

Classical CPU tools training at CERN

Performance Libraries Intel® oneAPI Threading Building Blocks Changes from previous versions

Aleksei Fedotov, Software Development Engineer at DSE PRE, Intel



intel®

Agenda

- Revamp of the Intel® Threading Building Blocks (TBB)
- Motivation of the Revamp
- What is Intel® oneAPI Threading Building Blocks (oneTBB)?
- How it Differs from Intel® Threading Building Blocks (TBB)
- Migration from TBB to oneTBB
- Conclusion

Revamp of Intel® Threading Building Blocks (TBB)

- The classic Intel® Threading Building Blocks (TBB) library had been revamped into Intel® oneAPI Threading Building Blocks (oneTBB)
- The latest version of the TBB library before the revamp is TBB 2020.3
- oneTBB is not binary compatible with TBB
- Among other things, name of the library had been changed to reflect its inclusion into Intel® oneAPI Toolkits
 - Its official name now is Intel® oneTBB Threading Building Blocks (oneTBB)
 - It can also be downloaded standalone

Motivation of the Library Revamp

- Compliance with the latest C++ standards
- Improve the usability and simplicity of the library
- Deprecation and eventual removal of legacy TBB features

Intel[®] oneAPI Threading Building Blocks (oneTBB)

- A portable C++ library for parallel programming
- Supports multiple platforms (some of them by community), e.g., x86, ARM, MIPS and others
- Distributed:
 - via open-source project at GitHub (<https://github.com/oneapi-src/oneTBB>)
 - as part of Intel oneAPI Base Toolkit
 - via package managers like apt and yum
 - via native channels like nuget.org and anaconda.org
- Qualified in accordance with ISO 26262 for Functional Safety and distributed to Mobileye

GitHub - oneapi-src/oneTBB: one x +

github.com/oneapi-src/oneTBB

Python 0.9% SWIG 0.1%
Starlark 0.1%

README.md

oneAPI Threading Building Blocks

license: Apache 2.0 oneTBB CI passing

oneTBB is a flexible C++ library that simplifies the work of adding parallelism to complex applications, even if you are not a threading expert.

The library lets you easily write parallel programs that take full advantage of the multi-core performance. Such programs are portable, composable and have a future-proof scalability. oneTBB provides you with functions, interfaces, and classes to parallelize and scale the code. All you have to do is to use the templates.

The library differs from typical threading packages in the following ways:

- oneTBB enables you to specify logical parallelism instead of threads.
- oneTBB targets threading for performance.
- oneTBB is compatible with other threading packages.
- oneTBB emphasizes scalable, data parallel programming.
- oneTBB relies on generic programming.

Refer to oneTBB [examples](#) and [samples](#) to see how you can use the library.

oneTBB is a part of [oneAPI](#). The current branch implements version 1.1 of oneAPI Specification.

Release Information

Here are [Release Notes](#) and [System Requirements](#).

