

# Intel Thread Building Blocks, Part II

SPD course 2012-13

Massimo Coppola

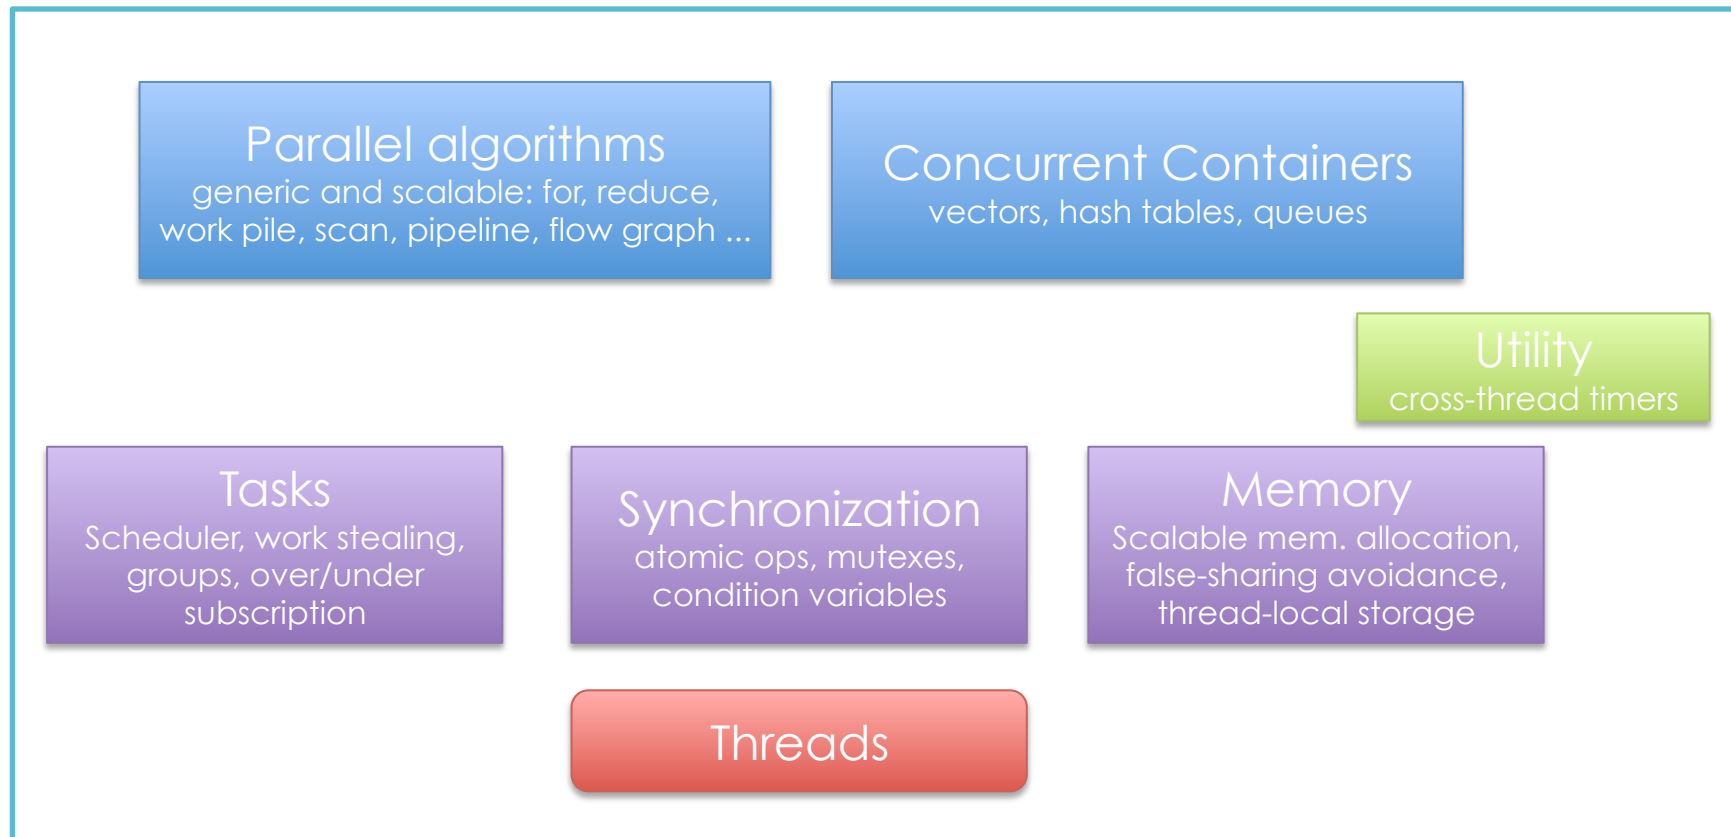
10/12/2012

# TBB Recap

- Portable environment
  - Based on C++11 standard compilers
  - Extensive use of templates
- No vectorization support (portability)
  - use vector support from your specific compiler
- Full environment: compile time + runtime
- Runtime includes
  - memory allocation
  - synchronization
  - task management
- TBB supports patterns as well as other features
  - algorithms, containers, mutexes, tasks...
  - mix of high and low level mechanisms
  - programmer must choose wisely

# TBB “layers”

- All TBB architectural elements are present in the user API, **except** the actual threads



# Threads and composability

- Composing parallel patterns
  - a pipeline of farms of maps of farms
  - a parallel for nested in a parallel loop within a pipeline
  - each construct can express more potential parallelism
  - deep nesting → too many threads → overhead
- Potential parallelism should be expressed
  - difficult or impossible to extract for the compiler
- Actual parallelism should be flexibly tuned
  - messy to define and optimize for the programmer, performance hardly portable
- TBB solution
  - Potential parallelism = tasks
  - Actual parallelism = threads
  - Mapping tasks over threads is largely automated and performed at run-time

