

High-level parallel programming in C++

Dénes Mátételki

www.emerson.com

March 18, 2012

Table of contents

- 1 Theory
 - High level vs. low level
 - Algorithms
- 2 Code samples
 - Map
 - Reduce
 - Sort
- 3 Final thoughts
 - Grainsize
 - Convolution
 - Summary

Comparison

High level

- Auto scaling-up
- Threadpool handling, load balancing.
- Synchronization and mutexes are handled.

Low level

- Manual thread creation.
- Manual joins and mutex handling.
- Better for event and I/O based threading.
- Compiler and external library independend.

Compared softwares (performance, code complexity)

Used

- Standard c++ (serial examples)
- openMP[1]
- Intel Thread Building Blocks (TBB)[2]
- QtConcurrent[3]

Skipped

- std::thread, std::mutex (c++0x)[6]
- POSIX threads[7]
- QThread[8]

Co-existence[5]

Possible, but the separate threadpools can lead to oversubscription.

Comparison

openMP

- Compiler support needed.
- C, C++, fortran.
- Best for bounded loops.
- No need for big code re-write.
- Hard to debug.
- Managed by a non-profit organization.

Intel TBB

- Object oriented.
- Concurrent data types.
- Parallel algorithms.
- Work stealing: dynamic load sharing.
- Relies heavily on templates.
- Heavy code rewrite is needed.

QtConcurrent

- Object oriented
- Limited number of algorithms.
- ...

Used algorithms for testing

List

- Map - Applies a given function to each element of a container.
- Reduction - Combines the results of sub-parts.
- Sort - Puts elements of a list in a certain order.

Note

- The used container is an `std::vector<float>`
- Container size was 60 million with random floats [1, 1000]
- Execution times are the averages of 3 executions.
- Used hardware was an Intel Xeon 64-bit machine with 6 cores (12 threads), 3,4Mz.
- Compiled with `gcc-4.4` and use flags: `-O3 -ffast-math -fwhole-program -fomit-frame-pointer -march=native -m64`

Serial map

c++ code

```
1 float modify(float value)
2 {
3     return 13.37 * pow(sqrt(value), log(value));
4 }
5
6
7 void serialMap(std::vector<float>& data)
8 {
9     for (size_t i = 0; i < data.size(); i++)
10        modify(data[i]);
11 }
```

Note

- “chunksize” equals the size of the data.
- This modify function will be used by the parallel examples too.

openMP parallel map

c++ code

```
1 void openMpMap(std::vector<float>& data,  
2               const int numberOfThreads,  
3               const int chunkSize)  
4 {  
5     size_t i;  
6  
7     #pragma omp parallel for \  
8     default(shared) private(i) \  
9     schedule(dynamic, chunkSize) \  
10    num_threads(numberOfThreads)  
11  
12    for (i = 0; i < data.size(); i++)  
13        data[i] = modify(data[i]);  
14 }
```

Note

Making it run in parallel is just a single pragma line.

Intel TBB map

c++ code

```
1  class itbbMap {
2  public:
3
4      itbbMap(std::vector<float>& data)
5          : m_data(data) {}
6
7      void operator()(const tbb::blocked_range<size_t>& r) const {
8          for( size_t i = r.begin(); i != r.end(); i++ )
9              m_data[i] = modify(m_data[i]);
10     }
11
12 private:
13     std::vector<float>& m_data;
14 };
15
16
17 tbb::task_scheduler_init init(NUMBER_OF_THREADS);
18 itbbMap im(data);
19 tbb::parallel_for(tbb::blocked_range<size_t>(0, data.size(), CHUNK_SIZE), im);
```

Note

Running a functor on chunks in parallel.

QtConcurrent map


c++ code

```
1 void QtMap(std::vector<float>& data)
2 {
3     QtConcurrent::blockingMap(data, modify);
4 }
5
6 QThreadPool::globalInstance()->setMaxThreadCount(NUMBER_OF_THREADS);
```

Note

- Chunksize is 1.
- Blocks till the iterator reaches the end.

Map execution times



Note

Serial remained the fastest (memory bound?) - No need to parallelize.