Documentation

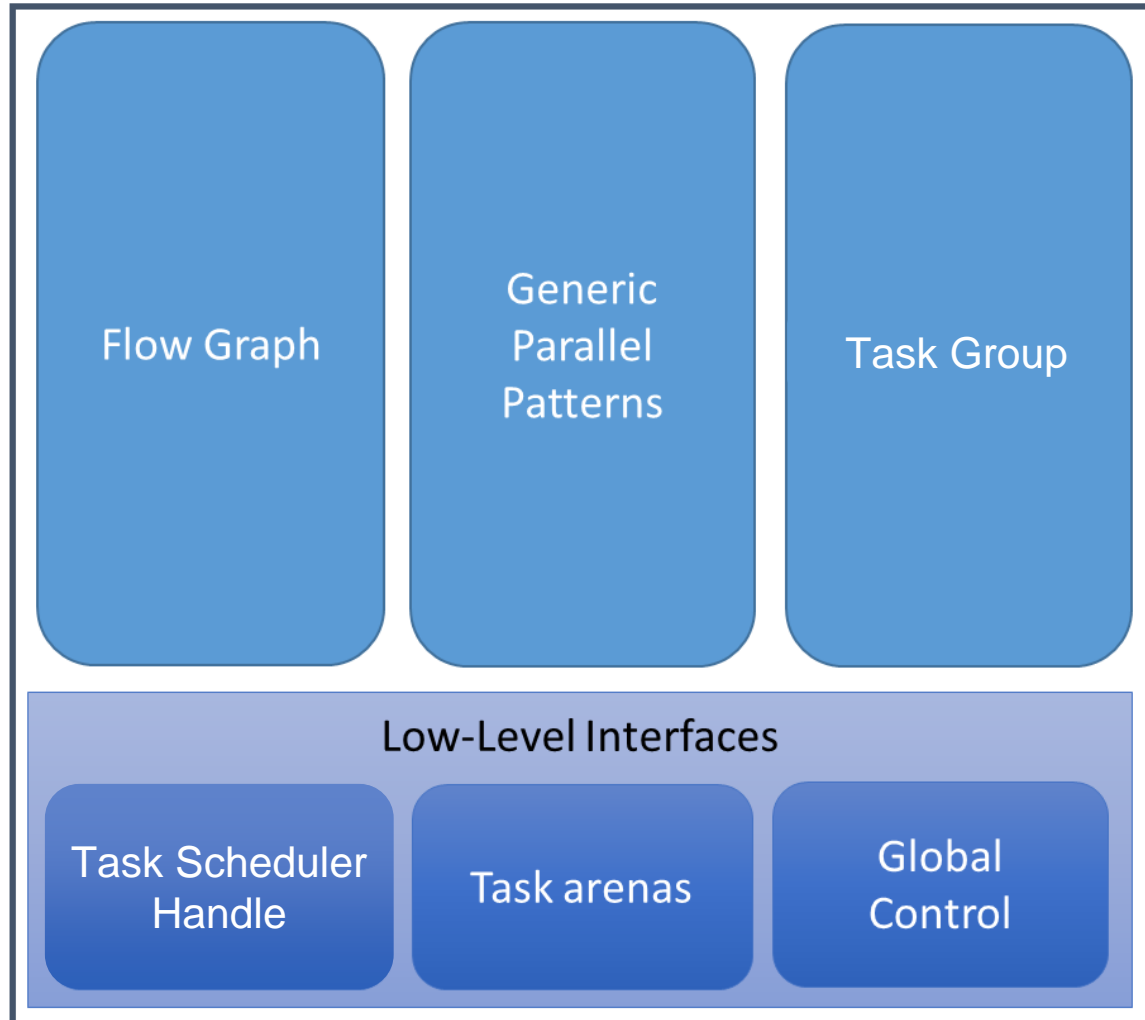
- [oneTBB Specification](#)
- [oneTBB Developer Guide and Reference](#)
- [Migrating from TBB to oneTBB](#)
- [README for the CMake build system](#)
- [Basic support for the Bazel build system](#)
- [oneTBB Discussions](#)

