profiling_worker_info`
- struct `starpu_profiling_bus_info`

## Macros

- #define `STARPU_PROFILING_DISABLE`
- #define `STARPU_PROFILING_ENABLE`
- #define `starpu_timespec_cmp`(a, b, CMP)

## Functions

- void `starpu_profiling_init` ()
- void `starpu_profiling_set_id` (int new\_id)
- int `starpu_profiling_status_set` (int status)
- int `starpu_profiling_status_get` (void)
- int `starpu_profiling_worker_get_info` (int workerid, struct `starpu_profiling_worker_info` \*worker\_info)
- int `starpu_bus_get_count` (void)
- int `starpu_bus_get_id` (int src, int dst)
- int `starpu_bus_get_src` (int busid)
- int `starpu_bus_get_dst` (int busid)
- int `starpu_bus_get_profiling_info` (int busid, struct `starpu_profiling_bus_info` \*bus\_info)

- static \_\_starpu\_inline void **starpu\_timespec\_clear** (struct timespec \*tsp)
- static \_\_starpu\_inline void **starpu\_timespec\_add** (struct timespec \*a, struct timespec \*b, struct timespec \*result)
- static \_\_starpu\_inline void **starpu\_timespec\_accumulate** (struct timespec \*result, struct timespec \*a)
- static \_\_starpu\_inline void **starpu\_timespec\_sub** (const struct timespec \*a, const struct timespec \*b, struct timespec \*result)
- double **starpu\_timing\_timespec\_delay\_us** (struct timespec \*start, struct timespec \*end)
- double **starpu\_timing\_timespec\_to\_us** (struct timespec \*ts)
- void **starpu\_profiling\_bus\_helper\_display\_summary** (void)
- void **starpu\_profiling\_worker\_helper\_display\_summary** (void)

## 19.17 starpu\_rand.h File Reference

```
#include <stdlib.h>
#include <starpu_config.h>
```

### Macros

- #define **starpu\_srand48**(seed)
- #define **starpu\_drand48**()
- #define **starpu\_erand48**(xsubi)
- #define **starpu\_srand48\_r**(seed, buffer)
- #define **starpu\_erand48\_r**(xsubi, buffer, result)

### Typedefs

- typedef int **starpu\_drand48\_data**

## 19.18 starpu\_sched\_ctx.h File Reference

```
#include <starpu.h>
```

### Functions

- unsigned **starpu\_sched\_ctx\_contains\_type\_of\_worker** (enum **starpu\_worker\_archtype** arch, unsigned sched\_ctx\_id)
- int **starpu\_sched\_get\_min\_priority** (void)
- int **starpu\_sched\_get\_max\_priority** (void)
- int **starpu\_sched\_set\_min\_priority** (int min\_prio)
- int **starpu\_sched\_set\_max\_priority** (int max\_prio)
- int **starpu\_sched\_ctx\_min\_priority\_is\_set** (unsigned sched\_ctx\_id)
- int **starpu\_sched\_ctx\_max\_priority\_is\_set** (unsigned sched\_ctx\_id)
- int **starpu\_get\_nready\_tasks\_of\_sched\_ctx** (unsigned sched\_ctx\_id)
- double **starpu\_get\_nready\_flops\_of\_sched\_ctx** (unsigned sched\_ctx\_id)

### Scheduling Context Worker Collection

- struct **starpu\_worker\_collection** \* **starpu\_sched\_ctx\_create\_worker\_collection** (unsigned sched\_ctx\_id, enum **starpu\_worker\_collection\_type** type)
- void **starpu\_sched\_ctx\_delete\_worker\_collection** (unsigned sched\_ctx\_id)
- struct **starpu\_worker\_collection** \* **starpu\_sched\_ctx\_get\_worker\_collection** (unsigned sched\_ctx\_id)

### Scheduling Context Link with Hypervisor

- void **starpu\_sched\_ctx\_set\_policy\_data** (unsigned sched\_ctx\_id, void \*policy\_data)
- void \* **starpu\_sched\_ctx\_get\_policy\_data** (unsigned sched\_ctx\_id)
- void **starpu\_sched\_ctx\_call\_pushed\_task\_cb** (int workerid, unsigned sched\_ctx\_id)

## Scheduling Contexts Basic API

- `#define STARPU_SCHED_CTX_POLICY_NAME`
- `#define STARPU_SCHED_CTX_POLICY_STRUCT`
- `#define STARPU_SCHED_CTX_POLICY_MIN_PRIO`
- `#define STARPU_SCHED_CTX_POLICY_MAX_PRIO`
- `unsigned starpu_sched_ctx_create (int *workerids_ctx, int nworkers_ctx, const char *sched_ctx_name,...)`
- `unsigned starpu_sched_ctx_create_inside_interval (const char *policy_name, const char *sched_name, int min_ncpus, int max_ncpus, int min_ngpus, int max_ngpus, unsigned allow_overlap)`
- `void starpu_sched_ctx_register_close_callback (unsigned sched_ctx_id, void(*close_callback)(unsigned sched_ctx_id, void *args), void *args)`
- `void starpu_sched_ctx_add_workers (int *workerids_ctx, int nworkers_ctx, unsigned sched_ctx_id)`
- `void starpu_sched_ctx_remove_workers (int *workerids_ctx, int nworkers_ctx, unsigned sched_ctx_id)`
- `void starpu_sched_ctx_delete (unsigned sched_ctx_id)`
- `void starpu_sched_ctx_set_inheritor (unsigned sched_ctx_id, unsigned inheritor)`
- `void starpu_sched_ctx_set_context (unsigned *sched_ctx_id)`
- `unsigned starpu_sched_ctx_get_context (void)`
- `void starpu_sched_ctx_stop_task_submission (void)`
- `void starpu_sched_ctx_finished_submit (unsigned sched_ctx_id)`
- `unsigned starpu_sched_ctx_get_workers_list (unsigned sched_ctx_id, int **workerids)`
- `unsigned starpu_sched_ctx_get_nworkers (unsigned sched_ctx_id)`
- `unsigned starpu_sched_ctx_get_nshared_workers (unsigned sched_ctx_id, unsigned sched_ctx_id2)`
- `unsigned starpu_sched_ctx_contains_worker (int workerid, unsigned sched_ctx_id)`
- `unsigned starpu_sched_ctx_worker_get_id (unsigned sched_ctx_id)`
- `unsigned starpu_sched_ctx_overlapping_ctxs_on_worker (int workerid)`

## Scheduling Context Priorities

- `#define STARPU_MIN_PRIO`
- `#define STARPU_MAX_PRIO`
- `#define STARPU_DEFAULT_PRIO`
- `int starpu_sched_ctx_get_min_priority (unsigned sched_ctx_id)`
- `int starpu_sched_ctx_get_max_priority (unsigned sched_ctx_id)`
- `int starpu_sched_ctx_set_min_priority (unsigned sched_ctx_id, int min_prio)`
- `int starpu_sched_ctx_set_max_priority (unsigned sched_ctx_id, int max_prio)`

## 19.19 starpu\_sched\_ctx\_hypervisor.h File Reference

### Data Structures

- struct `starpu_sched_ctx_performance_counters`

### Functions

#### Scheduling Context Link with Hypervisor

- `void starpu_sched_ctx_set_perf_counters (unsigned sched_ctx_id, void *perf_counters)`
- `void starpu_sched_ctx_notify_hypervisor_exists (void)`
- `unsigned starpu_sched_ctx_check_if_hypervisor_exists (void)`

## 19.20 starpu\_scheduler.h File Reference

```
#include <starpu.h>
```

## Data Structures

- struct `starpu_sched_policy`

## Functions

- struct `starpu_sched_policy` \*\* `starpu_sched_get_predefined_policies` ()
- void `starpu_worker_get_sched_condition` (int workerid, `starpu_pthread_mutex_t` \*\*`sched_mutex`, `starpu_pthread_cond_t` \*\*`sched_cond`)
- int `starpu_wakeup_worker` (int workerid, `starpu_pthread_cond_t` \*`cond`, `starpu_pthread_mutex_t` \*`mutex`)
- int `starpu_worker_can_execute_task` (unsigned workerid, struct `starpu_task` \*`task`, unsigned nimpl)
- int `starpu_push_local_task` (int workerid, struct `starpu_task` \*`task`, int back)
- int `starpu_push_task_end` (struct `starpu_task` \*`task`)
- int `starpu_combined_worker_assign_workerid` (int nworkers, int workerid\_array[])
- int `starpu_combined_worker_get_description` (int workerid, int \*`worker_size`, int \*\*`combined_workerid`)
- int `starpu_combined_worker_can_execute_task` (unsigned workerid, struct `starpu_task` \*`task`, unsigned nimpl)
- int `starpu_get_prefetch_flag` (void)
- int `starpu_prefetch_task_input_on_node` (struct `starpu_task` \*`task`, unsigned node)
- uint32\_t `starpu_task_footprint` (struct `starpu_perfmodel` \*`model`, struct `starpu_task` \*`task`, enum `starpu_perfmodel_archtype` arch, unsigned nimpl)
- double `starpu_task_expected_length` (struct `starpu_task` \*`task`, enum `starpu_perfmodel_archtype` arch, unsigned nimpl)
- double `starpu_worker_get_relative_speedup` (enum `starpu_perfmodel_archtype` perf\_archtype)
- double `starpu_task_expected_data_transfer_time` (unsigned memory\_node, struct `starpu_task` \*`task`)
- double `starpu_data_expected_transfer_time` (`starpu_data_handle_t` handle, unsigned memory\_node, enum `starpu_data_access_mode` mode)
- double `starpu_task_expected_power` (struct `starpu_task` \*`task`, enum `starpu_perfmodel_archtype` arch, unsigned nimpl)
- double `starpu_task_expected_conversion_time` (struct `starpu_task` \*`task`, enum `starpu_perfmodel_archtype` arch, unsigned nimpl)
- double `starpu_task_bundle_expected_length` (`starpu_task_bundle_t` bundle, enum `starpu_perfmodel_archtype` arch, unsigned nimpl)
- double `starpu_task_bundle_expected_data_transfer_time` (`starpu_task_bundle_t` bundle, unsigned memory\_node)
- double `starpu_task_bundle_expected_power` (`starpu_task_bundle_t` bundle, enum `starpu_perfmodel_archtype` arch, unsigned nimpl)
- void `starpu_sched_ctx_worker_shares_tasks_lists` (int workerid, int sched\_ctx\_id)

## 19.21 starpu\_stdlib.h File Reference

```
#include <starpu.h>
```

## Macros

- #define `STARPU_MALLOC_PINNED`
- #define `STARPU_MALLOC_COUNT`

## Functions

- void `starpu_malloc_set_align` (size\_t align)
- int `starpu_malloc` (void \*\*A, size\_t dim)
- int `starpu_free` (void \*A)
- int `starpu_malloc_flags` (void \*\*A, size\_t dim, int flags)
- int `starpu_free_flags` (void \*A, size\_t dim, int flags)
- starpu\_ssize\_t `starpu_memory_get_available` (unsigned node)

## 19.22 starpu\_task.h File Reference

```
#include <starpu.h>
#include <starpu_data.h>
#include <starpu_util.h>
#include <starpu_task_bundle.h>
#include <errno.h>
#include <cuda.h>
```

### Data Structures

- struct `starpu_codelet`
- struct `starpu_task`

### Macros

- #define `STARPU_CPU`
- #define `STARPU_CUDA`
- #define `STARPU_OPENCL`
- #define `STARPU_TASK_INVALID`
- #define `STARPU_MULTIPLE_CPU_IMPLEMENTATIONS`
- #define `STARPU_MULTIPLE_CUDA_IMPLEMENTATIONS`
- #define `STARPU_MULTIPLE_OPENCL_IMPLEMENTATIONS`
- #define `STARPU_TASK_INITIALIZER`
- #define `STARPU_TASK_GET_HANDLE(task, i)`
- #define `STARPU_TASK_SET_HANDLE(task, handle, i)`
- #define `STARPU_CODELET_GET_MODE(codelet, i)`
- #define `STARPU_CODELET_SET_MODE(codelet, mode, i)`

### Typedefs

- typedef uint64\_t `starpu_tag_t`
- typedef void(\* `starpu_cpu_func_t`)(void \*\*, void \*)
- typedef void(\* `starpu_cuda_func_t`)(void \*\*, void \*)
- typedef void(\* `starpu_opencl_func_t`)(void \*\*, void \*)

### Enumerations

- enum `starpu_codelet_type` { `STARPU_SEQ`, `STARPU_SPMD`, `STARPU_FORKJOIN` }
- enum `starpu_task_status` {
 `STARPU_TASK_INVALID`, `STARPU_TASK_BLOCKED`, `STARPU_TASK_BLOCKED_ON_TAG`, `STARPU_TASK_BLOCKED_ON_TASK`,  
`STARPU_TASK_RUNNING`, `STARPU_TASK_FINISHED`, `STARPU_TASK_BLOCKED_ON_DATA` }

### Functions

- void `starpu_tag_declare_deps` (`starpu_tag_t` id, unsigned ndeps,...)
- void `starpu_tag_declare_deps_array` (`starpu_tag_t` id, unsigned ndeps, `starpu_tag_t` \*array)
- void `starpu_task_declare_deps_array` (struct `starpu_task` \*task, unsigned ndeps, struct `starpu_task` \*task\_array[])
- int `starpu_tag_wait` (`starpu_tag_t` id)
- int `starpu_tag_wait_array` (unsigned ntags, `starpu_tag_t` \*id)
- void `starpu_tag_notify_from_apps` (`starpu_tag_t` id)
- void `starpu_tag_restart` (`starpu_tag_t` id)

- void `starpu_tag_remove` (`starpu_tag_t` id)
- void `starpu_task_init` (struct `starpu_task` \*task)
- void `starpu_task_clean` (struct `starpu_task` \*task)
- struct `starpu_task` \* `starpu_task_create` (void)
- void `starpu_task_destroy` (struct `starpu_task` \*task)
- int `starpu_task_submit` (struct `starpu_task` \*task) `STARPU_WARN_UNUSED_RESULT`
- int `starpu_task_submit_to_ctx` (struct `starpu_task` \*task, unsigned sched\_ctx\_id)
- int `starpu_task_wait` (struct `starpu_task` \*task) `STARPU_WARN_UNUSED_RESULT`
- int `starpu_task_wait_for_all` (void)
- int `starpu_task_wait_for_all_in_ctx` (unsigned sched\_ctx\_id)
- int `starpu_task_wait_for_no_ready` (void)
- int `starpu_task_nready` (void)
- int `starpu_task_nsubmitted` (void)
- void `starpu_codelet_init` (struct `starpu_codelet` \*cl)
- void `starpu_codelet_display_stats` (struct `starpu_codelet` \*cl)
- struct `starpu_task` \* `starpu_task_get_current` (void)
- void `starpu_parallel_task_barrier_init` (struct `starpu_task` \*task, int workerid)
- struct `starpu_task` \* `starpu_task_dup` (struct `starpu_task` \*task)
- void `starpu_task_set_implementation` (struct `starpu_task` \*task, unsigned impl)
- unsigned `starpu_task_get_implementation` (struct `starpu_task` \*task)

## 19.23 starpu\_task\_bundle.h File Reference

### Typedefs

- typedef struct  
  `_starpu_task_bundle` \* `starpu_task_bundle_t`

### Functions

- void `starpu_task_bundle_create` (`starpu_task_bundle_t` \*bundle)
- int `starpu_task_bundle_insert` (`starpu_task_bundle_t` bundle, struct `starpu_task` \*task)
- int `starpu_task_bundle_remove` (`starpu_task_bundle_t` bundle, struct `starpu_task` \*task)
- void `starpu_task_bundle_close` (`starpu_task_bundle_t` bundle)

## 19.24 starpu\_task\_list.h File Reference

```
#include <starpu_task.h>
#include <starpu_util.h>
```

### Data Structures

- struct `starpu_task_list`

### Functions

- static STARPU\_INLINE void `starpu_task_list_init` (struct `starpu_task_list` \*list)
- static STARPU\_INLINE void `starpu_task_list_push_front` (struct `starpu_task_list` \*list, struct `starpu_task` \*task)
- static STARPU\_INLINE void `starpu_task_list_push_back` (struct `starpu_task_list` \*list, struct `starpu_task` \*task)
- static STARPU\_INLINE struct  
  `starpu_task` \* `starpu_task_list_front` (struct `starpu_task_list` \*list)
- static STARPU\_INLINE struct  
  `starpu_task` \* `starpu_task_list_back` (struct `starpu_task_list` \*list)
- static STARPU\_INLINE int `starpu_task_list_empty` (struct `starpu_task_list` \*list)

- static STARPU\_INLINE void starpu\_task\_list\_erase (struct starpu\_task\_list \*list, struct starpu\_task \*task)
- static STARPU\_INLINE struct starpu\_task \* starpu\_task\_list\_pop\_front (struct starpu\_task\_list \*list)
- static STARPU\_INLINE struct starpu\_task \* starpu\_task\_list\_pop\_back (struct starpu\_task\_list \*list)
- static STARPU\_INLINE struct starpu\_task \* starpu\_task\_list\_begin (struct starpu\_task\_list \*list)
- static STARPU\_INLINE struct starpu\_task \* starpu\_task\_list\_end (struct starpu\_task\_list \*list STARPU\_ATTRIBUTE\_UNUSED)
- static STARPU\_INLINE struct starpu\_task \* starpu\_task\_list\_next (struct starpu\_task \*task)

## 19.25 starpu\_task\_util.h File Reference

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <assert.h>
#include <starpu.h>
```

### Macros

- #define STARPU\_VALUE
- #define STARPU\_CALLBACK
- #define STARPU\_CALLBACK\_WITH\_ARG
- #define STARPU\_CALLBACK\_ARG
- #define STARPU\_PRIORITY
- #define STARPU\_DATA\_ARRAY
- #define STARPU\_TAG
- #define STARPU\_HYPERVISOR\_TAG
- #define STARPU\_FLOPS
- #define STARPU\_SCHED\_CTX
- #define STARPU\_PROLOGUE\_CALLBACK
- #define STARPU\_PROLOGUE\_CALLBACK\_ARG

### MPI Insert Task

- #define STARPU\_EXECUTE\_ON\_NODE
- #define STARPU\_EXECUTE\_ON\_DATA

### Functions

- void starpu\_create\_sync\_task (starpu\_tag\_t sync\_tag, unsigned ndeps, starpu\_tag\_t \*deps, void(\*callback)(void \*), void \*callback\_arg)
- struct starpu\_task \* starpu\_task\_build (struct starpu\_codelet \*cl,...)
- int starpu\_insert\_task (struct starpu\_codelet \*cl,...)
- void starpu\_codelet\_unpack\_args (void \*cl\_arg,...)
- void starpu\_codelet\_pack\_args (void \*\*arg\_buffer, size\_t \*arg\_buffer\_size,...)

## 19.26 starpu\_thread.h File Reference