Serial reduce

c++ code

```
1 float serialReduce(std::vector<float>& data)
2 {
3     float min(FLT_MAX);
4     for (size_t i = 0; i < data.size(); i++)
5         if (data[i] < min)
6             min = data[i];
7
8     return min;
9 }
```

Note

- Minimum value search.
- Not actually a reduce.
- Following examples will try to achive this too.

openMP reduce

c++ code

```
1  int openMpReduce(std::vector<float>& data,
2      const int numberOfThreads,
3      const int chunkSize)
4  {
5      size_t i;
6      std::vector<float> separate_results(numberOfThreads, FLT_MAX);
7
8      #pragma omp parallel \
9          default(shared) private(i) \
10         num_threads(numberOfThreads)
11     {
12         int threadId = omp_get_thread_num();
13
14         #pragma omp for schedule(dynamic, chunkSize)
15
16         for (i = 0; i < data.size(); i++)
17             if (separate_results[threadId] < data[i])
18                 separate_results[threadId] = data[i];
19     }
20
21     float min(FLT_MAX);
22     for (i = 0; i < numberOfThreads; i++)
23         if (separate_results[i] < min)
24             min = separate_results[i];
25
26     return min;
27 }
```

Intel TBB reduce

c++ code


```
1  class itbbReduce {
2      const std::vector<float>& m_data;
3  public:
4      float m_min;
5
6      itbbReduce(std::vector<float>& data) : m_data(data) , m_min(FLT_MAX) {}
7      itbbReduce(itbbReduce& other, tbb::split) : m_data(other.m_data), m_min(FLT_MAX) {}
8
9      void operator()(const tbb::blocked_range<size_t>& r) {
10         float min = m_min;
11         for(size_t i = r.begin(); i != r.end(); i++)
12             if ( m_data[i] < min )
13                 min = m_data[i];
14
15         m_min = min;
16     }
17
18     void join(const itbbReduce& other) {
19         if ( other.m_min < m_min )
20             m_min = other.m_min;
21     }
22 };
23
24 itbbReduce mif(data);
25 tbb::parallel_reduce(tbb::blocked_range<size_t>(0, data.size(), CHUNK_SIZE), mif);
26 float min = mif.m_min;
```

QtConcurrent reduce

c++ code

```
1 void findMinimum(const std::vector<float>::const_iterator begin,
2                 const std::vector<float>::const_iterator end,
3                 float *result)
4 {
5     result = std::min_element(begin, end);
6 }
7
8
9 float QtReduce(std::vector<float>& data,
10               const int numberOfThreads,
11               const int chunkSize)
12 {
13     std::vector<float> separate_results(numberOfThreads, FLT_MAX);
14     QFutureSynchronizer<void> synchronizer;
15
16     for(int i = 0; i < numberOfThreads; i++)
17         synchronizer.addFuture(QtConcurrent::run(findLocalMinimum,
18                                                  data.begin()+i*chunkSize,
19                                                  data.begin()+(i+1)*chunkSize,
20                                                  separate_results.data()+i));
21
22     synchronizer.waitForFinished();
23
24     float min(FLT_MAX);
25     findMinimum(separate_results.begin(), separate_results.end(), min);
26     return min;
27 }
```

Reduce execution times



Note

No need for more than 4 threads.

Serial sort

c++ code

```
1 void serialSort(std::vector<float>& data)
2 {
3     std::sort(data.begin(), data.end());
4 }
```

Note: quicksort

- Pick a pivot point.
- Partition: Swap elements compared to pivot point.
- Recursively calls itself with the 2 new partitions.

openMP, Intel TBB sort

openMP c++ code

```
1  #include <parallel/algorithm>
2
3  void openMpSort(std::vector<float>& data)
4  {
5      __gnu_parallel::sort(data.begin(), data.end());
6  }
```


Note

Some algorithms are already rewritten to work in parallel with openMP.

Intel TBB c++ code

```
1  void itbbSort(std::vector<float>& data)
2  {
3      tbb::parallel_sort(data.begin(), data.end());
4  }
```

Sort execution times



Note

No need for more than 6 threads.