Installation

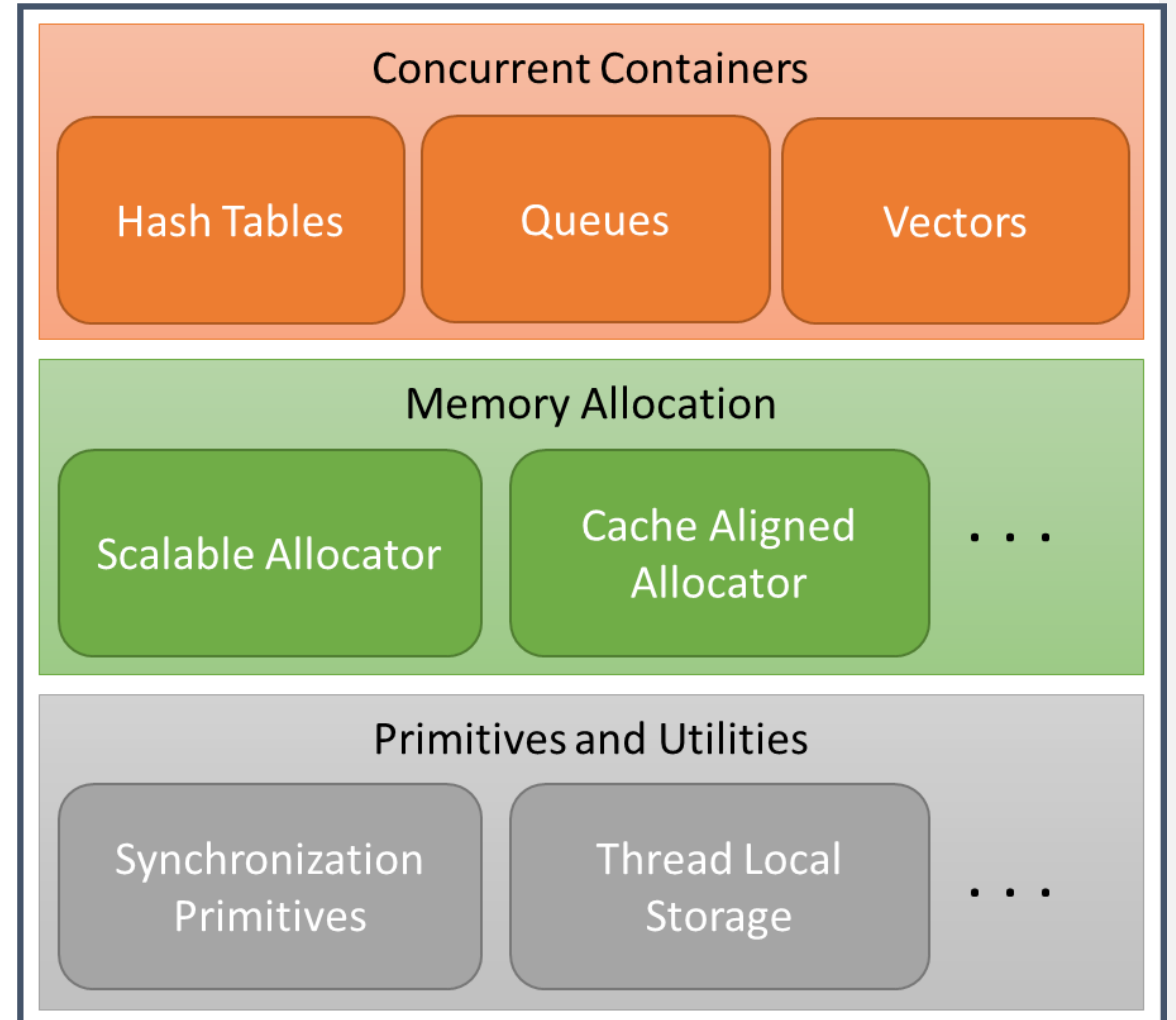
See [Installation from Sources](#) to learn how to install oneTBB.

Intel® oneAPI Threading Building Blocks (oneTBB)

Parallel Execution Interfaces



Interfaces Independent of Execution Model



Functionality Removed in oneTBB (Pre C++11 Compatibility API)

Deprecated/removed TBB functionality	Replacement
<code>tbb::atomic</code>	<code>std::atomic</code>
<code>tbb::flow::tuple</code> (incl. helper classes)	<code>std::tuple</code>
<code>tbb::mutex</code> , <code>tbb::critical_section</code> (incl. <code>tbb::improper_lock</code>)	<code>std::mutex</code>
<code>tbb::recursive_mutex</code>	<code>std::recursive_mutex</code>
<code>tbb::hash</code> (incl. <code>tbb::hasher</code>)	<code>std::hash</code>
<code>tbb::tbb_thread</code> / <code>std::thread</code> / <code>std::this_thread</code>	<code>std::thread</code> / <code>std::this_thread</code> with possible minimal changes related to <code>std::chrono</code>
<code>std::lock_guard</code> / <code>std::unique_lock</code> (incl. helper classes)	Minimal changes related to <code>std::chrono</code> might be required
<code>std::condition_variable</code> (incl. <code>std::cv_status</code> , <code>std::timeout</code> , <code>std::no_timeout</code>)	Minimal changes related to <code>std::chrono</code> might be required
<code>tbb::aligned_space</code>	<code>std::aligned_storage</code>
<code>tbb::tbb_exception</code> / <code>tbb::captured_exception</code> / <code>tbb::movable_exception</code>	No more needed due to TBB exact exception propagation