```
#include <xbt/synchro_core.h>
#include <msg/msg.h>
```

## Macros

- #define STARPU\_PTHREAD\_MUTEX\_INITIALIZER
- #define STARPU\_PTHREAD\_COND\_INITIALIZER

## Typedefs

- typedef int **starpu\_pthread\_t**
- typedef int **starpu\_pthread\_attr\_t**
- typedef xbt\_mutex\_t **starpu\_pthread\_mutex\_t**
- typedef int **starpu\_pthread\_mutexattr\_t**
- typedef int **starpu\_pthread\_key\_t**
- typedef xbt\_cond\_t **starpu\_pthread\_cond\_t**
- typedef int **starpu\_pthread\_condattr\_t**
- typedef xbt\_mutex\_t **starpu\_pthread\_rwlock\_t**
- typedef int **starpu\_pthread\_rwlockattr\_t**

## Functions

- int **starpu\_pthread\_create\_on** (char \*name, starpu\_pthread\_t \*thread, const starpu\_pthread\_attr\_t \*attr, void \*(\*start\_routine)(void \*), void \*arg, int where)
- int **starpu\_pthread\_create** (starpu\_pthread\_t \*thread, const starpu\_pthread\_attr\_t \*attr, void \*(\*start\_routine)(void \*), void \*arg)
- int **starpu\_pthread\_join** (starpu\_pthread\_t thread, void \*\*retval)
- int **starpu\_pthread\_attr\_init** (starpu\_pthread\_attr\_t \*attr)
- int **starpu\_pthread\_attr\_destroy** (starpu\_pthread\_attr\_t \*attr)
- int **starpu\_pthread\_attr\_setdetachstate** (starpu\_pthread\_attr\_t \*attr, int detachstate)
- int **starpu\_pthread\_mutex\_init** (starpu\_pthread\_mutex\_t \*mutex, const starpu\_pthread\_mutexattr\_t \*mutexattr)
- int **starpu\_pthread\_mutex\_destroy** (starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_mutex\_lock** (starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_mutex\_unlock** (starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_mutex\_trylock** (starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_key\_create** (starpu\_pthread\_key\_t \*key, void(\*destr\_function)(void \*))
- int **starpu\_pthread\_key\_delete** (starpu\_pthread\_key\_t key)
- int **starpu\_pthread\_setspecific** (starpu\_pthread\_key\_t key, const void \*pointer)
- void \* **starpu\_pthread\_getspecific** (starpu\_pthread\_key\_t key)
- int **starpu\_pthread\_cond\_init** (starpu\_pthread\_cond\_t \*cond, starpu\_pthread\_condattr\_t \*cond\_attr)
- int **starpu\_pthread\_cond\_signal** (starpu\_pthread\_cond\_t \*cond)
- int **starpu\_pthread\_cond\_broadcast** (starpu\_pthread\_cond\_t \*cond)
- int **starpu\_pthread\_cond\_wait** (starpu\_pthread\_cond\_t \*cond, starpu\_pthread\_mutex\_t \*mutex)
- int **starpu\_pthread\_cond\_timedwait** (starpu\_pthread\_cond\_t \*cond, starpu\_pthread\_mutex\_t \*mutex, const struct timespec \*abstime)
- int **starpu\_pthread\_cond\_destroy** (starpu\_pthread\_cond\_t \*cond)
- int **starpu\_pthread\_rwlock\_init** (starpu\_pthread\_rwlock\_t \*rwlock, const starpu\_pthread\_rwlockattr\_t \*attr)
- int **starpu\_pthread\_rwlock\_destroy** (starpu\_pthread\_rwlock\_t \*rwlock)
- int **starpu\_pthread\_rwlock\_rdlock** (starpu\_pthread\_rwlock\_t \*rwlock)
- int **starpu\_pthread\_rwlock\_tryrdlock** (starpu\_pthread\_rwlock\_t \*rwlock)
- int **starpu\_pthread\_rwlock\_wrlock** (starpu\_pthread\_rwlock\_t \*rwlock)
- int **starpu\_pthread\_rwlock\_trywrlock** (starpu\_pthread\_rwlock\_t \*rwlock)
- int **starpu\_pthread\_rwlock\_unlock** (starpu\_pthread\_rwlock\_t \*rwlock)

## 19.27 starpu\_thread\_util.h File Reference

```
#include <starpu_util.h>
#include <errno.h>
```

## Macros

- #define **STARPU\_PTHREAD\_CREATE\_ON**(name, thread, attr, routine, arg, where)
- #define **STARPU\_PTHREAD\_CREATE**(thread, attr, routine, arg)
- #define **STARPU\_PTHREAD\_MUTEX\_INIT**(mutex, attr)
- #define **STARPU\_PTHREAD\_MUTEX\_DESTROY**(mutex)
- #define **STARPU\_PTHREAD\_MUTEX\_LOCK**(mutex)
- #define **STARPU\_PTHREAD\_MUTEX\_TRYLOCK**(mutex)
- #define **STARPU\_PTHREAD\_MUTEX\_UNLOCK**(mutex)
- #define **STARPU\_PTHREAD\_KEY\_CREATE**(key, destr)
- #define **STARPU\_PTHREAD\_KEY\_DELETE**(key)
- #define **STARPU\_PTHREAD\_SETSPECIFIC**(key, ptr)
- #define **STARPU\_PTHREAD\_GETSPECIFIC**(key)
- #define **STARPU\_PTHREAD\_RWLOCK\_INIT**(rwlock, attr)
- #define **STARPU\_PTHREAD\_RWLOCK\_RDLOCK**(rwlock)
- #define **STARPU\_PTHREAD\_RWLOCK\_TRYRDLOCK**(rwlock)
- #define **STARPU\_PTHREAD\_RWLOCK\_WRLock**(rwlock)
- #define **STARPU\_PTHREAD\_RWLOCK\_TRYWRLock**(rwlock)
- #define **STARPU\_PTHREAD\_RWLOCK\_UNLOCK**(rwlock)
- #define **STARPU\_PTHREAD\_RWLOCK\_DESTROY**(rwlock)
- #define **STARPU\_PTHREAD\_COND\_INIT**(cond, attr)
- #define **STARPU\_PTHREAD\_COND\_DESTROY**(cond)
- #define **STARPU\_PTHREAD\_COND\_SIGNAL**(cond)
- #define **STARPU\_PTHREAD\_COND\_BROADCAST**(cond)
- #define **STARPU\_PTHREAD\_COND\_WAIT**(cond, mutex)
- #define **STARPU\_PTHREAD\_BARRIER\_INIT**(barrier, attr, count)
- #define **STARPU\_PTHREAD\_BARRIER\_DESTROY**(barrier)
- #define **STARPU\_PTHREAD\_BARRIER\_WAIT**(barrier)

## Functions

- static STARPU\_INLINE int **\_STARPU\_PTHREAD\_MUTEX\_TRYLOCK** (starpu\_thread\_mutex\_t \*mutex, char \*file, int line)
- static STARPU\_INLINE int **\_starpu\_thread\_rwlock\_tryrdlock** (starpu\_thread\_rwlock\_t \*rwlock, char \*file, int line)
- static STARPU\_INLINE int **\_starpu\_thread\_rwlock\_trywrlck** (starpu\_thread\_rwlock\_t \*rwlock, char \*file, int line)

## 19.28 starpu\_top.h File Reference

```
#include <starpu.h>
#include <stdlib.h>
#include <time.h>
```

## Data Structures

- struct **starpu\_top\_data**
- struct **starpu\_top\_param**

## Enumerations

- enum `starpu_top_data_type` { `STARPU_TOP_DATA_BOOLEAN`, `STARPU_TOP_DATA_INTEGER`, `STARPU_TOP_DATA_FLOAT` }
- enum `starpu_top_param_type` { `STARPU_TOP_PARAM_BOOLEAN`, `STARPU_TOP_PARAM_INTEGER`, `STARPU_TOP_PARAM_FLOAT`, `STARPU_TOP_PARAM_ENUM` }
- enum `starpu_top_message_type` {  
`TOP_TYPE_GO`, `TOP_TYPE_SET`, `TOP_TYPE_CONTINUE`, `TOP_TYPE_ENABLE`,  
`TOP_TYPE_DISABLE`, `TOP_TYPE_DEBUG`, `TOP_TYPE_UNKNOW` }

## Functions

### Functions to call before the initialisation

- struct `starpu_top_data` \* `starpu_top_add_data_boolean` (const char \*data\_name, int active)
- struct `starpu_top_data` \* `starpu_top_add_data_integer` (const char \*data\_name, int minimum\_value, int maximum\_value, int active)
- struct `starpu_top_data` \* `starpu_top_add_data_float` (const char \*data\_name, double minimum\_value, double maximum\_value, int active)
- struct `starpu_top_param` \* `starpu_top_register_parameter_boolean` (const char \*param\_name, int \*parameter\_field, void(\*callback)(struct `starpu_top_param` \*))
- struct `starpu_top_param` \* `starpu_top_register_parameter_integer` (const char \*param\_name, int \*parameter\_field, int minimum\_value, int maximum\_value, void(\*callback)(struct `starpu_top_param` \*))
- struct `starpu_top_param` \* `starpu_top_register_parameter_float` (const char \*param\_name, double \*parameter\_field, double minimum\_value, double maximum\_value, void(\*callback)(struct `starpu_top_param` \*))
- struct `starpu_top_param` \* `starpu_top_register_parameter_enum` (const char \*param\_name, int \*parameter\_field, char \*\*values, int nb\_values, void(\*callback)(struct `starpu_top_param` \*))

### Initialisation

- void `starpu_top_init_and_wait` (const char \*server\_name)

### To call after initialisation

- void `starpu_top_update_parameter` (const struct `starpu_top_param` \*param)
- void `starpu_top_update_data_boolean` (const struct `starpu_top_data` \*data, int value)
- void `starpu_top_update_data_integer` (const struct `starpu_top_data` \*data, int value)
- void `starpu_top_update_data_float` (const struct `starpu_top_data` \*data, double value)
- void `starpu_top_task_prevision` (struct `starpu_task` \*task, int devid, unsigned long long start, unsigned long long end)
- void `starpu_top_debug_log` (const char \*message)
- void `starpu_top_debug_lock` (const char \*message)

## 19.29 starpu.util.h File Reference

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <assert.h>
#include <starpu_config.h>
#include <starpu_task.h>
```

## Macros

- #define `STARPU_GNUC_PREREQ`(maj, min)
- #define `STARPU_UNLIKELY`(expr)
- #define `STARPU_LIKELY`(expr)
- #define `STARPU_ATTRIBUTE_UNUSED`

- `#define STARPU_ATTRIBUTE_INTERNAL`
- `#define STARPU_ATTRIBUTE_MALLOC`
- `#define STARPU_ATTRIBUTE_WARN_UNUSED_RESULT`
- `#define STARPU_ATTRIBUTE_PURE`
- `#define STARPU_ATTRIBUTE_ALIGNED(size)`
- `#define STARPU_INLINE`
- `#define endif`
- `#define STARPU_WARN_UNUSED_RESULT`
- `#define STARPU_POISON_PTR`
- `#define STARPU_MIN(a, b)`
- `#define STARPU_MAX(a, b)`
- `#define STARPU_ASSERT(x)`
- `#define STARPU_ASSERT_MSG(x, msg,...)`
- `#define STARPU_ABORT()`
- `#define STARPU_ABORT_MSG(msg,...)`
- `#define STARPU_CHECK_RETURN_VALUE(err, message,...)`
- `#define STARPU_CHECK_RETURN_VALUE_IS(err, value, message,...)`
- `#define STARPU_ATOMIC_SOMETHING(name, expr)`
- `#define STARPU_RMB()`
- `#define STARPU_WMB()`

## Functions

- `static __starpu_inline int starpu_get_env_number (const char *str)`
- `void starpu_trace_user_event (unsigned long code)`
- `void starpu_execute_on_each_worker (void(*func)(void *), void *arg, uint32_t where)`
- `void starpu_execute_on_each_worker_ex (void(*func)(void *), void *arg, uint32_t where, const char *name)`
- `void starpu_execute_on_specific_workers (void(*func)(void *), void *arg, unsigned num_workers, unsigned *workers, const char *name)`
- `int starpu_data_cpy (starpu_data_handle_t dst_handle, starpu_data_handle_t src_handle, int asynchronous, void(*callback_func)(void *), void *callback_arg)`
- `double starpu_timing_now (void)`

## 19.30 starpu\_worker.h File Reference

```
#include <stdlib.h>
#include <starpu_config.h>
```

## Data Structures

- struct `starpu_sched_ctx_iterator`
- struct `starpu_worker_collection`

## Enumerations

- enum `starpu_worker_archtype { STARPU_ANY_WORKER, STARPU_CPU_WORKER, STARPU_CUDA_WORKER, STARPU_OPENCL_WORKER }`
- enum `starpu_worker_collection_type { STARPU_WORKER_LIST }`

## Functions

- `unsigned starpu_worker_get_count (void)`
- `unsigned starpu_combined_worker_get_count (void)`
- `unsigned starpu_worker_is_combined_worker (int id)`
- `unsigned starpu_cpu_worker_get_count (void)`
- `unsigned starpu_cuda_worker_get_count (void)`
- `unsigned starpu_opencl_worker_get_count (void)`
- `int starpu_worker_get_id (void)`
- `int starpu_combined_worker_get_id (void)`
- `int starpu_combined_worker_get_size (void)`
- `int starpu_combined_worker_get_rank (void)`
- `enum starpu_worker_archtype starpu_worker_get_type (int id)`
- `int starpu_worker_get_count_by_type (enum starpu_worker_archtype type)`
- `int starpu_worker_get_ids_by_type (enum starpu_worker_archtype type, int *workerids, int maxsize)`
- `int starpu_worker_get_by_type (enum starpu_worker_archtype type, int num)`
- `int starpu_worker_get_by_devid (enum starpu_worker_archtype type, int devid)`
- `void starpu_worker_get_name (int id, char *dst, size_t maxlen)`
- `int starpu_worker_get_devid (int id)`
- `int starpu_worker_get_nsched_ctxs (int workerid)`

## 19.31 starpu\_mpi.h File Reference

```
#include <starpu.h>
#include <mpi.h>
```

## Typedefs

- `typedef void * starpu_mpi_req`

## Functions

### Communication

- `int starpu_mpi_isend (starpu_data_handle_t data_handle, starpu_mpi_req *req, int dest, int mpi_tag, MPI_Comm comm)`
- `int starpu_mpi_irecv (starpu_data_handle_t data_handle, starpu_mpi_req *req, int source, int mpi_tag, MPI_Comm comm)`
- `int starpu_mpi_send (starpu_data_handle_t data_handle, int dest, int mpi_tag, MPI_Comm comm)`
- `int starpu_mpi_recv (starpu_data_handle_t data_handle, int source, int mpi_tag, MPI_Comm comm, MPI_Status *status)`
- `int starpu_mpi_isend_detached (starpu_data_handle_t data_handle, int dest, int mpi_tag, MPI_Comm comm, void(*callback)(void *), void *arg)`
- `int starpu_mpi_irecv_detached (starpu_data_handle_t data_handle, int source, int mpi_tag, MPI_Comm comm, void(*callback)(void *), void *arg)`
- `int starpu_mpi_wait (starpu_mpi_req *req, MPI_Status *status)`
- `int starpu_mpi_test (starpu_mpi_req *req, int *flag, MPI_Status *status)`
- `int starpu_mpi_barrier (MPI_Comm comm)`
- `int starpu_mpi_isend_detached_unlock_tag (starpu_data_handle_t data_handle, int dest, int mpi_tag, MPI_Comm comm, starpu_tag_t tag)`
- `int starpu_mpi_irecv_detached_unlock_tag (starpu_data_handle_t data_handle, int source, int mpi_tag, MPI_Comm comm, starpu_tag_t tag)`
- `int starpu_mpi_isend_array_detached_unlock_tag (unsigned array_size, starpu_data_handle_t *data_handle, int *dest, int *mpi_tag, MPI_Comm *comm, starpu_tag_t tag)`
- `int starpu_mpi_irecv_array_detached_unlock_tag (unsigned array_size, starpu_data_handle_t *data_handle, int *source, int *mpi_tag, MPI_Comm *comm, starpu_tag_t tag)`

### Initialisation

- int `starpu_mpi_init` (int \*argc, char \*\*\*argv, int initialize\_mpi)
- int `starpu_mpi_initialize` (void)
- int `starpu_mpi_initialize_extended` (int \*rank, int \*world\_size)
- int `starpu_mpi_shutdown` (void)
- void `starpu_mpi_comm_amounts_retrieve` (size\_t \*comm\_amounts)

### MPI Insert Task

- int `starpu_mpi_insert_task` (MPI\_Comm comm, struct `starpu_codelet` \*codelet,...)
- void `starpu_mpi_get_data_on_node` (MPI\_Comm comm, `starpu_data_handle_t` data\_handle, int node)
- void `starpu_mpi_get_data_on_node_detached` (MPI\_Comm comm, `starpu_data_handle_t` data\_handle, int node, void(\*callback)(void \*), void \*arg)

### Collective Operations

- void `starpu_mpi_redux_data` (MPI\_Comm comm, `starpu_data_handle_t` data\_handle)
- int `starpu_mpi_scatter_detached` (`starpu_data_handle_t` \*data\_handles, int count, int root, MPI\_Comm comm, void(\*scallback)(void \*), void \*sarg, void(\*rcallback)(void \*), void \*rarg)
- int `starpu_mpi_gather_detached` (`starpu_data_handle_t` \*data\_handles, int count, int root, MPI\_Comm comm, void(\*scallback)(void \*), void \*sarg, void(\*rcallback)(void \*), void \*rarg)

### Communication Cache

- void `starpu_mpi_cache_flush` (MPI\_Comm comm, `starpu_data_handle_t` data\_handle)
- void `starpu_mpi_cache_flush_all_data` (MPI\_Comm comm)

## 19.32 sc\_hypervisor.h File Reference

```
#include <starpu.h>
#include <starpu_sched_ctx_hypervisor.h>
#include <sc_hypervisor_config.h>
#include <sc_hypervisor_monitoring.h>
#include <math.h>
```

### Data Structures

- struct `sc_hypervisor_policy`

### Functions

- void \* `sc_hypervisor_init` (struct `sc_hypervisor_policy` \*policy)
- void `sc_hypervisor_shutdown` (void)
- void `sc_hypervisor_register_ctx` (unsigned sched\_ctx, double total\_flops)
- void `sc_hypervisor_unregister_ctx` (unsigned sched\_ctx)
- void `sc_hypervisor_post_resize_request` (unsigned sched\_ctx, int task.tag)
- void `sc_hypervisor_resize_ctxs` (unsigned \*sched\_ctxs, int nsched\_ctxs, int \*workers, int nworkers)
- void `sc_hypervisor_stop_resize` (unsigned sched\_ctx)
- void `sc_hypervisor_start_resize` (unsigned sched\_ctx)
- const char \* `sc_hypervisor_get_policy` ()
- void `sc_hypervisor_add_workers_to_sched_ctx` (int \*workers\_to\_add, unsigned nworkers\_to\_add, unsigned sched\_ctx)
- void `sc_hypervisor_remove_workers_from_sched_ctx` (int \*workers\_to\_remove, unsigned nworkers\_to\_remove, unsigned sched\_ctx, unsigned now)
- void `sc_hypervisor_move_workers` (unsigned sender\_sched\_ctx, unsigned receiver\_sched\_ctx, int \*workers\_to\_move, unsigned nworkers\_to\_move, unsigned now)
- void `sc_hypervisor_size_ctxs` (unsigned \*sched\_ctxs, int nsched\_ctxs, int \*workers, int nworkers)

- unsigned `sc_hypervisor_get_size_req` (unsigned \*\*sched\_ctxs, int \*nsched\_ctxs, int \*\*workers, int \*nworkers)
- void `sc_hypervisor_save_size_req` (unsigned \*sched\_ctxs, int nsched\_ctxs, int \*workers, int nworkers)
- void `sc_hypervisor_free_size_req` (void)
- unsigned `sc_hypervisor_can_resize` (unsigned sched\_ctx)
- void `sc_hypervisor_set_type_of_task` (struct `starpu_codelet` \*cl, unsigned sched\_ctx, uint32\_t footprint, size\_t data\_size)
- void `sc_hypervisor_update_diff_total_flops` (unsigned sched\_ctx, double diff\_total\_flops)
- void `sc_hypervisor_update_diff_elapsed_flops` (unsigned sched\_ctx, double diff\_task\_flops)
- void `sc_hypervisor_update_resize_interval` (unsigned \*sched\_ctxs, int nsched\_ctxs)

## Variables

- `starpu_pthread_mutex_t act_hypervisor_mutex`

## 19.33 sc\_hypervisor\_config.h File Reference

```
#include <sc_hypervisor.h>
```

## Data Structures

- struct `sc_hypervisor_policy_config`

## Macros

- #define `SC_HYPERVISOR_MAX_IDLE`
- #define `SC_HYPERVISOR_MIN_WORKING`
- #define `SC_HYPERVISOR_PRIORITY`
- #define `SC_HYPERVISOR_MIN_WORKERS`
- #define `SC_HYPERVISOR_MAX_WORKERS`
- #define `SC_HYPERVISOR_GRANULARITY`
- #define `SC_HYPERVISOR_FIXED_WORKERS`
- #define `SC_HYPERVISOR_MIN_TASKS`
- #define `SC_HYPERVISOR_NEW_WORKERS_MAX_IDLE`
- #define `SC_HYPERVISOR_TIME_TO_APPLY`
- #define `SC_HYPERVISOR_NULL`
- #define `SC_HYPERVISOR_ISPEED_W_SAMPLE`
- #define `SC_HYPERVISOR_ISPEED_CTX_SAMPLE`
- #define `SC_HYPERVISOR_TIME_SAMPLE`
- #define `MAX_IDLE_TIME`
- #define `MIN_WORKING_TIME`

## Functions

- void `sc_hypervisor_set_config` (unsigned sched\_ctx, void \*config)
- struct `sc_hypervisor_policy_config` \* `sc_hypervisor_get_config` (unsigned sched\_ctx)
- void `sc_hypervisor_ctl` (unsigned sched\_ctx,...)