# Tasks vs threads

- Task is a unit of computation in TBB
  - can be executed in parallel with other tasks
  - the computation is carried on by a thread
  - task mapping onto threads is a choice of the runtime
    - the TBB user can provide hints on mapping
- Effects
  - Allow **Hierarchical Pattern Composability**
  - raise the level of abstraction
    - avoid dealing with different thread semantics
  - increase run-time portability across different architectures
    - adapt to different number of cores/threads per core

# TBB 4 Algorithms (1)

Over time, the distinction between parallel patterns and algorithms may become blurred

- **parallel\_for**
  - iteration over a range, can choose partitioner
- **parallel\_for\_each**
  - iteration via simple iterator, no partitioner choice
- **parallel\_do**
  - iteration over a set, may add items
- **parallel\_reduce**
  - reduction over a range, can choose partitioner, has deterministic variant
- **parallel\_scan**
  - parallel prefix over a range, can choose partitioner

# TBB 4 Algorithms (2)

- **parallel\_scan**
  - parallel prefix over a range, can choose partitioner
- **parallel\_while**
  - iteration over a stream, may add items
- **parallel\_sort**
  - sort over a set (via a RandomAccessIterator and compare function)
- **pipeline** and **filter**
  - runs a pipeline of filter stages, tasks in = tasks out
- **parallel\_invoke**
  - execute a group of tasks in parallel
- **thread\_bound\_filter**
  - a filter explicitly bound to a serving thread

# Parallel For / For each

```
void tbb::parallel_for_each (InputIterator first,  
                             InputIterator last, const Function &f)
```

- simple case, employs iterators
- was a special case of for in previous TBB
- equivalent to:

```
for (auto i=first; i<last; ++i) f(i);
```

- there is a variant specifying the context in which the f() is evaluated



```
parallel_for (  
    tbb::blocked_range<size_t> (begin, end,  
    GRAIN_SIZE), tbb_parallel_task());
```

- one way of specifying it, where `tbb_parallel_task` is a class
- TBB can use a class for parallel loop implementations.
  - The actual loop "chunks" are performed using the `()` operator of the class
  - data are passed via the class and the range
  - the computing function will receive a range as parameter
- The computing function can also be defined in-place via lambda expressions

```
parallel_for (  
    tbb::blocked_range<size_t> (begin, end,  
    GRAIN_SIZE), tbb_parallel_task(), partitioner);
```

- the partitioner is one of those specified by TBB (simple, auto, affinity)
- no real choice, just allocate a const partitioner and pass it to the parallel loops:

```
tbb::affinity_partitioner ap;
```

– (unless you want to define your own)

# TBB Range classes

- Range classes express intervals of parameter values and their decomposability
  - **recursively** splitting intervals to produce parallel work
- The Range concept relies on five methods
  - copy constructor
  - destructor
  - `is_divisible()` true if range not too small
  - `empty()` true if range empty
  - `split()` split the range in two parts

# The Range concept

Class R implementing the concept of range must define:

```
R::R( const R& );  
R::~~R();  
bool R::is_divisible() const;  
bool R::empty() const;  
R::R( R& r, split );
```

Split range R into two subranges.

One is returned via the parameter, the other one is the range itself, accordingly reduced

# Blocked Range

- TBB 4 has implementations of the range concept as templates for 1D, 2D and 3D blocked ranges
  - 3 nested parallel for are functionally equivalent to a simple parallel over a 3D range
  - the 2D and 3D range will likely exploit the caches better, due to the explicit 2D/3D tiling

```
tbb::blocked_range< Value > Class
```

```
tbb::blocked_range2d< RowValue, ColValue > Class
```

```
tbb::blocked_range3d< PageValue,  
                    RowValue, ColValue > Class
```

# partitioners

- simple
  - generate tasks by dividing the range as much as possible (remember about the grain size!)
- auto
  - divide into large chunks, divide further if more tasks are required
- affinity
  - carries state inside, will assign the tasks according to range locality to better exploit caches

# Combining the elements

- Apply a range template to your elementary data type
- Define a class computing the proper for-body over elements of a range
- Call the `parallel_for` passing at least the range and the function
- specify a partitioner and/or a grain size to tune task creation for load balancing

# Example (with lambda)

```
void relax( double *a, double *b,
           size_t n, int iterations)
{
    tbb::affinity_partitioner ap;
    for (size_t t=0; t<iterations; ++t) {
        tbb::parallel_for(
            tbb::blocked_range<size_t>(1, n-1),
            [=] ( tbb::blocked_range<size_t> r) {
                size_t e = r.end();
                for (size_t i=r.begin(), i<e; ++i)
                    /*do work on a[i], b[i] */;
            },
            ap);
        std::swap(a,b); // always read from a, write to b
    }
}
```



# Updated References

- Download docs and code from <http://threadingbuildingblocks.org/>
- Since TBB 4
  - many of the accompanying PDF (tutorial, reference) are no longer made available on the web site. Either
  - ask the teacher for TBB 3.0 copies
  - resort to books
- TBB Accompanying docs
  - download the full TBB source archive, it contains
    - an **example** directory with TBB examples and their description
    - a **doc** directory with full html reference docs
  - Getting started – install and compile examples **← TRY IT**
- Quick summary to lambda expressions in C++
  - [http://www.nacad.ufrj.br/online/intel/Documentation/en\\_US/compiler\\_c/main\\_cls/cref\\_cls/common/cppref\\_lambda\\_desc.htm](http://www.nacad.ufrj.br/online/intel/Documentation/en_US/compiler_c/main_cls/cref_cls/common/cppref_lambda_desc.htm)