Compatibility with Microsoft* Parallel Patterns Library (PPL)

Deprecated/removed TBB functionality	Replacement
<code>Concurrency::critical_section</code>	<code>std::mutex</code>
<code>Concurrency::reader_writer_lock</code> (incl. <code>Concurrency::improper_lock</code>)	<code>std::shared_mutex</code>
<code>Concurrency::parallel_invoke</code>	<code>tbb::parallel_invoke</code>
<code>Concurrency::parallel_for</code> (first, last, f)	<code>tbb::parallel_for</code> (first, last, f)
<code>Concurrency::parallel_for_each</code>	<code>tbb::parallel_for_each</code>
<code>Concurrency::task_group</code> (incl. helper classes)	<code>tbb::task_group</code>
<code>Concurrency::structured_task_group</code> (incl. helper classes)	<code>tbb::task_group</code>

Other Functionality

Deprecated/removed TBB functionality	Replacement
Task API (<code>tbb::task</code> , <code>tbb::empty_task</code> , <code>tbb::task_list</code> and related functions)	No direct replacement, the majority use cases can be covered with <code>tbb::task_group</code> , <code>tbb::flow::graph</code> . Task priorities can be covered with Flow graph node priorities and static arena-level priorities.
<code>tbb::task_scheduler_init</code>	<code>tbb::task_arena</code> , <code>tbb::global_control</code> , <code>task_scheduler_handle</code>
<code>tbb::pipeline</code> (incl. <code>tbb::filter</code> , <code>tbb::thread_bound_filter</code>)	<code>tbb::parallel_pipeline</code> , <code>tbb::flow_async_node</code> , resumable tasks
<code>tbb::flow::sender/receiver/continue_receiver</code>	Remain as unspecified base types for flow graph classes
Allocator template parameter for the flow graph nodes	No replacement is planned
<code>tbb::flow::async_msg</code> , <code>tbb::flow::streaming_node</code> , <code>tbb::flow::opencl_node</code>	No replacement is planned. To interact with asynchronous/heterogeneous activity use <code>tbb::flow::async_node</code> or resumable tasks
(preview) <code>tbb::serial::parallel_for</code>	Limit the number of threads to 1 with <code>task_arena</code> or <code>global_control</code>
(preview) <code>runtime_loader</code> (aka <code>tbbproxy</code> library)	No replacement is planned
<code>tbb::structured_task_group</code> (incl. helper classes)	<code>tbb::task_group</code>
<code>tbb::parallel_do</code>	<code>tbb::parallel_for_each</code>
<code>tbb::flow::source node</code>	<code>tbb::flow::input_node</code>
<code>tbb::reader_writer_lock</code>	<code>std::shared_mutex</code>

Migration to oneTBB

- Dedicated page that helps in migrating
 - https://oneapi-src.github.io/oneTBB/main/tbb_userguide/Migration_Guide.html
 - Regularly populated with new and not yet covered use cases
 - Now includes explanations on how to:
 - Migrate from `tbb::task_scheduler_init`
 - Migrate from low-level tasking API
 - Mix to runtimes
- PDF document describing classic interfaces to replace with (<https://www.intel.com/content/www/us/en/developer/articles/technical/tbb-revamp.html>)