## 19.34 sc\_hypervisor\_ip.h File Reference

```
#include <sc_hypervisor.h>
#include <starpu_config.h>
```

## Functions

- double `sc_hypervisor_lp_get_nworkers_per_ctx` (int nsched\_ctxs, int ntypes\_of\_workers, double res[nsched\_ctxs][ntypes\_of\_workers], int total\_nw[ntypes\_of\_workers], struct `types_of_workers` \*tw)
- double `sc_hypervisor_lp_get_tmax` (int nw, int \*workers)
- void `sc_hypervisor_lp_round_double_to_int` (int ns, int nw, double res[ns][nw], int res\_rounded[ns][nw])
- void `sc_hypervisor_lp_redistribute_resources_in_ctxs` (int ns, int nw, int res\_rounded[ns][nw], double res[ns][nw], unsigned \*sched\_ctxs, struct `types_of_workers` \*tw)
- void `sc_hypervisor_lp_distribute_resources_in_ctxs` (unsigned \*sched\_ctxs, int ns, int nw, int res\_rounded[ns][nw], double res[ns][nw], int \*workers, int nworkers, struct `types_of_workers` \*tw)
- void `sc_hypervisor_lp_place_resources_in_ctx` (int ns, int nw, double w\_in\_s[ns][nw], unsigned \*sched\_ctxs, int \*workers, unsigned do\_size, struct `types_of_workers` \*tw)
- void `sc_hypervisor_lp_share_remaining_resources` (int ns, unsigned \*sched\_ctxs, int nworkers, int \*workers)
- double `sc_hypervisor_lp_find_tmax` (double t1, double t2)
- unsigned `sc_hypervisor_lp_execute_dichotomy` (int ns, int nw, double w\_in\_s[ns][nw], unsigned solve\_lp\_integer, void \*specific\_data, double tmin, double tmax, double smallest\_tmax, double(\*lp\_estimated\_distrib\_func)(int ns, int nw, double draft\_w\_in\_s[ns][nw], unsigned is\_integer, double tmax, void \*specific\_data))

## 19.35 sc\_hypervisor\_monitoring.h File Reference

```
#include <sc_hypervisor.h>
```

## Data Structures

- struct `sc_hypervisor_resize_ack`
- struct `sc_hypervisor_wrapper`

## Functions

- struct `sc_hypervisor_wrapper` \* `sc_hypervisor_get_wrapper` (unsigned sched\_ctx)
- unsigned \* `sc_hypervisor_get_sched_ctxs` ()
- int `sc_hypervisor_get_nsched_ctxs` ()
- int `sc_hypervisor_get_nworkers_ctx` (unsigned sched\_ctx, enum `starpu_worker_archtype` arch)
- double `sc_hypervisor_get_elapsed_flops_per_sched_ctx` (struct `sc_hypervisor_wrapper` \*sc\_w)
- double `sc_hypervisor_get_total_elapsed_flops_per_sched_ctx` (struct `sc_hypervisor_wrapper` \*sc\_w)
- double `sc_hypervisor_get_speed_per_worker_type` (struct `sc_hypervisor_wrapper` \*sc\_w, enum `starpu_worker_archtype` arch)
- double `sc_hypervisor_get_speed` (struct `sc_hypervisor_wrapper` \*sc\_w, enum `starpu_worker_archtype` arch)

## 19.36 sc\_hypervisor\_policy.h File Reference

```
#include <sc_hypervisor.h>
```

## Data Structures

- struct `types_of_workers`
- struct `sc_hypervisor_policy_task_pool`

## Macros

- #define HYPERVISOR\_REDIM\_SAMPLE
- #define HYPERVISOR\_START\_REDIM\_SAMPLE
- #define SC\_NOTHING
- #define SC\_IDLE
- #define SC\_SPEED

## Functions

- void **sc\_hypervisor\_policy.add\_task\_to\_pool** (struct `starpu_codelet` \*cl, unsigned sched\_ctx, uint32\_t footprint, struct `sc_hypervisor_policy_task_pool` \*\*task\_pools, size\_t data\_size)
- void **sc\_hypervisor\_policy.remove\_task\_from\_pool** (struct `starpu_task` \*task, uint32\_t footprint, struct `sc_hypervisor_policy_task_pool` \*\*task\_pools)
- struct `sc_hypervisor_policy_task_pool` \* **sc\_hypervisor\_policy\_clone\_task\_pool** (struct `sc_hypervisor_policy_task_pool` \*tp)
- void **sc\_hypervisor\_get\_tasks\_times** (int nw, int nt, double times[nw][nt], int \*workers, unsigned size\_ctxs, struct `sc_hypervisor_policy_task_pool` \*task\_pools)
- unsigned **sc\_hypervisor\_find\_lowest\_prio\_sched\_ctx** (unsigned req\_sched\_ctx, int nworkers\_to\_move)
- int \* **sc\_hypervisor\_get\_idlest\_workers** (unsigned sched\_ctx, int \*nworkers, enum `starpu_worker_archtype` arch)
- int \* **sc\_hypervisor\_get\_idlest\_workers\_in\_list** (int \*start, int \*workers, int nall\_workers, int \*nworkers, enum `starpu_worker_archtype` arch)
- int **sc\_hypervisor\_get\_movable\_nworkers** (struct `sc_hypervisor_policy_config` \*config, unsigned sched\_ctx, enum `starpu_worker_archtype` arch)
- int **sc\_hypervisor\_compute\_nworkers\_to\_move** (unsigned req\_sched\_ctx)
- unsigned **sc\_hypervisor\_policy\_resize** (unsigned sender\_sched\_ctx, unsigned receiver\_sched\_ctx, unsigned force\_resize, unsigned now)
- unsigned **sc\_hypervisor\_policy\_resize\_to\_unknown\_receiver** (unsigned sender\_sched\_ctx, unsigned now)
- double **sc\_hypervisor\_get\_ctx\_speed** (struct `sc_hypervisor_wrapper` \*sc\_w)
- double **sc\_hypervisor\_get\_slowest\_ctx\_exec\_time** (void)
- double **sc\_hypervisor\_get\_fastest\_ctx\_exec\_time** (void)
- double **sc\_hypervisor\_get\_speed\_per\_worker** (struct `sc_hypervisor_wrapper` \*sc\_w, unsigned worker)
- double **sc\_hypervisor\_get\_speed\_per\_worker\_type** (struct `sc_hypervisor_wrapper` \*sc\_w, enum `starpu_worker_archtype` arch)
- double **sc\_hypervisor\_get\_ref\_speed\_per\_worker\_type** (struct `sc_hypervisor_wrapper` \*sc\_w, enum `starpu_worker_archtype` arch)
- void **sc\_hypervisor\_group\_workers\_by\_type** (struct `types_of_workers` \*tw, int \*total\_nw)
- enum `starpu_worker_archtype` **sc\_hypervisor\_get\_arch\_for\_index** (unsigned w, struct `types_of_workers` \*tw)
- unsigned **sc\_hypervisor\_get\_index\_for\_arch** (enum `starpu_worker_archtype` arch, struct `types_of_workers` \*tw)
- unsigned **sc\_hypervisor\_criteria\_fulfilled** (unsigned sched\_ctx, int worker)
- unsigned **sc\_hypervisor\_check\_idle** (unsigned sched\_ctx, int worker)
- unsigned **sc\_hypervisor\_check\_speed\_gap\_btwn\_ctxs** (void)
- unsigned **sc\_hypervisor\_get\_resize\_criteria** ()
- struct `types_of_workers` \* **sc\_hypervisor\_get\_types\_of\_workers** (int \*workers, unsigned nworkers)

### 19.36.1 Data Structure Documentation

`struct types_of_workers`

**Data Fields**

unsigned	ncpus	
unsigned	ncuda	
unsigned	nw	

# Chapter 20

## Deprecated List

### Global `starpu_codelet::cpu_func`

Optional field which has been made deprecated. One should use instead the field `starpu_codelet::cpu_funcs`.

### Global `starpu_codelet::cuda_func`

Optional field which has been made deprecated. One should use instead the `starpu_codelet::cuda_funcs` field.

### Global `starpu_codelet::opencl_func`

Optional field which has been made deprecated. One should use instead the `starpu_codelet::opencl_funcs` field.

### Global `starpu_data.free_pinned_if_possible`

Equivalent to `starpu_free()`. This macro is provided to avoid breaking old codes.

### Global `starpu_data_malloc_pinned_if_possible`

Equivalent to `starpu_malloc()`. This macro is provided to avoid breaking old codes.

### Global `starpu_mpi_initialize (void)`

This function has been made deprecated. One should use instead the function `starpu_mpi_init()`. This function does not call MPI\_Init(), it should be called beforehand.

### Global `starpu_mpi_initialize_extended (int *rank, int *world_size)`

This function has been made deprecated. One should use instead the function `starpu_mpi_init()`. MPI will be initialized by starpumpi by calling `MPI_Init_Thread(argc, argv, MPI_THREAD_SERIALIZED, ...)`.

### Global `STARPU_MULTIPLE_CPU_IMPLEMENTATIONS`

Setting the field `starpu_codelet::cpu_func` with this macro indicates the codelet will have several implementations. The use of this macro is deprecated. One should always only define the field `starpu_codelet::cpu_funcs`.

### Global `STARPU_MULTIPLE_CUDA_IMPLEMENTATIONS`

Setting the field `starpu_codelet::cuda_func` with this macro indicates the codelet will have several implementations. The use of this macro is deprecated. One should always only define the field `starpu_codelet::cuda_funcs`.

### Global `STARPU_MULTIPLE_OPENCL_IMPLEMENTATIONS`

Setting the field `starpu_codelet::opencl_func` with this macro indicates the codelet will have several implementations. The use of this macro is deprecated. One should always only define the field `starpu_codelet::opencl_funcs`.

### Global `starpu_perfmodel::cost_model (struct starpu_data_descr *)`

This field is deprecated. Use instead the field `starpu_perfmodel::cost_function` field.

### Global `starpu_perfmodel_per_arch::cost_model (struct starpu_data_descr *t)`

This field is deprecated. Use instead the field `starpu_perfmodel_per_arch::cost_function`.

### Global `starpu_task::buffers [STARPU_NMAXBUFS]`

This field has been made deprecated. One should use instead the field `starpu_task::handles` to specify the data handles accessed by the task. The access modes are now defined in the field `starpu_codelet::modes`.

# Index

## .Configure Options

--disable-asynchronous-copy, 70  
--disable-asynchronous-cuda-copy, 70  
--disable-asynchronous-opencl-copy, 70  
--disable-build-doc, 69  
--disable-cpu, 69  
--disable-cuda, 70  
--disable-cuda-memcpy-peer, 70  
--disable-gcc-extensions, 70  
--disable-opencl, 70  
--disable-socl, 70  
--disable-starpu-top, 70  
--enable-allocation-cache, 71  
--enable-coverage, 69  
--enable-debug, 69  
--enable-fast, 69  
--enable-long-check, 69  
--enable-max-sched-ctxs, 70  
--enable-maxbuffers, 71  
--enable-maxcpus, 69  
--enable-maxcudadev, 69  
--enable-maximplementations, 70  
--enable-maxopencldev, 70  
--enable-model-debug, 71  
--enable-mpi-progression-hook, 70  
--enable-opencl-simulator, 70  
--enable-opengl-render, 71  
--enable-perf-debug, 71  
--enable-quick-check, 69  
--enable-spinlock-check, 69  
--enable-stats, 71  
--enable-verbose, 69  
--with-cuda-dir, 70  
--with-cuda-include-dir, 70  
--with-cuda-lib-dir, 70  
--with-hwloc, 69  
--with-mpicc, 70  
--with-opencl-dir, 70  
--with-opencl-include-dir, 70  
--with-opencl-lib-dir, 70  
--without-hwloc, 69

## .Environment Variables

OCL\_ICD\_VENDORS, 65  
SC\_HYPERVISOR.LAZY.RESIZE, 67  
SC\_HYPERVISOR.MAX\_SPEED\_GAP, 67  
SC\_HYPERVISOR.POLICY, 66  
SC\_HYPERVISOR.SAMPLE\_CRITERIA, 67  
SC\_HYPERVISOR.START\_RESIZE, 66  
SC\_HYPERVISOR.STOP\_PRINT, 67

SC\_HYPERVISOR\_TRIGGER\_RESIZE, 66  
SOCL\_OCL\_LIB\_OPENCL, 65  
STARPU\_BUS\_CALIBRATE, 65  
STARPU\_BUS\_STATS, 66  
STARPU\_CALIBRATE, 64  
STARPU\_COMM\_STATS, 65  
STARPU\_DISABLE\_ASYNCNROUNOUS\_COPY, 64  
STARPU\_DISABLE\_ASYNCNROUNOUS\_CUDA\_-  
COPY, 64  
STARPU\_DISABLE\_ASYNCNROUNOUS\_OPENCL\_-  
COPY, 64  
STARPU\_ENABLE\_CUDA\_GPU\_GPU\_DIRECT, 64  
STARPU\_FXT\_PREFIX, 65  
STARPU\_GENERATE\_TRACE, 66  
STARPU\_HOME, 65  
STARPU\_HOSTNAME, 65  
STARPU\_IDLE\_POWER, 65  
STARPU\_LIMIT\_CPU\_MEM, 66  
STARPU\_LIMIT\_CUDA\_devid\_MEM, 66  
STARPU\_LIMIT\_CUDA\_MEM, 66  
STARPU\_LIMIT\_OPENCL\_devid\_MEM, 66  
STARPU\_LIMIT\_OPENCL\_MEM, 66  
STARPU\_LOGFILENAME, 65  
STARPU\_MAX\_WORKERSIZE, 64  
STARPU\_MEMORY\_STATS, 66  
STARPU\_MIN\_WORKERSIZE, 64  
STARPU\_MPI\_CACHE, 65  
STARPU\_NCPU, 63  
STARPU\_NCPUS, 63  
STARPU\_NCUDA, 63  
STARPU\_NOOPENCL, 63  
STARPU\_OPENCL\_ON\_CPUS, 63  
STARPU\_OPENCL\_ONLY\_ON\_CPUS, 63  
STARPU\_OPENCL\_PROGRAM\_DIR, 65  
STARPU\_PREFETCH, 65  
STARPU\_PROFILING, 65  
STARPU\_SCHED, 64  
STARPU\_SCHED\_ALPHA, 65  
STARPU\_SCHED\_BETA, 65  
STARPU\_SCHED\_GAMMA, 65  
STARPU\_SILENT, 65  
STARPU\_SINGLE\_COMBINED\_WORKER, 64  
STARPU\_STATS, 66  
STARPU\_SYNTHESIZE\_ARITY\_COMBINED\_-  
WORKER, 64  
STARPU\_WATCHDOG\_CRASH, 66  
STARPU\_WATCHDOG\_TIMEOUT, 66  
STARPU\_WORKER\_STATS, 66  
STARPU\_WORKERS\_CPUID, 63

STARPU\_WORKERS\_CUDAID, 64  
STARPU\_WORKERS\_NOBIND, 63  
STARPU\_WORKERS\_OPENCLID, 64

add  
    starpu\_worker\_collection, 89

add\_workers  
    starpu\_sched\_policy, 176

allocate\_data\_on\_node  
    starpu\_data\_interface\_ops, 99

any\_to\_any  
    starpu\_data\_copy\_methods, 101

benchmarking  
    starpu\_perfmodel, 137

buffers  
    starpu\_task, 125

bundle  
    starpu\_task, 127

CUDA Extensions, 145  
    STARPU\_MAXCUDADEVS, 145  
    STARPU\_USE\_CUDA, 145  
    starpu\_cublas\_init, 146  
    starpu\_cublas\_report\_error, 146  
    starpu\_cublas\_shutdown, 146  
    starpu\_cuda\_copy\_async\_sync, 146  
    starpu\_cuda\_get\_device\_properties, 146  
    starpu\_cuda\_get\_local\_stream, 146  
    starpu\_cuda\_report\_error, 146  
    starpu\_cuda\_set\_device, 146

callback  
    starpu\_top\_param, 167

callback\_arg  
    starpu\_task, 126

callback\_arg\_free  
    starpu\_task, 126

callback\_func  
    starpu\_task, 126

can\_execute  
    starpu\_codelet, 122

cl  
    starpu\_task, 125

cl\_arg  
    starpu\_task, 125

cl\_arg\_free  
    starpu\_task, 125

cl\_arg\_size  
    starpu\_task, 125

Codelet And Tasks, 121  
    STARPU\_FORKJOIN, 129  
    STARPU\_SEQ, 129  
    STARPU\_SPMD, 129  
    STARPU\_TASK\_BLOCKED, 129  
    STARPU\_TASK\_BLOCKED\_ON\_DATA, 129  
    STARPU\_TASK\_BLOCKED\_ON\_TAG, 129  
    STARPU\_TASK\_BLOCKED\_ON\_TASK, 129  
    STARPU\_TASK\_FINISHED, 129  
    STARPU\_TASK\_READY, 129