Custom QtConcurrent sort

c++ code

```

1  template <class SortType>
2  long QsPartition(SortType outputArray[], long left, long right) { ... }
3
4  template <class SortType>
5  void QsSequential(SortType array[], const long left, const long right) { ... }
6
7  template <class SortType>
8  void QuickSortTask (SortType array[], const long left, const long right, const int deep)
9  {
10     if (left < right) {
11         if (deep) {
12             const long part = QsPartition(array, left, right);
13             QtConcurrent::run(QuickSortTask<SortType>, array, part + 1, right, deep - 1);
14             QtConcurrent::run(QuickSortTask<SortType>, array, left, part - 1, deep - 1);
15         } else {
16             const long part = QsPartition(array, left, right);
17             QsSequential(array, part + 1, right);
18             QsSequential(array, left, part - 1);
19         }
20     }
21 }
22
23 void QtSort(std::vector<float>& data)
24 {
25     QtConcurrent::run(QuickSortTask<float>, data.data(), 0, data.size() - 1, 6);
26     QThreadPool::globalInstance()->waitForDone();
27 }


```

Custom openMP sort

c++ code

```
1 void sample_qsort(float* begin, float* end) { ... }
2
3 void sample_qsort_serial(float* begin, float* end) { ... }
4
5 void sample_qsort_adaptive(float* begin, float* end, const long nthreshold)
6 {
7     if (begin != end) {
8         // partition ...
9         if (end - begin + 1 <= nthreshold) {
10            sample_qsort_serial(begin, middle);
11            sample_qsort_serial(++middle, ++end);
12        } else {
13            #pragma omp task
14            sample_qsort_adaptive(begin, middle, nthreshold);
15            #pragma omp task
16            sample_qsort_adaptive(++middle, ++end, nthreshold);
17        }
18    }
19 }
20
21 void sample_qsort_adaptive(float* begin, float* end)
22 {
23     long nthreshold = ceil(sqrt(end - begin + 1)) / 2;
24     #pragma omp parallel
25     #pragma omp single nowait
26     sample_qsort_adaptive(begin, end, nthreshold);
27 }
```

Sort times of custom algorithms



Note

Container size is 6M - miserable...

Two quicksort approach to

Treshold

```

1 void qsort(float* begin,
2           float* end,
3           const long nthreshold)
4 {
5     if (begin != end) {
6         // partition ...
7         if (end-begin+1 <= nthreshold) {
8             // serial sort ...
9         } else {
10            // parallel sort ...
11        }
12    }
13 }
14
15 long deep =
16     ceil(sqrt(end - begin + 1)) / 2;

```

Depth

```


1 void qsort(float* begin,
2           float* end,
3           const int deep)
4 {
5     if (begin != end) {
6         // partition ...
7         if (deep) {
8             // serial sort ...
9         } else {
10            // parallel sort with deep-1
11        }
12    }
13 }
14
15 long deep = 15;

```

Note

Depth seems simpler yet faster.


Chunk size



Case A



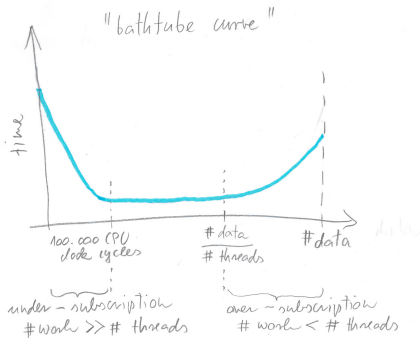
Case B

 : overhead
white area inside is the same

Note

- Unit is loop interaction per chunk. Default value is 1.
- Too small chunks can introduce more overhead than useful work.

Grain size



Note

- Unit is CPU cycles.
- Should be at least 100,000.

Task stealing - Intel TBB

Task stealing

- Each thread has a queue of tasks.
- If a thread has no more tasks then it “steals” from another.
- Think about tasks, not about threads when programming.

Threadpool

A threadpool with a common concurrent queue of tasks is a common practice in networking servers.


Work stealing

Another implementation is Cilk[4] - where each processor has a stack of frames.

1D gaussian filter

c++ code

```
1 void serialConvolution(std::vector<float>& output,  
2                       const std::vector<float>& input,  
3                       const std::vector<float>& kernel)  
4 {  
5     // skipping the edges: separate loops, paddings  
6     // output.size == input.size()-kernel.size()-1;  
7  
8     for (size_t i = 0; i < output.size(); i++) {  
9         float sum = 0;  
10        for (size_t j = 0; j <= kernel.size(); j++)  
11            sum += input[i+j] * kernel[j];  
12  
13        output[i] = sum;  
14    }  
15 }
```



Note

```
float kernel[7] = { 0.06, 0.061, 0.242, 0.383, 0.242, 0.061, 0.06 }
```