Example 1: Spawning of Individual Tasks Using classic TBB Task API

```
#include <tbb/task.h>

int main() {
    // Assuming RootTask, ChildTask1, and ChildTask2 are defined.
    RootTask& root = *new(tbb::task::allocate_root()) RootTask{};
    ChildTask1& child1 = *new(root.allocate_child()) ChildTask1{/*params*/};
    ChildTask2& child2 = *new(root.allocate_child()) ChildTask2{/*params*/};
    root.set_ref_count(3);
    tbb::task::spawn(child1);
    tbb::task::spawn(child2);
    root.wait_for_all();
}
```

Example 1: Spawning of Individual Tasks Using classic TBB Task API

```
#include <tbb/task.h>

int main() {
    // Assuming RootTask, ChildTask1, and ChildTask2 are defined.
    RootTask& root = *new(tbb::task::allocate_root()) RootTask{};
    ChildTask1& child1 = *new(root.allocate_child()) ChildTask1{/*params*/};
    ChildTask2& child2 = *new(root.allocate_child()) ChildTask2{/*params*/};
    root.set_ref_count(3);
    tbb::task::spawn(child1);
    tbb::task::spawn(child2);
    root.wait_for_all();
}
```

Example 1: Spawning of Individual Tasks Using `oneapi::tbb::task_group`

```
#include <oneapi/tbb/task_group.h>
```

```
int main() {  
    // Assuming ChildTask1, and ChildTask2 are defined.  
    oneapi::tbb::task_group tg;  
    tg.run(ChildTask1{/*params*/});  
    tg.run(ChildTask2{/*params*/});  
    tg.wait();  
}
```

Example 1: Spawning of Individual Tasks

Using `oneapi::tbb::parallel_invoke`

```
#include <oneapi/tbb/parallel_invoke.h>
```

```
int main() {  
    // Assuming ChildTask1, and ChildTask2 are defined.  
    oneapi::tbb::parallel_invoke(  
        ChildTask1{/*params*/},  
        ChildTask2{/*params*/}  
    );  
}
```

Example 2: Adding More Work During Task Execution in Classic TBB

```
#include <tbb/task.h>
// Assuming necessary entities are defined and implement required interfaces
struct Task : public tbb::task {
    Task(tbb::task& root, int i) : m_root(root), m_i(i) {}
    tbb::task* execute() override {
        // ... do some work for item m_i ...
        if (add_more_parallel_work) {
            tbb::task& child = *new(m_root.allocate_child()) OtherWork;
            tbb::task::spawn(child);
        }
        return nullptr;
    }
    // ...
};
```


Example 2: Adding More Work During Task Execution in Classic TBB

```
#include <tbb/task.h>
// Assuming necessary entities are defined and implement required interfaces
struct Task : public tbb::task {
    Task(tbb::task& root, int i) : m_root(root), m_i(i) {}
    tbb::task* execute() override {
        // ... do some work for item m_i ...
        if (add_more_parallel_work) {
            tbb::task& child = *new(m_root.allocate_child()) OtherWork;
            tbb::task::spawn(child);
        }
        return nullptr;
    }
    // ...
};
```

Example 2: Adding More Work During Task Execution Using `oneapi::tbb::parallel_for_each`

```
#include <vector>
#include <oneapi/tbb/parallel_for_each.h>

int main() {
    std::vector<int> items = { 0, 1, 2, 3, 4, 5, 6, 7 };
    oneapi::tbb::parallel_for_each(
        items.begin(), items.end(),
        [](int& i, tbb::feeder<int>& feeder) {
            // ... do some work for item i ...
            if (add_more_parallel_work)
                feeder.add(i);
        }
    );
}
```

Example 2: Adding More Work During Task Execution Using `oneapi::tbb::task_group`

```
#include <vector>
#include <oneapi/tbb/task_group.h>
int main() {
    std::vector<int> items = { 0, 1, 2, 3, 4, 5, 6, 7 };
    oneapi::tbb::task_group tg;
    for (std::size_t i = 0; i < items.size(); ++i) {
        tg.run([&i = items[i], &tg] {
            // ... do some work for item i ...
            if (add_more_parallel_work)
                tg.run(OtherWork{});
        });
    }
    tg.wait();
}
```