STARPU\_TASK\_RUNNING, 129  
STARPU\_CPU, 128  
STARPU\_CUDA, 128  
STARPU\_NMAXBUFS, 128  
STARPU\_OPENCL, 128  
STARPU\_TASK\_INVALID, 129  
starpu\_codelet\_display\_stats, 131  
starpu\_codelet\_init, 130  
starpu\_codelet\_type, 129  
starpu\_cpu\_func\_t, 129  
starpu\_create\_sync\_task, 131  
starpu\_cuda\_func\_t, 129  
starpu\_opencl\_func\_t, 129  
starpu\_task\_clean, 130  
starpu\_task\_create, 130  
starpu\_task\_destroy, 130  
starpu\_task\_dup, 130  
starpu\_task\_get\_current, 131  
starpu\_task\_get\_implementation, 131  
starpu\_task\_init, 130  
starpu\_task\_nready, 131  
starpu\_task\_nsubmitted, 131  
starpu\_task\_set\_implementation, 131  
starpu\_task\_status, 129  
starpu\_task\_submit, 130  
starpu\_task\_submit\_to\_ctx, 130  
starpu\_task\_wait, 130  
starpu\_task\_wait\_for\_all, 131  
starpu\_task\_wait\_for\_all\_in\_ctx, 131  
starpu\_task\_wait\_for\_no\_ready, 131

compare  
    starpu\_data\_interface\_ops, 99

copy\_methods  
    starpu\_data\_interface\_ops, 99

cost\_function  
    starpu\_perfmodel, 137  
    starpu\_perfmodel\_per\_arch, 138

cost\_model  
    starpu\_perfmodel, 137  
    starpu\_perfmodel\_per\_arch, 138

cpu\_func  
    starpu\_codelet, 123

cpu\_funcs  
    starpu\_codelet, 123

cuda\_func  
    starpu\_codelet, 123

cuda\_funcs  
    starpu\_codelet, 123

cuda\_to\_cuda  
    starpu\_data\_copy\_methods, 100

cuda\_to\_cuda\_async  
    starpu\_data\_copy\_methods, 101

cuda\_to\_opencl  
    starpu\_data\_copy\_methods, 101

cuda\_to\_ram  
    starpu\_data\_copy\_methods, 100

cuda\_to\_ram\_async  
    starpu\_data\_copy\_methods, 101

custom  
 sc\_hypervisor\_policy, 182

Data Interfaces, 95  
 STARPU\_BCSR\_INTERFACE\_ID, 108  
 STARPU\_BLOCK\_INTERFACE\_ID, 108  
 STARPU\_COO\_INTERFACE\_ID, 108  
 STARPU\_CSR\_INTERFACE\_ID, 108  
 STARPU\_MATRIX\_INTERFACE\_ID, 108  
 STARPU\_MAX\_INTERFACE\_ID, 108  
 STARPU\_MULTIFORMAT\_INTERFACE\_ID, 108  
 STARPU\_UNKNOWN\_INTERFACE\_ID, 108  
 STARPU\_VARIABLE\_INTERFACE\_ID, 108  
 STARPU\_VECTOR\_INTERFACE\_ID, 108  
 STARPU\_VOID\_INTERFACE\_ID, 108  
 STARPU\_BCSR\_GET\_NNZ, 106  
 STARPU\_BLOCK\_GET\_NX, 105  
 STARPU\_BLOCK\_GET\_NY, 106  
 STARPU\_BLOCK\_GET\_NZ, 106  
 STARPU\_COO\_GET\_NX, 108  
 STARPU\_COO\_GET\_NY, 108  
 STARPU\_COO\_GET\_ROWS, 107  
 STARPU\_CSR\_GET\_NNZ, 106  
 STARPU\_CSR\_GET\_NROW, 107  
 starpu\_bcsr\_data\_register, 109  
 starpu\_bcsr\_get\_c, 112  
 starpu\_bcsr\_get\_elemsize, 112  
 starpu\_bcsr\_get\_firstentry, 112  
 starpu\_bcsr\_get\_local\_colind, 112  
 starpu\_bcsr\_get\_local\_nzval, 112  
 starpu\_bcsr\_get\_local\_rowptr, 112  
 starpu\_bcsr\_get\_nnz, 112  
 starpu\_bcsr\_get\_nrow, 112  
 starpu\_bcsr\_get\_r, 112  
 starpu\_block\_data\_register, 109  
 starpu\_block\_get\_elemsize, 112  
 starpu\_block\_get\_local\_ldy, 111  
 starpu\_block\_get\_local\_ldz, 111  
 starpu\_block\_get\_local\_ptr, 112  
 starpu\_block\_get\_nx, 111  
 starpu\_block\_get\_ny, 111  
 starpu\_block\_get\_nz, 111  
 starpu\_coo\_data\_register, 110  
 starpu\_csr\_data\_register, 110  
 starpu\_csr\_get\_elemsize, 113  
 starpu\_csr\_get\_firstentry, 112  
 starpu\_csr\_get\_local\_colind, 113  
 starpu\_csr\_get\_local\_nzval, 112  
 starpu\_csr\_get\_local\_rowptr, 113  
 starpu\_csr\_get\_nnz, 112  
 starpu\_csr\_get\_nrow, 112  
 starpu\_data\_get\_interface\_id, 110  
 starpu\_data\_get\_interface\_on\_node, 110  
 starpu\_data\_get\_local\_ptr, 110  
 starpu\_data\_get\_size, 110  
 starpu\_data\_handle\_to\_pointer, 110  
 starpu\_data\_interface\_get\_next\_id, 113  
 starpu\_data\_interface\_id, 108  
 starpu\_data\_pack, 110

starpu\_data\_unpack, 110  
 starpu\_free\_on\_node, 113  
 starpu\_hash\_crc32c\_be, 113  
 starpu\_hash\_crc32c\_be\_n, 113  
 starpu\_hash\_crc32c\_string, 113  
 starpu\_interface\_copy, 113  
 starpu\_malloc\_on\_node, 113  
 starpu\_matrix\_data\_register, 109  
 starpu\_matrix\_get\_elemsize, 111  
 starpu\_matrix\_get\_local\_id, 111  
 starpu\_matrix\_get\_local\_ptr, 111  
 starpu\_matrix\_get\_nx, 111  
 starpu\_matrix\_get\_ny, 111  
 starpu\_variable\_data\_register, 109  
 starpu\_variable\_get\_elemsize, 110  
 starpu\_variable\_get\_local\_ptr, 111  
 starpu\_vector\_data\_register, 109  
 starpu\_vector\_get\_elemsize, 111  
 starpu\_vector\_get\_local\_ptr, 111  
 starpu\_vector\_get\_nx, 111  
 starpu\_void\_data\_register, 109

Data Management, 91  
 STARPU\_NONE, 93  
 STARPU\_R, 93  
 STARPU\_REDUX, 93  
 STARPU\_RW, 93  
 STARPU\_SCRATCH, 93  
 STARPU\_W, 93  
 starpu\_data\_access\_mode, 93  
 starpu\_data\_acquire, 95  
 starpu\_data\_acquire\_cb, 95  
 starpu\_data\_acquire\_on\_node, 95  
 starpu\_data\_acquire\_on\_node\_cb, 95  
 starpu\_data\_advise\_as\_important, 94  
 starpu\_data\_handle\_t, 93  
 starpu\_data\_invalidate, 94  
 starpu\_data\_invalidate\_submit, 94  
 starpu\_data\_lookup, 94  
 starpu\_data\_prefetch\_on\_node, 94  
 starpu\_data\_query\_status, 94  
 starpu\_data\_register, 93  
 starpu\_data\_register\_same, 93  
 starpu\_data\_release, 95  
 starpu\_data\_release\_on\_node, 95  
 starpu\_data\_request\_allocation, 94  
 starpu\_data\_set\_reduction\_methods, 94  
 starpu\_data\_set\_wt\_mask, 94  
 starpu\_data\_unregister, 93  
 starpu\_data\_unregister\_no\_coherency, 94  
 starpu\_data\_unregister\_submit, 94

Data Partition, 113  
 starpu\_bcsr\_filter\_canonical\_block, 118  
 starpu\_block\_filter\_block, 117  
 starpu\_block\_filter\_block\_shadow, 117  
 starpu\_block\_filter\_depth\_block, 118  
 starpu\_block\_filter\_depth\_block\_shadow, 118  
 starpu\_block\_filter\_vertical\_block, 117  
 starpu\_block\_filter\_vertical\_block\_shadow, 118

starpu\_csr\_filter\_vertical\_block, 118  
starpu\_data\_get\_child, 116  
starpu\_data\_get\_nb\_children, 116  
starpu\_data\_get\_sub\_data, 116  
starpu\_data\_map\_filters, 116  
starpu\_data\_partition, 115  
starpu\_data\_unpartition, 115  
starpu\_data\_vget\_sub\_data, 116  
starpu\_data\_vmap\_filters, 116  
starpu\_matrix\_filter\_block, 117  
starpu\_matrix\_filter\_block\_shadow, 117  
starpu\_matrix\_filter\_vertical\_block, 117  
starpu\_matrix\_filter\_vertical\_block\_shadow, 117  
starpu\_vector\_filter\_block, 116  
starpu\_vector\_filter\_block\_shadow, 116  
starpu\_vector\_filter\_divide\_in\_2, 117  
starpu\_vector\_filter\_list, 116  
deinit  
    starpu\_worker\_collection, 89  
deinit\_sched  
    starpu\_sched\_policy, 176  
destroy  
    starpu\_task, 127  
detach  
    starpu\_task, 127  
display  
    starpu\_data\_interface\_ops, 99  
double\_max\_value  
    starpu\_top\_param, 168  
double\_min\_value  
    starpu\_top\_param, 167  
dyn\_handles  
    starpu\_task, 125  
dyn\_interfaces  
    starpu\_task, 125  
dyn\_modes  
    starpu\_codelet, 123  
  
end\_ctx  
    sc\_hypervisor\_policy, 183  
enum\_values  
    starpu\_top\_param, 167  
execute\_on\_a\_specific\_worker  
    starpu\_task, 126  
Expert Mode, 165  
    starpu\_progression\_hook\_deregister, 165  
    starpu\_progression\_hook\_register, 165  
    starpu\_wake\_all\_blocked\_workers, 165  
Explicit Dependencies, 133  
    starpu\_tag\_declare\_deps, 134  
    starpu\_tag\_declare\_deps\_array, 134  
    starpu\_tag\_notify\_from\_apps, 135  
    starpu\_tag\_remove, 135  
    starpu\_tag\_restart, 134  
    starpu\_tag\_t, 134  
    starpu\_tag\_wait, 134  
    starpu\_tag\_wait\_array, 134  
    starpu\_task\_declare\_deps\_array, 134  
FFT Support, 154  
    starpuft\_cleanup, 155  
    starpuft\_destroy\_plan, 155  
    starpuft\_execute, 155  
    starpuft\_execute\_handle, 155  
    starpuft\_free, 154  
    starpuft\_malloc, 154  
    starpuft\_plan\_dft\_1d, 154  
    starpuft\_plan\_dft\_2d, 155  
    starpuft\_start, 155  
    starpuft\_start\_handle, 155  
filter\_arg  
    starpu\_data\_filter, 115  
filter\_arg\_ptr  
    starpu\_data\_filter, 115  
filter\_func  
    starpu\_data\_filter, 115  
flops  
    starpu\_task, 127  
footprint  
    starpu\_data\_interface\_ops, 99  
free\_data\_on\_node  
    starpu\_data\_interface\_ops, 99  
FxT Support, 152  
    starpu\_fxt\_generate\_trace, 154  
    starpu\_fxt\_options\_init, 154  
    starpu\_fxt\_start\_profiling, 154  
    starpu\_fxt\_stop\_profiling, 154  
  
get\_child\_ops  
    starpu\_data\_filter, 115  
get\_mf\_ops  
    starpu\_data\_interface\_ops, 99  
get\_nchildren  
    starpu\_data\_filter, 115  
get\_next  
    starpu\_worker\_collection, 89  
get\_size  
    starpu\_data\_interface\_ops, 99  
  
handle\_idle\_cycle  
    sc\_hypervisor\_policy, 182  
handle\_idle\_end  
    sc\_hypervisor\_policy, 183  
handle\_popped\_task  
    sc\_hypervisor\_policy, 183  
handle\_post\_exec\_hook  
    sc\_hypervisor\_policy, 183  
handle\_pushed\_task  
    sc\_hypervisor\_policy, 182  
handle\_submitted\_job  
    sc\_hypervisor\_policy, 183  
handle\_to\_pointer  
    starpu\_data\_interface\_ops, 99  
handles  
    starpu\_task, 125  
has\_next  
    starpu\_worker\_collection, 89  
history

starpu\_perfmodel\_per\_arch, 138  
 hypervisor\_tag  
     starpu\_task, 127  
  
 id  
     starpu\_top\_param, 167  
 Implicit Data Dependencies, 135  
     starpu\_data\_get\_default\_sequential\_consistency\_flag, 135  
     starpu\_data\_set\_default\_sequential\_consistency\_flag, 135  
     starpu\_data\_set\_sequential\_consistency\_flag, 135  
 init  
     starpu\_worker\_collection, 89  
 init\_iterator  
     starpu\_worker\_collection, 89  
 init\_sched  
     starpu\_sched\_policy, 176  
 Initialization and Termination, 75  
     starpu\_asynchronous\_copy\_disabled, 79  
     starpu\_asynchronous\_cuda\_copy\_disabled, 79  
     starpu\_asynchronous\_opencl\_copy\_disabled, 79  
     starpu\_conf\_init, 78  
     starpu\_init, 78  
     starpu\_shutdown, 79  
     starpu\_topology\_print, 79  
 Insert\_Task, 131  
     STARPU\_CALLBACK, 132  
     STARPU\_DATA\_ARRAY, 132  
     STARPU\_FLOPS, 132  
     STARPU\_PRIORITY, 132  
     STARPU\_SCHED\_CTX, 132  
     STARPU\_TAG, 132  
     STARPU\_VALUE, 132  
     starpu\_codelet\_pack\_args, 133  
     starpu\_codelet\_unpack\_args, 133  
     starpu\_insert\_task, 133  
     starpu\_task\_build, 133  
 int\_max\_value  
     starpu\_top\_param, 167  
 int\_min\_value  
     starpu\_top\_param, 167  
 interface\_size  
     starpu\_data\_interface\_ops, 99  
 interfaceid  
     starpu\_data\_interface\_ops, 99  
 interfaces  
     starpu\_task, 125  
 is\_loaded  
     starpu\_perfmodel, 137  
 is\_multiformat  
     starpu\_data\_interface\_ops, 99  
  
 list  
     starpu\_perfmodel\_per\_arch, 138  
  
 MPI Support, 155  
     STARPU\_USE\_MPI, 156  
     starpu\_data\_get\_rank, 159  
  
     starpu\_data\_get\_tag, 159  
     starpu\_data\_set\_rank, 159  
     starpu\_data\_set\_tag, 159  
     starpu\_mpi\_barrier, 158  
     starpu\_mpi\_cache\_flush, 158  
     starpu\_mpi\_cache\_flush\_all\_data, 158  
     starpu\_mpi\_comm\_amounts\_retrieve, 157  
     starpu\_mpi\_gather\_detached, 160  
     starpu\_mpi\_get\_data\_on\_node, 159  
     starpu\_mpi\_get\_data\_on\_node\_detached, 159  
     starpu\_mpi\_init, 157  
     starpu\_mpi\_initialize, 157  
     starpu\_mpi\_initialize\_extended, 157  
     starpu\_mpi\_insert\_task, 159  
     starpu\_mpi\_irecv, 157  
     starpu\_mpi\_irecv\_array\_detached\_unlock\_tag, 158  
     starpu\_mpi\_irecv\_detached, 158  
     starpu\_mpi\_irecv\_detached\_unlock\_tag, 158  
     starpu\_mpi\_isend, 157  
     starpu\_mpi\_isend\_array\_detached\_unlock\_tag, 158  
     starpu\_mpi\_isend\_detached, 157  
     starpu\_mpi\_isend\_detached\_unlock\_tag, 158  
     starpu\_mpi\_recv, 157  
     starpu\_mpi\_redux\_data, 159  
     starpu\_mpi\_scatter\_detached, 160  
     starpu\_mpi\_send, 157  
     starpu\_mpi\_shutdown, 157  
     starpu\_mpi\_test, 158  
     starpu\_mpi\_wait, 158  
 magic  
     starpu\_task, 127  
 max\_parallelism  
     starpu\_codelet, 122  
 mf\_skip  
     starpu\_task, 127  
 Miscellaneous Helpers, 152  
     starpu\_data\_cpy, 152  
     starpu\_execute\_on\_each\_worker, 152  
 model  
     starpu\_codelet, 123  
 model\_rwlock  
     starpu\_perfmodel, 137  
 modes  
     starpu\_codelet, 123  
 Multiformat Data Interface, 118  
     starpu\_multiformat\_data\_register, 120  
  
 name  
     sc\_hypervisor\_policy, 182  
     starpu\_codelet, 124  
     starpu\_top\_param, 167  
 nb\_values  
     starpu\_top\_param, 167  
 nbuffers  
     starpu\_codelet, 123  
 nchildren  
     starpu\_data\_filter, 115  
 next  
     starpu\_task, 127

starpu\_top\_param, 168  
notify\_delete\_context  
    starpu\_sched\_ctx\_performance\_counters, 171  
notify\_idle\_cycle  
    starpu\_sched\_ctx\_performance\_counters, 171  
notify\_poped\_task  
    starpu\_sched\_ctx\_performance\_counters, 171  
notify\_pushed\_task  
    starpu\_sched\_ctx\_performance\_counters, 171  
notify\_submitted\_job  
    starpu\_sched\_ctx\_performance\_counters, 171  
nworkers  
    starpu\_worker\_collection, 89

