

Introducing Parallel Pixie Dust: Advanced Library-Based Support for Parallelism in C++ BCS-APSG March 2011

Jason McGuiness¹

ppd-bcs-aps-g-11@hussar.demon.co.uk
<http://libjmmcg.sf.net/>

February 23, 2011

¹Copyright © J.M.McGuiness, 2011.

Sequence of Presentation.

- ▶ An introduction: why I am here.
- ▶ Why do we need multiple threads?
 - ▶ Multiple processors, multiple cores....
- ▶ How do we manage all these threads:
 - ▶ libraries, languages and compilers.
- ▶ My problem statement, an experiment if you will...
- ▶ An outline of Parallel Pixie Dust (PPD).
- ▶ Important theorems.
- ▶ Examples, and their motivation.
- ▶ Conclusion.
- ▶ Future directions.

Introduction.

Why yet another thread-library presentation?

- ▶ Because we still find it hard to write multi-threaded programs correctly.
 - ▶ According to programming folklore.
- ▶ We haven't successfully replaced the von Neumann architecture:
 - ▶ Stored program architectures are still prevalent.
- ▶ The memory wall still affects us:
 - ▶ The CPU-instruction retirement rate, i.e. rate at which programs require and generate data, exceeds the the memory bandwidth - a by product of Moore's Law.
 - ▶ Modern architectures add extra cores to CPUs, in this instance, extra memory buses which feed into those cores.

Why do We Need Multiple Threads?

The memory wall is a major influencing factor:

- ▶ A definition of the memory wall is the effect that one gets when the rate at which a processor may retire instructions exceeds the bandwidth of the memory sub-system.
- ▶ The use of Amdahl's law has been highly effective in concealing this issue: 90% of the time only 10% of the code (or data) is used.
- ▶ Another factor is locality of data both in space and time.
 - ▶ This means that cache hierarchies work very effectively.
 - ▶ But cache hierarchies do not overcome the latency of memory access, just the bandwidth.
 - ▶ As the gap between memory bandwidth and processor instruction retirement rates has widened, caches have become larger and the hierarchies deeper to cope.
 - ▶ So they have become more expensive.

How to Achieve Greater Instruction-Retirement Rates?

The increase in cost has meant alternatives have become viable:

- ▶ Such as the use of multiple processors & multiple cores.
 - ▶ Which may be logically viewed as multiple execution units retiring instructions.
- ▶ We now have n -cores retiring instructions, et voilà, we have increased the instruction retirement rate.
 - ▶ *If we can keep the n execution units running.*
- ▶ So now we have some sort of thread of execution, either architectural or OS-level, or a hybrid.
- ▶ These threads need to be managed. Therein the problems lie...

How many Threads?

The future appears to be laid out: inexorably more execution units.

- ▶ For example quad-core processors are becoming more popular.
 - ▶ 4-way blade frames with quad cores provide 16 execution units in close proximity.
 - ▶ [picoChip] (from Bath, U.K.) has hundreds of execution units per chip, with up to 16 in a network.
 - ▶ Future supercomputers will have many millions of execution units, e.g. IBM BlueGene/C [Cyclops64].

The Problems with Threading: a Brief Review.

Let's leave aside the maxim: "The first rule of optimization is don't do it. The second rule is don't do it yet."²

We'll presume we need to create threads. Creating threads is easy!!! (`_beginthreadex()`, `pthread_create()`, etc...)

- ▶ But synchronizing access to the data they mutate is hard:
 - ▶ Problem: race-conditions or deadlocks may arise. i.e. ensuring that the threaded code will run correctly for all "reasonable" data inputs is problematic, and effectively impossible.
 - ▶ Ensuring that the threaded code runs an efficient schedule is usually just ignored.
- ▶ Folklore not just in computer-science but in general programming, states that writing correct multi-threaded programs is a **black art**.

I do not disagree!

²Attributed to Michael Jackson,

How do We Solve these Problems?