Example 3: Deferred Task Creation Using Classic TBB

```
#include <tbb/task.h>
int main() {
    // Assuming RootTask, ChildTask, and CallbackTask are defined.
    RootTask& root = *new(tbb::task::allocate_root()) RootTask{};
    ChildTask& child = *new(root.allocate_child()) ChildTask{/*params*/};
    CallbackTask& cb_task = *new(root.allocate_child()) CallbackTask{/*params*/};
    root.set_ref_count(3);
    tbb::task::spawn(child);
    register_callback( [cb_task&]() { tbb::task::enqueue(cb_task); } );
    root.wait_for_all();
    // Control flow will reach here only after both ChildTask and CallbackTask are
    // executed, i.e. after the callback is called
}
```

Example 3: Deferred Task Creation Using Classic TBB

```
#include <tbb/task.h>
int main() {
    // Assuming RootTask, ChildTask, and CallbackTask are defined.
    RootTask& root = *new(tbb::task::allocate_root()) RootTask{};
    ChildTask& child = *new(root.allocate_child()) ChildTask{/*params*/};
    CallbackTask& cb_task = *new(root.allocate_child()) CallbackTask{/*params*/};
    root.set_ref_count(3);
    tbb::task::spawn(child);
    register_callback( [cb_task&]() { tbb::task::enqueue(cb_task); } );
    root.wait_for_all();
    // Control flow will reach here only after both ChildTask and CallbackTask are
    // executed, i.e. after the callback is called
}
```

Example 3: Deferred Task Creation Using `oneapi::tbb::task_group`

```
#include <oneapi/tbb/task_group.h>
int main(){
    oneapi::tbb::task_group tg;
    oneapi::tbb::task_arena arena;
    // Assuming ChildTask and CallbackTask are defined.

    auto cb = tg.defer(CallbackTask{/*params*/});
    register_callback( [&tg, c = std::move(cb), &arena]{ arena.enqueue(c); } );

    tg.run(ChildTask{/*params*/});
    tg.wait();
    // Control flow gets here once both ChildTask and CallbackTask are executed
    // i.e. after the callback is called
}
```

Example 3: Deferred Task Creation Using `oneapi::tbb::task_group` (alternative)

```
#include <oneapi/tbb/task_group.h>
int main(){
    oneapi::tbb::task_group tg;
    // Assuming ChildTask and CallbackTask are defined.

    auto cb = tg.defer(CallbackTask{/*params*/});
    register_callback( [&tg, c = std::move(cb)]{ tbb::this_task_arena::enqueue(c); } );

    tg.run(ChildTask{/*params*/});
    tg.wait();
    // Control flow gets here once both ChildTask and CallbackTask are executed
    // i.e. after the callback is called
}
```

Other Changes: New Community Preview Features

- Task arena interface extension to support Hybrid CPUs
- Collaborative call once
- Extended the high-level task API to simplify migration from TBB to oneTBB
- Task scheduler handle to support waiting of worker thread termination
- Added heterogeneous lookup, erase, insert in concurrent hash map
- etc.

Other Changes: New Features

- Defer execution of a task in the task group
- Specify arbitrary `task_group_context` for the `task_group`
- Task arena interface extension to specify priority of the arena
- Support of latest C++ standards and compilers
- Support of Address Sanitizer and Thread Sanitizer
- Several community preview features are now officially supported
 - Concurrent ordered containers
 - Task arena interface extension for NUMA
 - Flow Graph API to support relative priorities for functional nodes
 - Resumable tasks

Conclusion

- oneTBB is a revamped version of TBB
 - Included into Intel® oneAPI Toolkits
 - Migration Guide (https://oneapi-src.github.io/oneTBB/main/tbb_userguide/Migration_Guide.html)
 - oneTBB team can help in migration!
- Available on GitHub (<https://github.com/oneapi-src/oneTBB>)
 - Anyone can open an Issue, suggest a Pull Request
- Includes CMake support and is included into popular package managers
- Backward compatible with previous releases of oneTBB

intel®