OpenCL Extensions, 146  
    STARPU\_USE\_OPENCL, 149  
    starpu\_opencl\_allocate\_memory, 151  
    starpu\_opencl\_collect\_stats, 151  
    starpu\_opencl\_compile\_opencl\_from\_file, 150  
    starpu\_opencl\_compile\_opencl\_from\_string, 150  
    starpu\_opencl\_copy\_async\_sync, 152  
    starpu\_opencl\_copy\_opencl\_to\_opencl, 151  
    starpu\_opencl\_copy\_opencl\_to\_ram, 151  
    starpu\_opencl\_copy\_ram\_to\_opencl, 151  
    starpu\_opencl\_display\_error, 151  
    starpu\_opencl\_error\_string, 151  
    starpu\_opencl\_get\_context, 149  
    starpu\_opencl\_get\_current\_context, 149  
    starpu\_opencl\_get\_current\_queue, 149  
    starpu\_opencl\_get\_device, 149  
    starpu\_opencl\_get\_queue, 149  
    starpu\_opencl\_load\_binary\_opencl, 150  
    starpu\_opencl\_load\_kernel, 150  
    starpu\_opencl\_load\_opencl\_from\_file, 150  
    starpu\_opencl\_load\_opencl\_from\_string, 150  
    starpu\_opencl\_load\_program\_source, 150  
    starpu\_opencl\_release\_kernel, 151  
    starpu\_opencl\_report\_error, 151  
    starpu\_opencl\_set\_kernel\_args, 149  
    starpu\_opencl\_unload\_opencl, 150

opencl\_func  
    starpu\_codelet, 123

opencl\_funcs  
    starpu\_codelet, 123

opencl\_to\_cuda  
    starpu\_data\_copy\_methods, 101

opencl\_to\_opencl  
    starpu\_data\_copy\_methods, 101

opencl\_to\_opencl\_async  
    starpu\_data\_copy\_methods, 101

opencl\_to\_ram  
    starpu\_data\_copy\_methods, 101

opencl\_to\_ram\_async  
    starpu\_data\_copy\_methods, 101

pack\_data  
    starpu\_data\_interface\_ops, 100

Parallel Tasks, 164  
    starpu\_combined\_worker\_assign\_workerid, 164

starpu\_combined\_worker\_can\_execute\_task, 164  
starpu\_combined\_worker\_get\_count, 164  
starpu\_combined\_worker\_get\_description, 164  
starpu\_combined\_worker\_get\_id, 164  
starpu\_combined\_worker\_get\_rank, 164  
starpu\_combined\_worker\_get\_size, 164  
starpu\_parallel\_task\_barrier\_init, 164

per\_arch  
    starpu\_perfmodel, 137

per\_worker\_stats  
    starpu\_codelet, 124

Performance Model, 136  
    STARPU\_COMMON, 139  
    STARPU\_CPU\_DEFAULT, 139  
    STARPU\_CUDA\_DEFAULT, 139  
    STARPU\_HISTORY\_BASED, 139  
    STARPU\_NL\_REGRESSION\_BASED, 140  
    STARPU\_OPENCL\_DEFAULT, 139  
    STARPU\_PER\_ARCH, 139  
    STARPU\_REGRESSION\_BASED, 140  
    starpu\_bus\_print\_affinity, 140  
    starpu\_bus\_print\_bandwidth, 140  
    starpu\_perfmodel\_archtype, 139  
    starpu\_perfmodel\_debugfilepath, 140  
    starpu\_perfmodel\_directory, 140  
    starpu\_perfmodel\_get\_arch\_name, 140  
    starpu\_perfmodel\_list, 140  
    starpu\_perfmodel\_load\_symbol, 140  
    starpu\_perfmodel\_print, 140  
    starpu\_perfmodel\_print\_all, 140  
    starpu\_perfmodel\_type, 139  
    starpu\_perfmodel\_unload\_model, 140  
    starpu\_perfmodel\_update\_history, 140  
    starpu\_transfer\_bandwidth, 141  
    starpu\_transfer\_latency, 141  
    starpu\_transfer\_predict, 141  
    starpu\_worker\_get\_perf\_archtype, 140

policy\_description  
    starpu\_sched\_policy, 176

policy\_name  
    starpu\_sched\_policy, 176

pop\_every\_task  
    starpu\_sched\_policy, 176

pop\_task  
    starpu\_sched\_policy, 176

post\_exec\_hook  
    starpu\_sched\_policy, 176

power\_model  
    starpu\_codelet, 124

pre\_exec\_hook  
    starpu\_sched\_policy, 176

predicted  
    starpu\_task, 127

predicted\_transfer  
    starpu\_task, 127

prev  
    starpu\_task, 127

priority

**starpu\_task**, 126  
**Profiling**, 141  
  **starpu\_bus\_get\_count**, 143  
  **starpu\_bus\_get\_dst**, 143  
  **starpu\_bus\_get\_id**, 143  
  **starpu\_bus\_get\_profiling\_info**, 143  
  **starpu\_bus\_get\_src**, 143  
  **starpu\_data\_display\_memory\_stats**, 144  
  **starpu\_profiling\_bus\_helper\_display\_summary**, 144  
  **starpu\_profiling\_init**, 143  
  **starpu\_profiling\_set\_id**, 143  
  **starpu\_profiling\_status\_get**, 143  
  **starpu\_profiling\_status\_set**, 143  
  **starpu\_profiling\_worker\_get\_info**, 143  
  **starpu\_profiling\_worker\_helper\_display\_summary**, 144  
  **starpu\_timing\_timespec\_delay\_us**, 144  
  **starpu\_timing\_timespec\_to\_us**, 144  
**profiling\_info**  
  **starpu\_task**, 127  
**prologue\_callback\_arg**  
  **starpu\_task**, 126  
**prologue\_callback\_arg\_free**  
  **starpu\_task**, 126  
**prologue\_callback\_func**  
  **starpu\_task**, 126  
**push\_task**  
  **starpu\_sched\_policy**, 176  
**push\_task\_notify**  
  **starpu\_sched\_policy**, 176  
  
**ram\_to\_cuda**  
  **starpu\_data\_copy\_methods**, 100  
**ram\_to\_cuda\_async**  
  **starpu\_data\_copy\_methods**, 101  
**ram\_to\_opencl**  
  **starpu\_data\_copy\_methods**, 100  
**ram\_to\_opencl\_async**  
  **starpu\_data\_copy\_methods**, 101  
**ram\_to\_ram**  
  **starpu\_data\_copy\_methods**, 100  
**regenerate**  
  **starpu\_task**, 127  
**register\_data\_handle**  
  **starpu\_data\_interface\_ops**, 99  
**regression**  
  **starpu\_perfmodel\_per\_arch**, 138  
**remove**  
  **starpu\_worker\_collection**, 89  
**remove\_workers**  
  **starpu\_sched\_policy**, 176  
**resize\_ctxs**  
  **sc\_hypervisor\_policy**, 182  
**Running Drivers**, 165  
  **starpu\_driver\_deinit**, 165  
  **starpu\_driver\_init**, 165  
  **starpu\_driver\_run**, 165  
  **starpu\_driver\_run\_once**, 165  
  **starpu\_drivers\_request\_termination**, 165  
  
**STARPU\_ANY\_WORKER**  
  Workers' Properties, 90  
**STARPU\_BCSR\_INTERFACE\_ID**  
  Data Interfaces, 108  
**STARPU\_BLOCK\_INTERFACE\_ID**  
  Data Interfaces, 108  
**STARPU\_COMMON**  
  Performance Model, 139  
**STARPU\_COO\_INTERFACE\_ID**  
  Data Interfaces, 108  
**STARPU\_CPU\_DEFAULT**  
  Performance Model, 139  
**STARPU\_CPU\_RAM**  
  Workers' Properties, 89  
**STARPU\_CPU\_WORKER**  
  Workers' Properties, 90  
**STARPU\_CSR\_INTERFACE\_ID**  
  Data Interfaces, 108  
**STARPU\_CUDA\_DEFAULT**  
  Performance Model, 139  
**STARPU\_CUDA\_RAM**  
  Workers' Properties, 89  
**STARPU\_CUDA\_WORKER**  
  Workers' Properties, 90  
**STARPU\_FORKJOIN**  
  Codelet And Tasks, 129  
**STARPU\_HISTORY\_BASED**  
  Performance Model, 139  
**STARPU\_MATRIX\_INTERFACE\_ID**  
  Data Interfaces, 108  
**STARPU\_MAX\_INTERFACE\_ID**  
  Data Interfaces, 108  
**STARPU\_MULTIFORMAT\_INTERFACE\_ID**  
  Data Interfaces, 108  
**STARPU\_NL\_REGRESSION\_BASED**  
  Performance Model, 140  
**STARPU\_NONE**  
  Data Management, 93  
**STARPU\_OPENCL\_DEFAULT**  
  Performance Model, 139  
**STARPU\_OPENCL\_RAM**  
  Workers' Properties, 89  
**STARPU\_OPENCL\_WORKER**  
  Workers' Properties, 90  
**STARPU\_PER\_ARCH**  
  Performance Model, 139  
**STARPU\_R**  
  Data Management, 93  
**STARPU\_REDUX**  
  Data Management, 93  
**STARPU\_REGRESSION\_BASED**  
  Performance Model, 140  
**STARPU\_RW**  
  Data Management, 93  
**STARPU\_SCRATCH**  
  Data Management, 93  
**STARPU\_SEQ**  
  Codelet And Tasks, 129

STARPU\_SPM<sub>D</sub>  
    Codelet And Tasks, 129

STARPU\_TASK\_BLOCKED  
    Codelet And Tasks, 129

STARPU\_TASK\_BLOCKED\_ON\_DATA  
    Codelet And Tasks, 129

STARPU\_TASK\_BLOCKED\_ON\_TAG  
    Codelet And Tasks, 129

STARPU\_TASK\_BLOCKED\_ON\_TASK  
    Codelet And Tasks, 129

STARPU\_TASK\_FINISHED  
    Codelet And Tasks, 129

STARPU\_TASK\_READY  
    Codelet And Tasks, 129

STARPU\_TASK\_RUNNING  
    Codelet And Tasks, 129

STARPU\_TOP\_DATA\_BOOLEAN  
    StarPU-Top Interface, 168

STARPU\_TOP\_DATA\_FLOAT  
    StarPU-Top Interface, 168

STARPU\_TOP\_DATA\_INTEGER  
    StarPU-Top Interface, 168

STARPU\_TOP\_PARAM\_BOOLEAN  
    StarPU-Top Interface, 168

STARPU\_TOP\_PARAM\_ENUM  
    StarPU-Top Interface, 168

STARPU\_TOP\_PARAM\_FLOAT  
    StarPU-Top Interface, 168

STARPU\_TOP\_PARAM\_INTEGER  
    StarPU-Top Interface, 168

STARPU\_UNKNOWN\_INTERFACE\_ID  
    Data Interfaces, 108

STARPU\_UNUSED  
    Workers' Properties, 89

STARPU\_VARIABLE\_INTERFACE\_ID  
    Data Interfaces, 108

STARPU\_VECTOR\_INTERFACE\_ID  
    Data Interfaces, 108

STARPU\_VOID\_INTERFACE\_ID  
    Data Interfaces, 108

STARPU\_W  
    Data Management, 93

STARPU\_WORKER\_LIST  
    Workers' Properties, 90

SC\_HYPERVISOR\_NULL  
    Scheduling Context Hypervisor - Regular usage, 180

STARPU\_ABORT  
    Toolbox, 82

STARPU\_ABORT\_MSG  
    Toolbox, 82

STARPU\_ASSERT  
    Toolbox, 82

STARPU\_ASSERT\_MSG  
    Toolbox, 82

STARPU\_BCSR\_GET\_NNZ  
    Data Interfaces, 106

STARPU\_BLOCK\_GET\_NX  
    Data Interfaces, 105

STARPU\_BLOCK\_GET\_NY  
    Data Interfaces, 106

STARPU\_BLOCK\_GET\_NZ  
    Data Interfaces, 106

STARPU\_CALLBACK  
    Insert\_Task, 132

STARPU\_COO\_GET\_NX  
    Data Interfaces, 108

STARPU\_COO\_GET\_NY  
    Data Interfaces, 108

STARPU\_COO\_GET\_ROWS  
    Data Interfaces, 107

STARPU\_CPU  
    Codelet And Tasks, 128

STARPU\_CSR\_GET\_NNZ  
    Data Interfaces, 106

STARPU\_CSR\_GET\_NROW  
    Data Interfaces, 107

STARPU\_CUDA  
    Codelet And Tasks, 128

STARPU\_DATA\_ARRAY  
    Insert\_Task, 132

STARPU\_DEFAULT\_PRIO  
    Scheduling Contexts, 172

STARPU\_FLOPS  
    Insert\_Task, 132

STARPU\_GNUC\_PREREQ  
    Toolbox, 81

STARPU\_LIKELY  
    Toolbox, 81

STARPU\_MAJOR\_VERSION  
    Versioning, 75

STARPU\_MALLOC\_COUNT  
    Standard Memory Library, 80

STARPU\_MAX  
    Toolbox, 82

STARPU\_MAX\_PRIO  
    Scheduling Contexts, 172

STARPU\_MAXCUDADEVS  
    CUDA Extensions, 145

STARPU\_MAXOPENCLDEVS  
    OpenCL Extensions, 149

STARPU\_MIN  
    Toolbox, 82

STARPU\_MIN\_PRIO  
    Scheduling Contexts, 172

STARPU\_MINOR\_VERSION  
    Versioning, 75

STARPU\_NMAXBUFS  
    Codelet And Tasks, 128

STARPU\_NMAXWORKERS  
    Workers' Properties, 89

STARPU\_OPENCL  
    Codelet And Tasks, 128

STARPU\_POISON\_PTR  
    Toolbox, 82

STARPU\_PRIORITY

Insert\_Task, 132  
**STARPU\_RMB**  
 Toolbox, 82  
**STARPU\_SCHED\_CTX**  
 Insert\_Task, 132  
**STARPU\_TAG**  
 Insert\_Task, 132  
**STARPU\_TASK\_INVALID**  
 Codelet And Tasks, 129  
**STARPU\_UNLIKELY**  
 Toolbox, 81  
**STARPU\_USE\_CUDA**  
 CUDA Extensions, 145  
**STARPU\_USE\_MPI**  
 MPI Support, 156  
**STARPU\_USE\_OPENCL**  
 OpenCL Extensions, 149  
**STARPU\_VALUE**  
 Insert\_Task, 132  
**STARPU\_WMB**  
 Toolbox, 82  
**sc\_hypervisor\_policy**, 182  
**sc\_hypervisor\_policy\_config**, 183  
**sc\_hypervisor\_policy\_task\_pool**, 185  
**sc\_hypervisor\_resize\_ack**, 185  
**sc\_hypervisor\_wrapper**, 183  
**sc\_hypervisor.h**, 214  
**sc\_hypervisor\_add\_workers\_to\_sched\_ctx**  
 Scheduling Context Hypervisor - Regular usage, 181  
**sc\_hypervisor\_can\_resize**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_config.h**, 215  
**sc\_hypervisor\_ctl**  
 Scheduling Context Hypervisor - Regular usage, 181  
**sc\_hypervisor\_free\_size\_req**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_get\_config**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_get\_elapsed\_flops\_per\_sched\_ctx**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_get\_nsched\_ctxs**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_get\_policy**  
 Scheduling Context Hypervisor - Regular usage, 181  
**sc\_hypervisor\_get\_sched\_ctxs**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_get\_size\_req**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_get\_wrapper**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_init**  
 Scheduling Context Hypervisor - Regular usage, 180  
**sc\_hypervisor\_lp.h**, 215  
**sc\_hypervisor\_monitoring.h**, 216  
**sc\_hypervisor\_move\_workers**  
 Scheduling Context Hypervisor - Regular usage, 181  
**sc\_hypervisor\_policy**  
 custom, 182  
 end\_ctx, 183  
 handle\_idle\_cycle, 182  
 handle\_idle\_end, 183  
 handle\_popped\_task, 183  
 handle\_post\_exec\_hook, 183  
 handle\_pushed\_task, 182  
 handle\_submitted\_job, 183  
 name, 182  
 resize\_ctxs, 182  
 size\_ctxs, 182  
**sc\_hypervisor\_policy.h**, 216  
**sc\_hypervisor\_post\_resize\_request**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_register\_ctx**  
 Scheduling Context Hypervisor - Regular usage, 180  
**sc\_hypervisor\_remove\_workers\_from\_sched\_ctx**  
 Scheduling Context Hypervisor - Regular usage, 181  
**sc\_hypervisor\_save\_size\_req**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_set\_config**  
 Scheduling Context Hypervisor - Building a new resizing policy, 186  
**sc\_hypervisor\_set\_type\_of\_task**  
 Scheduling Context Hypervisor - Regular usage, 181  
**sc\_hypervisor\_shutdown**  
 Scheduling Context Hypervisor - Regular usage, 180  
**sc\_hypervisor\_size\_ctxs**  
 Scheduling Context Hypervisor - Regular usage, 181  
**sc\_hypervisor\_start\_resize**  
 Scheduling Context Hypervisor - Regular usage, 181  
**sc\_hypervisor\_stop\_resize**  
 Scheduling Context Hypervisor - Regular usage, 180  
**sc\_hypervisor\_unregister\_ctx**  
 Scheduling Context Hypervisor - Regular usage, 180  
**sc\_hypervisor\_update\_diff\_elapsed\_flops**

Scheduling Context Hypervisor - Regular usage, 181  
 sc\_hypervisor\_update\_diff\_total\_flops  
     Scheduling Context Hypervisor - Regular usage, 181  
 sched\_ctx  
     starpu\_task, 127  
 scheduled  
     starpu\_task, 127  
 Scheduling Context Hypervisor - Building a new resizing policy, 181  
     sc\_hypervisor\_can\_resize, 186  
     sc\_hypervisor\_free\_size\_req, 186  
     sc\_hypervisor\_get\_config, 186  
     sc\_hypervisor\_get\_elapsed\_flops\_per\_sched\_ctx, 186  
     sc\_hypervisor\_get\_nsched\_ctxs, 186  
     sc\_hypervisor\_get\_sched\_ctxs, 186  
     sc\_hypervisor\_get\_size\_req, 186  
     sc\_hypervisor\_get\_wrapper, 186  
     sc\_hypervisor\_post\_resize\_request, 186  
     sc\_hypervisor\_save\_size\_req, 186  
     sc\_hypervisor\_set\_config, 186  
 Scheduling Context Hypervisor - Regular usage, 178  
     SC\_HYPERVISOR\_NULL, 180  
     sc\_hypervisor\_add\_workers\_to\_sched\_ctx, 181  
     sc\_hypervisor\_ctl, 181  
     sc\_hypervisor\_get\_policy, 181  
     sc\_hypervisor\_init, 180  
     sc\_hypervisor\_move\_workers, 181  
     sc\_hypervisor\_register\_ctx, 180  
     sc\_hypervisor\_remove\_workers\_from\_sched\_ctx, 181  
     sc\_hypervisor\_set\_type\_of\_task, 181  
     sc\_hypervisor\_shutdown, 180  
     sc\_hypervisor\_size\_ctxs, 181  
     sc\_hypervisor\_start\_resize, 181  
     sc\_hypervisor\_stop\_resize, 180  
     sc\_hypervisor\_unregister\_ctx, 180  
     sc\_hypervisor\_update\_diff\_elapsed\_flops, 181  
     sc\_hypervisor\_update\_diff\_total\_flops, 181  
 Scheduling Contexts, 170  
     STARPU\_DEFAULT\_PRIO, 172  
     STARPU\_MAX\_PRIO, 172  
     STARPU\_MIN\_PRIO, 172  
     starpu\_sched\_ctx.add\_workers, 172  
     starpu\_sched\_ctx.call\_pushed\_task\_cb, 174  
     starpu\_sched\_ctx.check\_if\_hypervisor\_exists, 174  
     starpu\_sched\_ctx.contains\_worker, 173  
     starpu\_sched\_ctx.create, 172  
     starpu\_sched\_ctx.create\_inside\_interval, 172  
     starpu\_sched\_ctx.create\_worker\_collection, 174  
     starpu\_sched\_ctx.delete, 173  
     starpu\_sched\_ctx.delete\_worker\_collection, 174  
     starpu\_sched\_ctx.finished\_submit, 173  
     starpu\_sched\_ctx.get\_context, 173  
     starpu\_sched\_ctx.get\_max\_priority, 174  
     starpu\_sched\_ctx.get\_min\_priority, 174  
     starpu\_sched\_ctx.get\_nshared\_workers, 173  
     starpu\_sched\_ctx.get\_nworkers, 173  
     starpu\_sched\_ctx.get\_policy\_data, 174  
     starpu\_sched\_ctx.get\_worker\_collection, 174  
     starpu\_sched\_ctx.get\_workers\_list, 173  
     starpu\_sched\_ctx.notify\_hypervisor\_exists, 174  
     starpu\_sched\_ctx.overlapping\_ctxs\_on\_worker, 173  
     starpu\_sched\_ctx.register\_close\_callback, 172  
     starpu\_sched\_ctx.remove\_workers, 173  
     starpu\_sched\_ctx.set\_context, 173  
     starpu\_sched\_ctx.set\_inheritor, 173  
     starpu\_sched\_ctx.set\_max\_priority, 174  
     starpu\_sched\_ctx.set\_min\_priority, 174  
     starpu\_sched\_ctx.set\_perf\_counters, 174  
     starpu\_sched\_ctx.set\_policy\_data, 174  
     starpu\_sched\_ctx.stop\_task\_submission, 173  
     starpu\_sched\_ctx.worker\_get\_id, 173  
 Scheduling Policy, 175  
     starpu\_data\_expected\_transfer\_time, 178  
     starpu\_get\_prefetch\_flag, 178  
     starpu\_prefetch\_task\_input\_on\_node, 178  
     starpu\_push\_local\_task, 177  
     starpu\_push\_task\_end, 177  
     starpu\_sched\_ctx.worker\_shares\_tasks\_lists, 178  
     starpu\_sched\_get\_max\_priority, 177  
     starpu\_sched\_get\_min\_priority, 177  
     starpu\_sched\_get\_predefined\_policies, 177  
     starpu\_sched\_set\_max\_priority, 177  
     starpu\_sched\_set\_min\_priority, 177  
     starpu\_task\_expected\_conversion\_time, 178  
     starpu\_task\_expected\_data\_transfer\_time, 178  
     starpu\_task\_expected\_length, 177  
     starpu\_task\_expected\_power, 178  
     starpu\_task\_footprint, 177  
     starpu\_timing\_now, 177  
     starpu\_worker\_can\_execute\_task, 177  
     starpu\_worker\_get\_relative\_speedup, 178  
     starpu\_worker\_get\_sched\_condition, 177  
 sequential\_consistency  
     starpu\_task, 126  
 size\_base  
     starpu\_perfmodel, 137  
     starpu\_perfmodel\_per\_arch, 138  
 size\_ctxs  
     sc\_hypervisor\_policy, 182  
 Standard Memory Library, 79  
     STARPU\_MALLOC\_COUNT, 80  
     starpu\_data\_free\_pinned\_if\_possible, 79  
     starpu\_data\_malloc\_pinned\_if\_possible, 79  
     starpu\_free, 80  
     starpu\_free\_flags, 80  
     starpu\_malloc, 80  
     starpu\_malloc\_flags, 80  
     starpu\_malloc\_set\_align, 80  
     starpu\_memory\_get\_available, 80  
 StarPU-Top Interface  
     STARPU\_TOP\_DATA\_BOOLEAN, 168  
     STARPU\_TOP\_DATA\_FLOAT, 168