Optimized convolution


c++ code

```

1 void operator()(const tbb::blocked_range<size_t>& r) const
2 {
3     // skipping the edges, shall be done in separate task
4     const float* p = &m_input[0] + r.begin();
5     float* d = &m_output[0] + r.begin();
6
7     const size_t n = m_kernel.size();
8     float k[n]; // pre-read kernel
9     float c[n]; // pre-read values
10    k[0] = m_kernel[0];
11    for (size_t i = 1; i < n; ++i) {
12        c[i] = m_input[i-1];
13        k[i] = m_kernel[i];
14    }
15
16    // chunk size % kernel.size() != 0 should be handled...
17    for (size_t i = 0; i < r.size(); i += n) {
18        d[i+0] = (c[0] = p[i+0])*k[0]+c[1]*k[2]+c[2]*k[2]+c[3]*k[3]+c[4]*k[4]+c[5]*k[5]+c[6]*k[6];
19        d[i+1] = (c[6] = p[i+1])*k[0]+c[0]*k[2]+c[1]*k[2]+c[2]*k[3]+c[3]*k[4]+c[4]*k[5]+c[5]*k[6];
20        d[i+2] = (c[5] = p[i+2])*k[0]+c[6]*k[2]+c[0]*k[2]+c[1]*k[3]+c[2]*k[4]+c[3]*k[5]+c[4]*k[6];
21        d[i+3] = (c[4] = p[i+3])*k[0]+c[5]*k[2]+c[6]*k[2]+c[0]*k[3]+c[1]*k[4]+c[2]*k[5]+c[3]*k[6];
22        d[i+4] = (c[3] = p[i+4])*k[0]+c[4]*k[2]+c[5]*k[2]+c[6]*k[3]+c[0]*k[4]+c[1]*k[5]+c[2]*k[6];
23        d[i+5] = (c[2] = p[i+5])*k[0]+c[3]*k[2]+c[4]*k[2]+c[5]*k[3]+c[6]*k[4]+c[0]*k[5]+c[1]*k[6];
24        d[i+6] = (c[1] = p[i+6])*k[0]+c[2]*k[2]+c[3]*k[2]+c[4]*k[3]+c[5]*k[4]+c[6]*k[5]+c[0]*k[6];
25    }
26 }

```

Convolution running times



Note

Memory-read optimization can result the same performance improvements as parallelization.

Things to keep in mind

Checklist

- Pass primitive types by value.
- Pass objects by address.
- Have function-local copies of member variables.
- Avoid to read values multiple times.
- Choose correct chunk size.
- Instead of shared memory, consider reduction.
- Plan datastructures to avoid memory-boundings.*

Things to keep in mind

Checklist

- Pass primitive types by value.
- Pass objects by address.
- Have function-local copies of member variables.
- Avoid to read values multiple times.
- Choose correct chunk size.
- Instead of shared memory, consider reduction.
- Plan datastructures to avoid memory-boundings.*

*data-oriented design[9]

If only someone could tell us more about it...

Links



[openMP.http://openmp.org](http://openmp.org)



[Intel Thread Building Blocks.http://threadingbuildingblocks.org/](http://threadingbuildingblocks.org/)



[QtConcurrent.http://doc.qt.nokia.com/4.8-snapshot/qtconcurrent.html](http://doc.qt.nokia.com/4.8-snapshot/qtconcurrent.html)



[Cilk.http://software.intel.com/en-us/articles/intel-cilk-plus](http://software.intel.com/en-us/articles/intel-cilk-plus)



[Comparison of Intel TBB, openMP and native threads.http://software.intel.com/en-us/articles/intel-threading-building-blocks-openmp-or-native-threads/](http://software.intel.com/en-us/articles/intel-threading-building-blocks-openmp-or-native-threads/)



[std::thread in C++http://en.cppreference.com/w/cpp/thread](http://en.cppreference.com/w/cpp/thread)



[POSIX threads tutorial.http://www.yolinux.com/TUTORIALS/LinuxTutorialPosixThreads.html](http://www.yolinux.com/TUTORIALS/LinuxTutorialPosixThreads.html)



[Qt threads.http://qt-project.org/doc/qt-4.8/threads.html](http://qt-project.org/doc/qt-4.8/threads.html)



[Data oriented design.http://gamesfromwithin.com/data-oriented-design](http://gamesfromwithin.com/data-oriented-design)



[L^AT_EX beamer class for creating presentations.https://bitbucket.org/rivanvx/beamer/wiki/Home](https://bitbucket.org/rivanvx/beamer/wiki/Home)



[Gnuplot - An open source plotting software.http://www.gnuplot.info/](http://www.gnuplot.info/)