STARPU\_TOP\_DATA\_INTEGER, 168  
 STARPU\_TOP\_PARAM\_BOOLEAN, 168  
 STARPU\_TOP\_PARAM\_ENUM, 168  
 STARPU\_TOP\_PARAM\_FLOAT, 168  
 STARPU\_TOP\_PARAM\_INTEGER, 168  
 TOP\_TYPE\_CONTINUE, 168  
 TOP\_TYPE\_DEBUG, 168  
 TOP\_TYPE\_DISABLE, 168  
 TOP\_TYPE\_ENABLE, 168  
 TOP\_TYPE\_GO, 168  
 TOP\_TYPE\_SET, 168  
 TOP\_TYPE\_UNKNOW, 168  
 StarPU-Top Interface, 166  
     starpu\_top\_add\_data\_boolean, 168  
     starpu\_top\_add\_data\_float, 168  
     starpu\_top\_add\_data\_integer, 168  
     starpu\_top\_data\_type, 168  
     starpu\_top\_debug\_lock, 170  
     starpu\_top\_debug\_log, 169  
     starpu\_top\_init\_and\_wait, 169  
     starpu\_top\_message\_type, 168  
     starpu\_top\_param\_type, 168  
     starpu\_top\_register\_parameter\_boolean, 169  
     starpu\_top\_register\_parameter\_enum, 169  
     starpu\_top\_register\_parameter\_float, 169  
     starpu\_top\_register\_parameter\_integer, 169  
     starpu\_top\_task\_prevision, 169  
     starpu\_top\_update\_data\_boolean, 169  
     starpu\_top\_update\_data\_float, 169  
     starpu\_top\_update\_data\_integer, 169  
     starpu\_top\_update\_parameter, 169  
 starpu.h, 189  
 starpu\_bcsr\_interface, 103  
 starpu\_block\_interface, 103  
 starpu\_codelet, 122  
 starpu\_conf, 76  
 starpu\_coo\_interface, 104  
 starpu\_csr\_interface, 104  
 starpu\_data\_copy\_methods, 100  
 starpu\_data\_descr, 124  
 starpu\_data\_filter, 115  
 starpu\_data\_interface\_ops, 99  
 starpu\_driver, 76  
 starpu\_driver.id, 76  
 starpu\_fxt\_codelet\_event, 153  
 starpu\_fxt\_options, 153  
 starpu\_matrix\_interface, 103  
 starpu\_multiformat\_data\_interface\_ops, 119  
 starpu\_multiformat\_interface, 120  
 starpu\_opencl\_program, 148  
 starpu\_perfmodel, 136  
 starpu\_perfmodel\_history\_entry, 139  
 starpu\_perfmodel\_history\_list, 138  
 starpu\_perfmodel\_per\_arch, 138  
 starpu\_perfmodel\_regression\_model, 137  
 starpu\_profiling\_bus\_info, 142  
 starpu\_profiling\_task\_info, 141  
 starpu\_profiling\_worker\_info, 142  
     starpu\_sched\_ctx\_iterator, 89  
     starpu\_sched\_ctx\_performance\_counters, 171  
     starpu\_sched\_policy, 175  
     starpu\_task, 124  
     starpu\_task\_list, 162  
     starpu\_top\_data, 166  
     starpu\_top\_param, 167  
     starpu\_variable\_interface, 101  
     starpu\_vector\_interface, 76, 103  
     starpu\_worker\_collection, 88  
     starpu\_asynchronous\_copy\_disabled  
         Initialization and Termination, 79  
     starpu\_asynchronous\_cuda\_copy\_disabled  
         Initialization and Termination, 79  
     starpu\_asynchronous\_opencl\_copy\_disabled  
         Initialization and Termination, 79  
     starpu\_bcsr\_data\_register  
         Data Interfaces, 109  
     starpu\_bcsr\_filter\_canonical\_block  
         Data Partition, 118  
     starpu\_bcsr\_get\_c  
         Data Interfaces, 112  
     starpu\_bcsr\_get\_elemsize  
         Data Interfaces, 112  
     starpu\_bcsr\_get\_firstentry  
         Data Interfaces, 112  
     starpu\_bcsr\_get\_local\_colind  
         Data Interfaces, 112  
     starpu\_bcsr\_get\_local\_nzval  
         Data Interfaces, 112  
     starpu\_bcsr\_get\_local\_rowptr  
         Data Interfaces, 112  
     starpu\_bcsr\_get\_nnz  
         Data Interfaces, 112  
     starpu\_bcsr\_get\_nrow  
         Data Interfaces, 112  
     starpu\_bcsr\_get\_r  
         Data Interfaces, 112  
     starpu\_block\_data\_register  
         Data Interfaces, 109  
     starpu\_block\_filter\_block  
         Data Partition, 117  
     starpu\_block\_filter\_block\_shadow  
         Data Partition, 117  
     starpu\_block\_filter\_depth\_block  
         Data Partition, 118  
     starpu\_block\_filter\_depth\_block\_shadow  
         Data Partition, 118  
     starpu\_block\_filter\_vertical\_block  
         Data Partition, 117  
     starpu\_block\_filter\_vertical\_block\_shadow  
         Data Partition, 118  
     starpu\_block\_get\_elemsize  
         Data Interfaces, 112  
     starpu\_block\_get\_local\_ldy  
         Data Interfaces, 111  
     starpu\_block\_get\_local\_ldz  
         Data Interfaces, 111

starpu\_block\_get\_local\_ptr  
    Data Interfaces, 112  
starpu\_block\_get\_nx  
    Data Interfaces, 111  
starpu\_block\_get\_ny  
    Data Interfaces, 111  
starpu\_block\_get\_nz  
    Data Interfaces, 111  
starpu\_bound.h, 190  
starpu\_bound\_compute  
    Theoretical Lower Bound on Execution Time, 144  
starpu\_bound\_print  
    Theoretical Lower Bound on Execution Time, 145  
starpu\_bound\_print\_dot  
    Theoretical Lower Bound on Execution Time, 144  
starpu\_bound\_print\_lp  
    Theoretical Lower Bound on Execution Time, 145  
starpu\_bound\_print\_mps  
    Theoretical Lower Bound on Execution Time, 145  
starpu\_bound\_start  
    Theoretical Lower Bound on Execution Time, 144  
starpu\_bound\_stop  
    Theoretical Lower Bound on Execution Time, 144  
starpu\_bus\_get\_count  
    Profiling, 143  
starpu\_bus\_get\_dst  
    Profiling, 143  
starpu\_bus\_get\_id  
    Profiling, 143  
starpu\_bus\_get\_profiling\_info  
    Profiling, 143  
starpu\_bus\_get\_src  
    Profiling, 143  
starpu\_bus\_print\_affinity  
    Performance Model, 140  
starpu\_bus\_print\_bandwidth  
    Performance Model, 140  
starpu\_codelet  
    can\_execute, 122  
    cpu\_func, 123  
    cpu\_funcs, 123  
    cuda\_func, 123  
    cuda\_funcs, 123  
    dyn\_modes, 123  
    max\_parallelism, 122  
    model, 123  
    modes, 123  
    name, 124  
    nbuffers, 123  
    opencl\_func, 123  
    opencl\_funcs, 123  
    per\_worker\_stats, 124  
    power\_model, 124  
    type, 122  
    where, 122  
starpu\_codelet\_display\_stats  
    Codelet And Tasks, 131  
starpu\_codelet\_init  
    Codelet And Tasks, 130  
starpu\_codelet\_pack\_args  
    Insert\_Task, 133  
starpu\_codelet\_type  
    Codelet And Tasks, 129  
starpu\_codelet\_unpack\_args  
    Insert\_Task, 133  
starpu\_combined\_worker\_assign\_workerid  
    Parallel Tasks, 164  
starpu\_combined\_worker\_can\_execute\_task  
    Parallel Tasks, 164  
starpu\_combined\_worker\_get\_count  
    Parallel Tasks, 164  
starpu\_combined\_worker\_get\_description  
    Parallel Tasks, 164  
starpu\_combined\_worker\_get\_id  
    Parallel Tasks, 164  
starpu\_combined\_worker\_get\_rank  
    Parallel Tasks, 164  
starpu\_combined\_worker\_get\_size  
    Parallel Tasks, 164  
starpu\_conf\_init  
    Initialization and Termination, 78  
starpu\_config.h, 190  
starpu\_coo\_data\_register  
    Data Interfaces, 110  
starpu\_cpu\_func\_t  
    Codelet And Tasks, 129  
starpu\_cpu\_worker\_get\_count  
    Workers' Properties, 90  
starpu\_create\_sync\_task  
    Codelet And Tasks, 131  
starpu\_csr\_data\_register  
    Data Interfaces, 110  
starpu\_csr\_filter\_vertical\_block  
    Data Partition, 118  
starpu\_csr\_get\_elemsize  
    Data Interfaces, 113  
starpu\_csr\_get\_firstryentry  
    Data Interfaces, 112  
starpu\_csr\_get\_local\_colind  
    Data Interfaces, 113  
starpu\_csr\_get\_local\_nzval  
    Data Interfaces, 112  
starpu\_csr\_get\_local\_rowptr  
    Data Interfaces, 113  
starpu\_csr\_get\_nnz  
    Data Interfaces, 112  
starpu\_csr\_get\_nrow  
    Data Interfaces, 112  
starpu\_cublas.h, 191  
starpu\_cublas\_init  
    CUDA Extensions, 146  
starpu\_cublas\_report\_error  
    CUDA Extensions, 146  
starpu\_cublas\_shutdown  
    CUDA Extensions, 146  
starpu\_cuda.h, 191

**starpu\_cuda\_copy\_async\_sync**  
 CUDA Extensions, 146  
**starpu\_cuda\_func\_t**  
 Codelet And Tasks, 129  
**starpu\_cuda\_get\_device\_properties**  
 CUDA Extensions, 146  
**starpu\_cuda\_get\_local\_stream**  
 CUDA Extensions, 146  
**starpu\_cuda\_report\_error**  
 CUDA Extensions, 146  
**starpu\_cuda\_set\_device**  
 CUDA Extensions, 146  
**starpu\_cuda\_worker\_get\_count**  
 Workers' Properties, 90  
**starpu\_data.h**, 192  
**starpu\_data\_access\_mode**  
 Data Management, 93  
**starpu\_data\_acquire**  
 Data Management, 95  
**starpu\_data\_acquire\_cb**  
 Data Management, 95  
**starpu\_data\_acquire\_on\_node**  
 Data Management, 95  
**starpu\_data\_acquire\_on\_node\_cb**  
 Data Management, 95  
**starpu\_data\_advise\_as\_important**  
 Data Management, 94  
**starpu\_data\_copy\_methods**  
 any\_to\_any, 101  
 cuda\_to\_cuda, 100  
 cuda\_to\_cuda\_async, 101  
 cuda\_to\_opencl, 101  
 cuda\_to\_ram, 100  
 cuda\_to\_ram\_async, 101  
 opencl\_to\_cuda, 101  
 opencl\_to\_opencl, 101  
 opencl\_to\_opencl\_async, 101  
 opencl\_to\_ram, 101  
 opencl\_to\_ram\_async, 101  
 ram\_to\_cuda, 100  
 ram\_to\_cuda\_async, 101  
 ram\_to\_opencl, 100  
 ram\_to\_opencl\_async, 101  
 ram\_to\_ram, 100  
**starpu\_data\_cpy**  
 Miscellaneous Helpers, 152  
**starpu\_data\_display\_memory\_stats**  
 Profiling, 144  
**starpu\_data\_expected\_transfer\_time**  
 Scheduling Policy, 178  
**starpu\_data\_filter**  
 filter\_arg, 115  
 filter\_arg\_ptr, 115  
 filter\_func, 115  
 get\_child\_ops, 115  
 get\_nchildren, 115  
 nchildren, 115  
**starpu\_data\_filters.h**, 193  
**starpu\_data\_free\_pinned\_if\_possible**  
 Standard Memory Library, 79  
**starpu\_data\_get\_child**  
 Data Partition, 116  
**starpu\_data\_get\_default\_sequential\_consistency\_flag**  
 Implicit Data Dependencies, 135  
**starpu\_data\_get\_interface\_id**  
 Data Interfaces, 110  
**starpu\_data\_get\_interface\_on\_node**  
 Data Interfaces, 110  
**starpu\_data\_get\_local\_ptr**  
 Data Interfaces, 110  
**starpu\_data\_get\_nb\_children**  
 Data Partition, 116  
**starpu\_data\_get\_rank**  
 MPI Support, 159  
**starpu\_data\_get\_size**  
 Data Interfaces, 110  
**starpu\_data\_get\_sub\_data**  
 Data Partition, 116  
**starpu\_data\_get\_tag**  
 MPI Support, 159  
**starpu\_data\_handle\_t**  
 Data Management, 93  
**starpu\_data\_handle\_to\_pointer**  
 Data Interfaces, 110  
**starpu\_data\_interface\_get\_next\_id**  
 Data Interfaces, 113  
**starpu\_data\_interface\_id**  
 Data Interfaces, 108  
**starpu\_data\_interface\_ops**  
 allocate\_data\_on\_node, 99  
 compare, 99  
 copy\_methods, 99  
 display, 99  
 footprint, 99  
 free\_data\_on\_node, 99  
 get\_mf\_ops, 99  
 get\_size, 99  
 handle\_to\_pointer, 99  
 interface\_size, 99  
 interfaceid, 99  
 is\_multiformat, 99  
 pack\_data, 100  
 register\_data\_handle, 99  
 unpack\_data, 100  
**starpu\_data\_interfaces.h**, 195  
**starpu\_data\_invalidate**  
 Data Management, 94  
**starpu\_data\_invalidate\_submit**  
 Data Management, 94  
**starpu\_data\_lookup**  
 Data Management, 94  
**starpu\_data\_malloc\_pinned\_if\_possible**  
 Standard Memory Library, 79  
**starpu\_data\_map\_filters**  
 Data Partition, 116  
**starpu\_data\_pack**

Data Interfaces, 110  
starpu\_data\_partition  
    Data Partition, 115  
starpu\_data\_prefetch\_on\_node  
    Data Management, 94  
starpu\_data\_query\_status  
    Data Management, 94  
starpu\_data\_register  
    Data Management, 93  
starpu\_data\_register\_same  
    Data Management, 93  
starpu\_data\_release  
    Data Management, 95  
starpu\_data\_release\_on\_node  
    Data Management, 95  
starpu\_data\_request\_allocation  
    Data Management, 94  
starpu\_data\_set\_default\_sequential\_consistency\_flag  
    Implicit Data Dependencies, 135  
starpu\_data\_set\_rank  
    MPI Support, 159  
starpu\_data\_set\_reduction\_methods  
    Data Management, 94  
starpu\_data\_set\_sequential\_consistency\_flag  
    Implicit Data Dependencies, 135  
starpu\_data\_set\_tag  
    MPI Support, 159  
starpu\_data\_set\_wt\_mask  
    Data Management, 94  
starpu\_data\_unpack  
    Data Interfaces, 110  
starpu\_data\_unpartition  
    Data Partition, 115  
starpu\_data\_unregister  
    Data Management, 93  
starpu\_data\_unregister\_no\_coherency  
    Data Management, 94  
starpu\_data\_unregister\_submit  
    Data Management, 94  
starpu\_data\_vget\_sub\_data  
    Data Partition, 116  
starpu\_data\_vmap\_filters  
    Data Partition, 116  
starpu\_deprecated\_api.h, 199  
starpu\_driver.h, 199  
starpu\_driver\_deinit  
    Running Drivers, 165  
starpu\_driver\_init  
    Running Drivers, 165  
starpu\_driver\_run  
    Running Drivers, 165  
starpu\_driver\_run\_once  
    Running Drivers, 165  
starpu\_drivers\_request\_termination  
    Running Drivers, 165  
starpu\_execute\_on\_each\_worker  
    Miscellaneous Helpers, 152  
starpu\_expert.h, 199  
starpu\_free  
    Standard Memory Library, 80  
starpu\_free\_flags  
    Standard Memory Library, 80  
starpu\_free\_on\_node  
    Data Interfaces, 113  
starpu\_fxt.h, 199  
starpu\_fxt\_generate\_trace  
    Fxt Support, 154  
starpu\_fxt\_options\_init  
    Fxt Support, 154  
starpu\_fxt\_start\_profiling  
    Fxt Support, 154  
starpu\_fxt\_stop\_profiling  
    Fxt Support, 154  
starpu\_get\_env\_number  
    Toolbox, 82  
starpu\_get\_prefetch\_flag  
    Scheduling Policy, 178  
starpu\_get\_version  
    Versioning, 75  
starpu\_hash.h, 200  
starpu\_hash\_crc32c\_be  
    Data Interfaces, 113  
starpu\_hash\_crc32c\_be\_n  
    Data Interfaces, 113  
starpu\_hash\_crc32c\_string  
    Data Interfaces, 113  
starpu\_init  
    Initialization and Termination, 78  
starpu\_insert\_task  
    Insert\_Task, 133  
starpu\_interface\_copy  
    Data Interfaces, 113  
starpu\_malloc  
    Standard Memory Library, 80  
starpu\_malloc\_flags  
    Standard Memory Library, 80  
starpu\_malloc\_on\_node  
    Data Interfaces, 113  
starpu\_malloc\_set\_align  
    Standard Memory Library, 80  
starpu\_matrix\_data\_register  
    Data Interfaces, 109  
starpu\_matrix\_filter\_block  
    Data Partition, 117  
starpu\_matrix\_filter\_block\_shadow  
    Data Partition, 117  
starpu\_matrix\_filter\_vertical\_block  
    Data Partition, 117  
starpu\_matrix\_filter\_vertical\_block\_shadow  
    Data Partition, 117  
starpu\_matrix\_get\_elemsize  
    Data Interfaces, 111  
starpu\_matrix\_get\_local\_id  
    Data Interfaces, 111  
starpu\_matrix\_get\_local\_ptr  
    Data Interfaces, 111

**starpu\_matrix\_get\_nx**  
 Data Interfaces, 111  
**starpu\_matrix\_get\_ny**  
 Data Interfaces, 111  
**starpu\_memory\_get\_available**  
 Standard Memory Library, 80  
**starpu\_mpi.h**, 213  
**starpu\_mpi\_barrier**  
 MPI Support, 158  
**starpu\_mpi\_cache\_flush**  
 MPI Support, 158  
**starpu\_mpi\_cache\_flush\_all\_data**  
 MPI Support, 158  
**starpu\_mpi\_comm\_amounts\_retrieve**  
 MPI Support, 157  
**starpu\_mpi\_gather\_detached**  
 MPI Support, 160  
**starpu\_mpi\_get\_data\_on\_node**  
 MPI Support, 159  
**starpu\_mpi\_get\_data\_on\_node\_detached**  
 MPI Support, 159  
**starpu\_mpi\_init**  
 MPI Support, 157  
**starpu\_mpi\_initialize**  
 MPI Support, 157  
**starpu\_mpi\_initialize\_extended**  
 MPI Support, 157  
**starpu\_mpi\_insert\_task**  
 MPI Support, 159  
**starpu\_mpi\_irecv**  
 MPI Support, 157  
**starpu\_mpi\_irecv\_array\_detached\_unlock\_tag**  
 MPI Support, 158  
**starpu\_mpi\_irecv\_detached**  
 MPI Support, 158  
**starpu\_mpi\_irecv\_detached\_unlock\_tag**  
 MPI Support, 158  
**starpu\_mpi\_isend**  
 MPI Support, 157  
**starpu\_mpi\_isend\_array\_detached\_unlock\_tag**  
 MPI Support, 158  
**starpu\_mpi\_isend\_detached**  
 MPI Support, 157  
**starpu\_mpi\_isend\_detached\_unlock\_tag**  
 MPI Support, 158  
**starpu\_mpi\_recv**  
 MPI Support, 157  
**starpu\_mpi\_redux\_data**  
 MPI Support, 159  
**starpu\_mpi\_scatter\_detached**  
 MPI Support, 160  
**starpu\_mpi\_send**  
 MPI Support, 157  
**starpu\_mpi\_shutdown**  
 MPI Support, 157  
**starpu\_mpi\_test**  
 MPI Support, 158  
**starpu\_mpi\_wait**

**starpu\_multiformat\_data\_register**  
 Multiformat Data Interface, 120  
**starpu\_node\_get\_kind**  
 Workers' Properties, 91  
**starpu\_node\_kind**  
 Workers' Properties, 89  
**starpu\_opencl.h**, 200  
**starpu\_opencl\_allocate\_memory**  
 OpenCL Extensions, 151  
**starpu\_opencl\_collect\_stats**  
 OpenCL Extensions, 151  
**starpu\_opencl\_compile\_opencl\_from\_file**  
 OpenCL Extensions, 150  
**starpu\_opencl\_compile\_opencl\_from\_string**  
 OpenCL Extensions, 150  
**starpu\_opencl\_copy\_async\_sync**  
 OpenCL Extensions, 152  
**starpu\_opencl\_copy\_opencl\_to\_opencl**  
 OpenCL Extensions, 151  
**starpu\_opencl\_copy\_opencl\_to\_ram**  
 OpenCL Extensions, 151  
**starpu\_opencl\_copy\_ram\_to\_opencl**  
 OpenCL Extensions, 151  
**starpu\_opencl\_display\_error**  
 OpenCL Extensions, 151  
**starpu\_opencl\_error\_string**  
 OpenCL Extensions, 151  
**starpu\_opencl\_func\_t**  
 Codelet And Tasks, 129  
**starpu\_opencl\_get\_context**  
 OpenCL Extensions, 149  
**starpu\_opencl\_get\_current\_context**  
 OpenCL Extensions, 149  
**starpu\_opencl\_get\_current\_queue**  
 OpenCL Extensions, 149  
**starpu\_opencl\_get\_device**  
 OpenCL Extensions, 149  
**starpu\_opencl\_get\_queue**  
 OpenCL Extensions, 149  
**starpu\_opencl\_load\_binary\_opencl**  
 OpenCL Extensions, 150  
**starpu\_opencl\_load\_kernel**  
 OpenCL Extensions, 150  
**starpu\_opencl\_load\_opencl\_from\_file**  
 OpenCL Extensions, 150  
**starpu\_opencl\_load\_opencl\_from\_string**  
 OpenCL Extensions, 150  
**starpu\_opencl\_load\_program\_source**  
 OpenCL Extensions, 150  
**starpu\_opencl\_release\_kernel**  
 OpenCL Extensions, 151  
**starpu\_opencl\_report\_error**  
 OpenCL Extensions, 151  
**starpu\_opencl\_set\_kernel\_args**  
 OpenCL Extensions, 149  
**starpu\_opencl\_unload\_opencl**  
 OpenCL Extensions, 150

starpu\_opencl\_worker\_get\_count  
    Workers' Properties, 90  
starpu\_parallel\_task\_barrier\_init  
    Parallel Tasks, 164  
starpu\_perfmodel  
    benchmarking, 137  
    cost\_function, 137  
    cost\_model, 137  
    is\_loaded, 137  
    model\_rwlock, 137  
    per\_arch, 137  
    size\_base, 137  
    symbol, 137  
    type, 137  
starpu\_perfmodel.h, 201  
starpu\_perfmodel\_archtype  
    Performance Model, 139  
starpu\_perfmodel\_debugfilepath  
    Performance Model, 140  
starpu\_perfmodel\_directory  
    Performance Model, 140  
starpu\_perfmodel\_get\_arch\_name  
    Performance Model, 140  
starpu\_perfmodel\_list  
    Performance Model, 140  
starpu\_perfmodel\_load\_symbol  
    Performance Model, 140  
starpu\_perfmodel\_per\_arch  
    cost\_function, 138  
    cost\_model, 138  
    history, 138  
    list, 138  
    regression, 138  
    size\_base, 138  
starpu\_perfmodel\_print  
    Performance Model, 140  
starpu\_perfmodel\_print\_all  
    Performance Model, 140  
starpu\_perfmodel\_type  
    Performance Model, 139  
starpu\_perfmodel\_unload\_model  
    Performance Model, 140  
starpu\_perfmodel\_update\_history  
    Performance Model, 140  
starpu\_prefetch\_task\_input\_on\_node  
    Scheduling Policy, 178  
starpu\_private  
    starpu\_task, 127  
starpu\_profiling.h, 202  
starpu\_profiling\_bus\_helper\_display\_summary  
    Profiling, 144  
starpu\_profiling\_init  
    Profiling, 143  
starpu\_profiling\_set\_id  
    Profiling, 143  
starpu\_profiling\_status\_get  
    Profiling, 143  
starpu\_profiling\_status\_set

    Profiling, 143  
starpu\_profiling\_worker\_get\_info  
    Profiling, 143  
starpu\_profiling\_worker\_helper\_display\_summary  
    Profiling, 144  
starpu\_progression\_hook\_deregister  
    Expert Mode, 165  
starpu\_progression\_hook\_register  
    Expert Mode, 165  
starpu\_pthread\_attr\_destroy  
    Threads, 86  
starpu\_pthread\_attr\_init  
    Threads, 85  
starpu\_pthread\_attr\_setdetachstate  
    Threads, 86  
starpu\_pthread\_cond\_broadcast  
    Threads, 87  
starpu\_pthread\_cond\_destroy  
    Threads, 87  
starpu\_pthread\_cond\_init  
    Threads, 87  
starpu\_pthread\_cond\_signal  
    Threads, 87  
starpu\_pthread\_cond\_timedwait  
    Threads, 87  
starpu\_pthread\_cond\_wait  
    Threads, 87  
starpu\_pthread\_create  
    Threads, 85  
starpu\_pthread\_getspecific  
    Threads, 86  
starpu\_pthread\_join  
    Threads, 85  
starpu\_pthread\_key\_create  
    Threads, 86  
starpu\_pthread\_key\_delete  
    Threads, 86  
starpu\_pthread\_mutex\_destroy  
    Threads, 86  
starpu\_pthread\_mutex\_init  
    Threads, 86  
starpu\_pthread\_mutex\_lock  
    Threads, 86  
starpu\_pthread\_mutex\_trylock  
    Threads, 86  
starpu\_pthread\_mutex\_unlock  
    Threads, 86  
starpu\_pthread\_rwlock\_destroy  
    Threads, 87  
starpu\_pthread\_rwlock\_init  
    Threads, 87  
starpu\_pthread\_rwlock\_rdlock  
    Threads, 87  
starpu\_pthread\_rwlock\_unlock  
    Threads, 87  
starpu\_pthread\_rwlock\_wrlock  
    Threads, 87  
starpu\_pthread\_setspecific

starpu.push.local.task  
     Scheduling Policy, 177  
 starpu.push.task.end  
     Scheduling Policy, 177  
 starpu.rand.h, 203  
 starpu.sched.ctx.h, 203  
 starpu.sched.ctx.add.workers  
     Scheduling Contexts, 172  
 starpu.sched.ctx.call.pushed\_task.cb  
     Scheduling Contexts, 174  
 starpu.sched.ctx.check\_if.hypervisor.exists  
     Scheduling Contexts, 174  
 starpu.sched.ctx.contains.worker  
     Scheduling Contexts, 173  
 starpu.sched.ctx.create  
     Scheduling Contexts, 172  
 starpu.sched.ctx.create.inside.interval  
     Scheduling Contexts, 172  
 starpu.sched.ctx.create.worker.collection  
     Scheduling Contexts, 174  
 starpu.sched.ctx.delete  
     Scheduling Contexts, 173  
 starpu.sched.ctx.delete.worker.collection  
     Scheduling Contexts, 174  
 starpu.sched.ctx.finished.submit  
     Scheduling Contexts, 173  
 starpu.sched.ctx.get.context  
     Scheduling Contexts, 173  
 starpu.sched.ctx.get.max.priority  
     Scheduling Contexts, 174  
 starpu.sched.ctx.get.min.priority  
     Scheduling Contexts, 174  
 starpu.sched.ctx.get.nshared.workers  
     Scheduling Contexts, 173  
 starpu.sched.ctx.get.nworkers  
     Scheduling Contexts, 173  
 starpu.sched.ctx.get.policy.data  
     Scheduling Contexts, 174  
 starpu.sched.ctx.get.worker.collection  
     Scheduling Contexts, 174  
 starpu.sched.ctx.get.workers.list  
     Scheduling Contexts, 173  
 starpu.sched.ctx.hypervisor.h, 204  
 starpu.sched.ctx.notify.hypervisor.exists  
     Scheduling Contexts, 174  
 starpu.sched.ctx.overlapping.ctxs.on.worker  
     Scheduling Contexts, 173  
 starpu.sched.ctx.performance.counters  
     notify.delete.context, 171  
     notify.idle.cycle, 171  
     notify.poped.task, 171  
     notify.pushed.task, 171  
     notify.submitted.job, 171  
 starpu.sched.ctx.register.close.callback  
     Scheduling Contexts, 172  
 starpu.sched.ctx.remove.workers  
     Scheduling Contexts, 173  
 starpu.sched.ctx.set.context  
     Scheduling Contexts, 173  
 starpu.sched.ctx.set.inheritor  
     Scheduling Contexts, 173  
 starpu.sched.ctx.set.max.priority  
     Scheduling Contexts, 174  
 starpu.sched.ctx.set.min.priority  
     Scheduling Contexts, 174  
 starpu.sched.ctx.set.perf.counters  
     Scheduling Contexts, 174  
 starpu.sched.ctx.set.policy.data  
     Scheduling Contexts, 174  
 starpu.sched.ctx.stop.task.submission  
     Scheduling Contexts, 173  
 starpu.sched.ctx.worker.get.id  
     Scheduling Contexts, 173  
 starpu.sched.ctx.worker.shares.tasks.lists  
     Scheduling Policy, 178  
 starpu.sched.get.max.priority  
     Scheduling Policy, 177  
 starpu.sched.get.min.priority  
     Scheduling Policy, 177  
 starpu.sched.get.predefined.policies  
     Scheduling Policy, 177  
 starpu.sched.policy  
     add.workers, 176  
     deinit.sched, 176  
     init.sched, 176  
     policy.description, 176  
     policy.name, 176  
     pop.every.task, 176  
     pop.task, 176  
     post.exec.hook, 176  
     pre.exec.hook, 176  
     push.task, 176  
     push.task.notify, 176  
     remove.workers, 176  
 starpu.sched.set.max.priority  
     Scheduling Policy, 177  
 starpu.sched.set.min.priority  
     Scheduling Policy, 177  
 starpu.scheduler.h, 204  
 starpu.shutdown  
     Initialization and Termination, 79  
 starpu.stdlib.h, 205  
 starpu.tag.declare.deps  
     Explicit Dependencies, 134  
 starpu.tag.declare.deps.array  
     Explicit Dependencies, 134  
 starpu.tag.notify.from.apps  
     Explicit Dependencies, 135  
 starpu.tag.remove  
     Explicit Dependencies, 135  
 starpu.tag.restart  
     Explicit Dependencies, 134  
 starpu.tag.t  
     Explicit Dependencies, 134  
 starpu.tag.wait

Explicit Dependencies, 134  
starpu\_tag\_wait\_array  
    Explicit Dependencies, 134  
starpu\_task  
    buffers, 125  
    bundle, 127  
    callback\_arg, 126  
    callback\_arg\_free, 126  
    callback\_func, 126  
    cl, 125  
    cl\_arg, 125  
    cl\_arg\_free, 125  
    cl\_arg\_size, 125  
    destroy, 127  
    detach, 127  
    dyn\_handles, 125  
    dyn\_interfaces, 125  
    execute\_on\_a\_specific\_worker, 126  
    flops, 127  
    handles, 125  
    hypervisor\_tag, 127  
    interfaces, 125  
    magic, 127  
    mf\_skip, 127  
    next, 127  
    predicted, 127  
    predicted\_transfer, 127  
    prev, 127  
    priority, 126  
    profiling\_info, 127  
    prologue\_callback\_arg, 126  
    prologue\_callback\_arg\_free, 126  
    prologue\_callback\_func, 126  
    regenerate, 127  
    sched\_ctx, 127  
    scheduled, 127  
    sequential\_consistency, 126  
    starpu\_private, 127  
    status, 127  
    synchronous, 126  
    tag\_id, 126  
    use\_tag, 126  
    workerid, 127  
starpu\_task.h, 206  
starpu\_task\_build  
    Insert\_Task, 133  
starpu\_task\_bundle.h, 207  
starpu\_task\_bundle\_close  
    Task Bundles, 161  
starpu\_task\_bundle\_create  
    Task Bundles, 160  
starpu\_task\_bundle\_expected\_data\_transfer\_time  
    Task Bundles, 161  
starpu\_task\_bundle\_expected\_length  
    Task Bundles, 161  
starpu\_task\_bundle\_expected\_power  
    Task Bundles, 161  
starpu\_task\_bundle\_insert  
    Task Bundles, 160  
starpu\_task\_bundle\_remove  
    Task Bundles, 161  
starpu\_task\_bundle\_t  
    Task Bundles, 160  
starpu\_task\_clean  
    Codelet And Tasks, 130  
starpu\_task\_create  
    Codelet And Tasks, 130  
starpu\_task\_declare\_deps\_array  
    Explicit Dependencies, 134  
starpu\_task\_destroy  
    Codelet And Tasks, 130  
starpu\_task\_dup  
    Codelet And Tasks, 130  
starpu\_task\_expected\_conversion\_time  
    Scheduling Policy, 178  
starpu\_task\_expected\_data\_transfer\_time  
    Scheduling Policy, 178  
starpu\_task\_expected\_length  
    Scheduling Policy, 177  
starpu\_task\_expected\_power  
    Scheduling Policy, 178  
starpu\_task\_footprint  
    Scheduling Policy, 177  
starpu\_task\_get\_current  
    Codelet And Tasks, 131  
starpu\_task\_get\_implementation  
    Codelet And Tasks, 131  
starpu\_task\_init  
    Codelet And Tasks, 130  
starpu\_task\_list.h, 207  
starpu\_task\_list\_back  
    Task Lists, 163  
starpu\_task\_list\_begin  
    Task Lists, 163  
starpu\_task\_list\_empty  
    Task Lists, 163  
starpu\_task\_list\_erase  
    Task Lists, 163  
starpu\_task\_list\_front  
    Task Lists, 163  
starpu\_task\_list\_init  
    Task Lists, 163  
starpu\_task\_list\_next  
    Task Lists, 163  
starpu\_task\_list\_pop\_back  
    Task Lists, 163  
starpu\_task\_list\_pop\_front  
    Task Lists, 163  
starpu\_task\_list\_push\_back  
    Task Lists, 163  
starpu\_task\_list\_push\_front  
    Task Lists, 163  
starpu\_task\_nready  
    Codelet And Tasks, 131  
starpu\_task\_nsubmitted  
    Codelet And Tasks, 131

**starpu\_task\_set\_implementation**  
 Codelet And Tasks, 131  
**starpu\_task\_status**  
 Codelet And Tasks, 129  
**starpu\_task\_submit**  
 Codelet And Tasks, 130  
**starpu\_task\_submit\_to\_ctx**  
 Codelet And Tasks, 130  
**starpu\_task\_util.h**, 208  
**starpu\_task\_wait**  
 Codelet And Tasks, 130  
**starpu\_task\_wait\_for\_all**  
 Codelet And Tasks, 131  
**starpu\_task\_wait\_for\_all\_in\_ctx**  
 Codelet And Tasks, 131  
**starpu\_task\_wait\_for\_no\_ready**  
 Codelet And Tasks, 131  
**starpu\_thread.h**, 208  
**starpu\_thread\_util.h**, 209  
**starpu\_timing\_now**  
 Scheduling Policy, 177  
**starpu\_timing\_timespec\_delay\_us**  
 Profiling, 144  
**starpu\_timing\_timespec\_to\_us**  
 Profiling, 144  
**starpu\_top.h**, 210  
**starpu\_top\_add\_data\_boolean**  
 StarPU-Top Interface, 168  
**starpu\_top\_add\_data\_float**  
 StarPU-Top Interface, 168  
**starpu\_top\_add\_data\_integer**  
 StarPU-Top Interface, 168  
**starpu\_top\_data\_type**  
 StarPU-Top Interface, 168  
**starpu\_top\_debug\_lock**  
 StarPU-Top Interface, 170  
**starpu\_top\_debug\_log**  
 StarPU-Top Interface, 169  
**starpu\_top\_init\_and\_wait**  
 StarPU-Top Interface, 169  
**starpu\_top\_message\_type**  
 StarPU-Top Interface, 168  
**starpu\_top\_param**  
 callback, 167  
 double\_max\_value, 168  
 double\_min\_value, 167  
 enum\_values, 167  
 id, 167  
 int\_max\_value, 167  
 int\_min\_value, 167  
 name, 167  
 nb\_values, 167  
 next, 168  
 type, 167  
 value, 167  
**starpu\_top\_param\_type**  
 StarPU-Top Interface, 168  
**starpu\_top\_register\_parameter\_boolean**  
 StarPU-Top Interface, 169  
**starpu\_top\_register\_parameter\_enum**  
 StarPU-Top Interface, 169  
**starpu\_top\_register\_parameter\_float**  
 StarPU-Top Interface, 169  
**starpu\_top\_register\_parameter\_integer**  
 StarPU-Top Interface, 169  
**starpu\_top\_task\_prevision**  
 StarPU-Top Interface, 169  
**starpu\_top\_update\_data\_boolean**  
 StarPU-Top Interface, 169  
**starpu\_top\_update\_data\_float**  
 StarPU-Top Interface, 169  
**starpu\_top\_update\_data\_integer**  
 StarPU-Top Interface, 169  
**starpu\_top\_update\_parameter**  
 StarPU-Top Interface, 169  
**starpu\_topology\_print**  
 Initialization and Termination, 79  
**starpu\_transfer\_bandwidth**  
 Performance Model, 141  
**starpu\_transfer\_latency**  
 Performance Model, 141  
**starpu\_transfer\_predict**  
 Performance Model, 141  
**starpu\_util.h**, 211  
**starpu\_variable\_data\_register**  
 Data Interfaces, 109  
**starpu\_variable\_get\_elemsize**  
 Data Interfaces, 110  
**starpu\_variable\_get\_local\_ptr**  
 Data Interfaces, 111  
**starpu\_vector\_data\_register**  
 Data Interfaces, 109  
**starpu\_vector\_filter\_block**  
 Data Partition, 116  
**starpu\_vector\_filter\_block\_shadow**  
 Data Partition, 116  
**starpu\_vector\_filter\_divide\_in\_2**  
 Data Partition, 117  
**starpu\_vector\_filter\_list**  
 Data Partition, 116  
**starpu\_vector\_get\_elemsize**  
 Data Interfaces, 111  
**starpu\_vector\_get\_local\_ptr**  
 Data Interfaces, 111  
**starpu\_vector\_get\_nx**  
 Data Interfaces, 111  
**starpu\_void\_data\_register**  
 Data Interfaces, 109  
**starpu\_wake\_all\_blocked\_workers**  
 Expert Mode, 165  
**starpu\_worker.h**, 212  
**starpu\_worker\_archtype**  
 Workers' Properties, 89  
**starpu\_worker\_can\_execute\_task**  
 Scheduling Policy, 177  
**starpu\_worker\_collection**

add, 89  
 deinit, 89  
 get\_next, 89  
 has\_next, 89  
 init, 89  
 init\_iterator, 89  
 nworkers, 89  
 remove, 89  
 type, 89  
 workerids, 89  
**starpu\_worker\_collection\_type**  
     Workers' Properties, 90  
**starpu\_worker\_get\_by\_devid**  
     Workers' Properties, 91  
**starpu\_worker\_get\_by\_type**  
     Workers' Properties, 90  
**starpu\_worker\_get\_count**  
     Workers' Properties, 90  
**starpu\_worker\_get\_count\_by\_type**  
     Workers' Properties, 90  
**starpu\_worker\_get\_devid**  
     Workers' Properties, 91  
**starpu\_worker\_get\_id**  
     Workers' Properties, 90  
**starpu\_worker\_get\_ids\_by\_type**  
     Workers' Properties, 90  
**starpu\_worker\_get\_memory\_node**  
     Workers' Properties, 91  
**starpu\_worker\_get\_name**  
     Workers' Properties, 91  
**starpu\_worker\_get\_perf\_archtype**  
     Performance Model, 140  
**starpu\_worker\_get\_relative\_speedup**  
     Scheduling Policy, 178  
**starpu\_worker\_get\_sched\_condition**  
     Scheduling Policy, 177  
**starpu\_worker\_get\_type**  
     Workers' Properties, 91  
**starpufft\_cleanup**  
     FFT Support, 155  
**starpufft\_destroy\_plan**  
     FFT Support, 155  
**starpufft\_execute**  
     FFT Support, 155  
**starpufft\_execute\_handle**  
     FFT Support, 155  
**starpufft\_free**  
     FFT Support, 154  
**starpufft\_malloc**  
     FFT Support, 154  
**starpufft\_plan\_dft\_1d**  
     FFT Support, 154  
**starpufft\_plan\_dft\_2d**  
     FFT Support, 155  
**starpufft\_start**  
     FFT Support, 155  
**starpufft\_start\_handle**  
     FFT Support, 155  
**status**  
     starpu\_task, 127  
**symbol**  
     starpu\_perfmodel, 137  
**synchronous**  
     starpu\_task, 126  
**TOP\_TYPE\_CONTINUE**  
     StarPU-Top Interface, 168  
**TOP\_TYPE\_DEBUG**  
     StarPU-Top Interface, 168  
**TOP\_TYPE\_DISABLE**  
     StarPU-Top Interface, 168  
**TOP\_TYPE\_ENABLE**  
     StarPU-Top Interface, 168  
**TOP\_TYPE\_GO**  
     StarPU-Top Interface, 168  
**TOP\_TYPE\_SET**  
     StarPU-Top Interface, 168  
**TOP\_TYPE\_UNKNOW**  
     StarPU-Top Interface, 168  
**tag\_id**  
     starpu\_task, 126  
**Task Bundles, 160**  
     **starpu\_task\_bundle\_close**, 161  
     **starpu\_task\_bundle\_create**, 160  
     **starpu\_task\_bundle\_expected\_data\_transfer\_time**,  
         161  
     **starpu\_task\_bundle\_expected\_length**, 161  
     **starpu\_task\_bundle\_expected\_power**, 161  
     **starpu\_task\_bundle\_insert**, 160  
     **starpu\_task\_bundle\_remove**, 161  
     **starpu\_task\_bundle\_t**, 160  
**Task Lists, 161**  
     **starpu\_task\_list\_back**, 163  
     **starpu\_task\_list\_begin**, 163  
     **starpu\_task\_list\_empty**, 163  
     **starpu\_task\_list\_erase**, 163  
     **starpu\_task\_list\_front**, 163  
     **starpu\_task\_list\_init**, 163  
     **starpu\_task\_list\_next**, 163  
     **starpu\_task\_list\_pop\_back**, 163  
     **starpu\_task\_list\_pop\_front**, 163  
     **starpu\_task\_list\_push\_back**, 163  
     **starpu\_task\_list\_push\_front**, 163  
**Theoretical Lower Bound on Execution Time, 144**  
     **starpu\_bound\_compute**, 144  
     **starpu\_bound\_print**, 145  
     **starpu\_bound\_print\_dot**, 144  
     **starpu\_bound\_print\_lp**, 145  
     **starpu\_bound\_print\_mps**, 145  
     **starpu\_bound\_start**, 144  
     **starpu\_bound\_stop**, 144  
**Threads, 82**  
     **starpu\_pthread\_attr\_destroy**, 86  
     **starpu\_pthread\_attr\_init**, 85  
     **starpu\_pthread\_attr\_setdetachstate**, 86  
     **starpu\_pthread\_cond\_broadcast**, 87  
     **starpu\_pthread\_cond\_destroy**, 87

starpu\_pthread\_cond\_init, 87  
 starpu\_pthread\_cond\_signal, 87  
 starpu\_pthread\_cond\_timedwait, 87  
 starpu\_pthread\_cond\_wait, 87  
 starpu\_pthread\_create, 85  
 starpu\_pthread\_getspecific, 86  
 starpu\_pthread\_join, 85  
 starpu\_pthread\_key\_create, 86  
 starpu\_pthread\_key\_delete, 86  
 starpu\_pthread\_mutex\_destroy, 86  
 starpu\_pthread\_mutex\_init, 86  
 starpu\_pthread\_mutex\_lock, 86  
 starpu\_pthread\_mutex\_trylock, 86  
 starpu\_pthread\_mutex\_unlock, 86  
 starpu\_pthread\_rwlock\_destroy, 87  
 starpu\_pthread\_rwlock\_init, 87  
 starpu\_pthread\_rwlock\_rdlock, 87  
 starpu\_pthread\_rwlock\_unlock, 87  
 starpu\_pthread\_rwlock\_wrlock, 87  
 starpu\_pthread\_setspecific, 86  
**Toolbox**, 80  
 STARPU\_ABORT, 82  
 STARPU\_ABORT\_MSG, 82  
 STARPU\_ASSERT, 82  
 STARPU\_ASSERT\_MSG, 82  
 STARPU\_GNUC\_PREREQ, 81  
 STARPU\_LIKELY, 81  
 STARPU\_MAX, 82  
 STARPU\_MIN, 82  
 STARPU\_POISON\_PTR, 82  
 STARPU\_RMB, 82  
 STARPU\_UNLIKELY, 81  
 STARPU\_WMB, 82  
 starpu\_get\_env\_number, 82  
**type**  
 starpu\_codelet, 122  
 starpu\_perfmodel, 137  
 starpu\_top\_param, 167  
 starpu\_worker\_collection, 89  
**types\_of\_workers**, 217  
  
**unpack\_data**  
 starpu\_data\_interface\_ops, 100  
**use\_tag**  
 starpu\_task, 126  
  
**value**  
 starpu\_top\_param, 167  
**Versioning**, 75  
 STARPU\_MAJOR\_VERSION, 75  
 STARPU\_MINOR\_VERSION, 75  
 starpu\_get\_version, 75  
  
**where**  
 starpu\_codelet, 122  
**workerid**  
 starpu\_task, 127  
**workerids**  
 starpu\_worker\_collection, 89  
  
**Workers' Properties**, 88  
 STARPU\_ANY\_WORKER, 90  
 STARPU\_CPU\_RAM, 89  
 STARPU\_CPU\_WORKER, 90  
 STARPU\_CUDA\_RAM, 89  
 STARPU\_CUDA\_WORKER, 90  
 STARPU\_OPENCL\_RAM, 89  
 STARPU\_OPENCL\_WORKER, 90  
 STARPU\_UNUSED, 89  
 STARPU\_WORKER\_LIST, 90  
 STARPU\_NMAXWORKERS, 89  
 starpu\_cpu\_worker\_get\_count, 90  
 starpu\_cuda\_worker\_get\_count, 90  
 starpu\_node\_get\_kind, 91  
 starpu\_node\_kind, 89  
 starpu\_opencl\_worker\_get\_count, 90  
 starpu\_worker\_archtype, 89  
 starpu\_worker\_collection\_type, 90  
 starpu\_worker\_get\_by\_devid, 91  
 starpu\_worker\_get\_by\_type, 90  
 starpu\_worker\_get\_count, 90  
 starpu\_worker\_get\_count\_by\_type, 90  
 starpu\_worker\_get\_devid, 91  
 starpu\_worker\_get\_id, 90  
 starpu\_worker\_get\_ids\_by\_type, 90  
 starpu\_worker\_get\_memory\_node, 91  
 starpu\_worker\_get\_name, 91  
 starpu\_worker\_get\_type, 91

## **Part III**

# **Appendix**



# Chapter 21

## Full Source Code for the 'Scaling a Vector' Example

### 21.1 Main Application

```
/*
 * This example demonstrates how to use StarPU to scale an array by a factor.
 * It shows how to manipulate data with StarPU's data management library.
 * 1- how to declare a piece of data to StarPU (starpu_vector_data_register)
 * 2- how to describe which data are accessed by a task (task->handles[0])
 * 3- how a kernel can manipulate the data (buffers[0].vector.ptr)
 */
#include <starpu.h>

#define NX 2048

extern void scal_cpu_func(void *buffers[], void *_args);
extern void scal_sse_func(void *buffers[], void *_args);
extern void scal_cuda_func(void *buffers[], void *_args);
extern void scal_opencl_func(void *buffers[], void *_args);

static struct starpu_codelet cl = {
    .where = STARPU_CPU | STARPU_CUDA | STARPU_OPENCL,
    /* CPU implementation of the codelet */
    .cpu_funcs = { scal_cpu_func, scal_sse_func, NULL },
#ifndef STARPU_USE_CUDA
    /* CUDA implementation of the codelet */
    .cuda_funcs = { scal_cuda_func, NULL },
#endif
#ifndef STARPU_USE_OPENCL
    /* OpenCL implementation of the codelet */
    .opencl_funcs = { scal_opencl_func, NULL },
#endif
    .nbuffers = 1,
    .modes = { STARPU_RW }
};

#ifndef STARPU_USE_OPENCL
struct starpu_opencl_program programs;
#endif

int main(int argc, char **argv)
{
    /* We consider a vector of float that is initialized just as any of C
     * data */
    float vector[NX];
    unsigned i;
    for (i = 0; i < NX; i++)
        vector[i] = 1.0f;

    fprintf(stderr, "BEFORE: First element was %f\n", vector[0]);

    /* Initialize StarPU with default configuration */
    starpu_init(NULL);

#ifndef STARPU_USE_OPENCL
    starpu_opencl_load_opencl_from_file(
        "examples/basic-examples/vector-scal-opencl/kernel.cl", &programs, NULL);
#endif

    /* Tell StarPU to associate the "vector" vector with the "vector_handle"
     * identifier. When a task needs to access a piece of data, it should
     * refer to the handle that is associated to it.
```

```

 * In the case of the "vector" data interface:
 * - the first argument of the registration method is a pointer to the
 *   handle that should describe the data
 * - the second argument is the memory node where the data (ie. "vector")
 *   resides initially: 0 stands for an address in main memory, as
 *   opposed to an address on a GPU for instance.
 * - the third argument is the address of the vector in RAM
 * - the fourth argument is the number of elements in the vector
 * - the fifth argument is the size of each element.
 */
starpu_data_handle_t vector_handle;
starpu_vector_data_register(&vector_handle, 0, (uintptr_t)vector,
                           NX, sizeof(vector[0]));

float factor = 3.14;

/* create a synchronous task: any call to starpu_task_submit will block
 * until it is terminated */
struct starpu_task *task = starpu_task_create();
task->synchronous = 1;

task->cl = &cl;

/* the codelet manipulates one buffer in RW mode */
task->handles[0] = vector_handle;

/* an argument is passed to the codelet, beware that this is a
 * READ-ONLY buffer and that the codelet may be given a pointer to a
 * COPY of the argument */
task->cl.arg = &factor;
task->cl.arg.size = sizeof(factor);

/* execute the task on any eligible computational ressource */
starpu_task_submit(task);

/* StarPU does not need to manipulate the array anymore so we can stop
 * monitoring it */
starpu_data_unregister(vector_handle);

#ifndef STARPU_USE_OPENCL
    starpu_opencl_unload_opencl(&programs);
#endif

/* terminate StarPU, no task can be submitted after */
starpu_shutdown();

fprintf(stderr, "AFTER First element is %f\n", vector[0]);

return 0;
}

```

## 21.2 CPU Kernel

```

#include <starpu.h>
#include <xmmmintrin.h>

/* This kernel takes a buffer and scales it by a constant factor */
void scal_cpu_func(void *buffers[], void *cl_arg)
{
    unsigned i;
    float *factor = cl_arg;

    /*
     * The "buffers" array matches the task->handles array: for instance
     * task->handles[0] is a handle that corresponds to a data with
     * vector "interface", so that the first entry of the array in the
     * codelet is a pointer to a structure describing such a vector (ie.
     * struct starpu_vector_interface *). Here, we therefore manipulate
     * the buffers[0] element as a vector: nx gives the number of elements
     * in the array, ptr gives the location of the array (that was possibly
     * migrated/replicated), and elemsize gives the size of each elements.
     */
    struct starpu_vector_interface *vector = buffers[0];

    /* length of the vector */
    unsigned n = STARPU_VECTOR_GET_NX(vector);

    /* get a pointer to the local copy of the vector: note that we have to
     * cast it in (float *) since a vector could contain any type of
     * elements so that the .ptr field is actually a uintptr_t */
    float *val = (float *)STARPU_VECTOR_GET_PTR(vector);

    /* scale the vector */
    for (i = 0; i < n; i++)

```

```

        val[i] *= *factor;
    }

void scal_sse_func(void *buffers[], void *cl_arg)
{
    float *vector = (float *) STARPU_VECTOR_GET_PTR(buffers[0]);
    unsigned int n = STARPU_VECTOR_GET_NX(buffers[0]);
    unsigned int n_iterations = n/4;

    __m128 *VECTOR = (__m128*) vector;
    __m128 FACTOR = STARPU_ATTRIBUTE_ALIGNED(16);
    float factor = *(float *) cl_arg;
    FACTOR = _mm.set1_ps(factor);

    unsigned int i;
    for (i = 0; i < n_iterations; i++)
        VECTOR[i] = _mm.mul_ps(FACTOR, VECTOR[i]);

    unsigned int remainder = n%4;
    if (remainder != 0)
    {
        unsigned int start = 4 * n_iterations;
        for (i = start; i < start+remainder; ++i)
        {
            vector[i] = factor * vector[i];
        }
    }
}

```

## 21.3 CUDA Kernel

```

#include <starpu.h>

static __global__ void vector_mult_cuda(unsigned n, float *val,
                                         float factor)
{
    unsigned i = blockIdx.x*blockDim.x + threadIdx.x;
    if (i < n)
        val[i] *= factor;
}

extern "C" void scal_cuda_func(void *buffers[], void *_args)
{
    float *factor = (float *)_args;

    /* length of the vector */
    unsigned n = STARPU_VECTOR_GET_NX(buffers[0]);
    /* local copy of the vector pointer */
    float *val = (float *)STARPU_VECTOR_GET_PTR(buffers[0]);
    unsigned threads_per_block = 64;
    unsigned nblocks = (n + threads_per_block-1) / threads_per_block;

    vector_mult_cuda<<<nblocks,threads_per_block, 0, starpu.cuda.get_local_stream()>>>
        (n, val, *factor);

    cudaStreamSynchronize(starpu.cuda.get_local_stream());
}

```

## 21.4 OpenCL Kernel

### 21.4.1 Invoking the Kernel

```

#include <starpu.h>

extern struct starpu_opencl_program programs;

void scal_opencl_func(void *buffers[], void *_args)
{
    float *factor = _args;
    int id, devid, err;
    cl_kernel kernel; /* OpenCL specific code */
    cl_command_queue queue; /* OpenCL specific code */
    cl_event event; /* OpenCL specific code */

    /* length of the vector */
    unsigned n = STARPU_VECTOR_GET_NX(buffers[0]);
    /* OpenCL copy of the vector pointer */
    cl_mem val = (cl_mem)STARPU_VECTOR_GET_DEV_HANDLE(buffers[0]);

    { /* OpenCL specific code */
        id = starpu_worker_get_id();
    }
}

```

```

devid = starpu_worker_get_devid(id);

err = starpu_opencl_load_kernel(&kernel,
                               &queue,
                               &programs,
                               "vector_mult_opencl", /* Name of the codelet */
                               devid);
if (err != CL_SUCCESS) STARPU_OPENCL_REPORT_ERROR(err);

err = clSetKernelArg(kernel, 0, sizeof(n), &n);
err |= clSetKernelArg(kernel, 1, sizeof(val), &val);
err |= clSetKernelArg(kernel, 2, sizeof(*factor), factor);
if (err) STARPU_OPENCL_REPORT_ERROR(err);
}

{ /* OpenCL specific code */
size_t global=n;
size_t local;
size_t s;
cl_device_id device;

starpu_opencl_get_device(devid, &device);
err = clGetKernelWorkGroupInfo (kernel, device, CL_KERNEL_WORK_GROUP_SIZE,
                               sizeof(local), &local, &s);
if (err != CL_SUCCESS) STARPU_OPENCL_REPORT_ERROR(err);
if (local > global) local=global;

err = clEnqueueNDRangeKernel(queue, kernel, 1, NULL, &global, &local, 0,
                             NULL, &event);
if (err != CL_SUCCESS) STARPU_OPENCL_REPORT_ERROR(err);
}

{ /* OpenCL specific code */
clFinish(queue);
starpu_opencl_collect_stats(event);
clReleaseEvent(event);

starpu_opencl_release_kernel(kernel);
}
}

```

### 21.4.2 Source of the Kernel

```

__kernel void vector_mult_opencl(int nx, __global float* val, float factor)
{
    const int i = get_global_id(0);
    if (i < nx) {
        val[i] *= factor;
    }
}

```

## Chapter 22

# GNU Free Documentation License

Version 1.3, 3 November 2008

### Copyright

2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc. <http://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

### 1. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

### 2. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a worldwide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

### 3. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

### 4. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document’s license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a

computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

## 5. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- (a) Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- (b) List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- (c) State on the Title page the name of the publisher of the Modified Version, as the publisher.
- (d) Preserve all the copyright notices of the Document.
- (e) Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- (f) Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- (g) Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- (h) Include an unaltered copy of this License.
- (i) Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- (j) Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- (k) For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- (l) Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- (m) Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- (n) Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- (o) Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements”, provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

## 6. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements.”

## 7. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

## 8. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

## 9. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

## 10. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

## 11. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

## 12. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

## 22.1 ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (C) year your name. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

with the Invariant Sections being *list their titles*, with the Front-Cover Texts being *list*, and with the Back-Cover Texts being *list*.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.