- ▶ There are various attempts to solve the threading problem:
 - ▶ Language level:
 - ▶ e.g. HPF[1] and UPC[2] or Threaded-C: highly successful but mainly restricted to the domain of HPC.
 - ▶ Functional Java - an emergent language.
 - ▶ Library level:
 - ▶ e.g. PThreads, Win32, Open-MP and MPI: very low-level, usually supporting very primitive threading-related operations.
 - ▶ [Intel Thread Building Blocks, 3] and [Concurrent Collections] - very interesting.
 - ▶ [Boost.Thread, 4] - influenced by the C++0x standardization effort, and Anthony Williams' excellent [just::thread], amongst others.
 - ▶ Auto-threading compilers:
 - ▶ e.g. MSVC and IBM VisualAge currently support a limited form of Open-MP relating to loop scheduling.
 - ▶ e.g. HPF, EARTH-C. These generate guaranteed correct and optimal schedules! But few outside the scientific community use them. EARTH-C even solves the hardest problems to parallelize, but for the IBM EARTH-MANNA architecture [5].

A Quick Review of Some Threading Models:

- ▶ Restrict ourselves to library-based techniques.
- ▶ Raw library calls.
- ▶ The “thread as an active class”.
- ▶ The “thread pool” containing those objects.
- ▶ The issue: classes used to implement business logic also implement the operations to be run on the threads. This often means that the locking is intimately entangled with those objects, and possibly even the threading logic.
 - ▶ This makes it much harder to debug these applications. The question of how to implement multi-threaded debuggers correctly is still an open question.

Problems Arising from these Models.

In the face of tens, let alone millions, of threads this approach is untenable:

- ▶ Testing the business logic is a nightmare.
- ▶ Testing the threading in the face of maintenance upgrades is a practical impossibility - the code-base becomes locked into libraries and compilers.
- ▶ Performance can become sub-linear due to excessively conservative locking!
- ▶ The business cost of maintaining these balls of mud becomes untenable.
- ▶ We face the failure of using multiple execution units as a technique to overcome the memory wall.

And there's no other commercial architecture out there to compete..... *YIKES!!!!*

What am I looking for?

I decided to set myself a challenge, an experiment if you will, which may be summarized in this problem statement:

- ▶ Is it possible to create a cross-platform thread library that can guarantee a schedule that is deadlock & race-condition free, is optimal (mathematically), and also assists in debugging the business logic, in a general-purpose, imperative language that has little or no native threading support? Moreover the library should offer considerable flexibility, e.g. reflect the various types of thread-related costs and techniques.

Further Constraints I set Myself.

- ▶ The imperative language used is C++: I know it, and it doesn't have the library that I want, yet.
- ▶ I focused on *general purpose threading*:
 - ▶ The situation when loop-bounds or conditions are not known at compile-time.
 - ▶ If these are statically computable, then conditions may be removed, and loops may be statically unrolled, rescheduled, etc as is well-known in the literature and well implemented in HPF, much more efficiently.
 - ▶ This type of scheduling may be viewed as dynamic scheduling (run-time) as opposed to static scheduling (potentially compile-time), where the operations are statically assigned to execution pipelines, with relatively fixed timing constraints.
 - ▶ The idea is that it should be useful in general-purpose code, which could be written in a data-flow style.
 - ▶ I specifically avoided investigating issues and threading relating to asynchronous I/O.

What is Parallel Pixie Dust (PPD)?

The main design features of PPD are:

- ▶ It is based upon futures: in fact it implements a software-simulation of data-flow.
- ▶ Although “thread-as-an-active-class” exists, the primary method of controlling threads is thread pools of many different types (think traits).
 - ▶ Futures and thread pools can be used to implement a tree-like thread schedule.
- ▶ Parallel versions of certain binary function objects, so that each operand may be executed in parallel.
- ▶ Adaptors are provided for the STL collections to assist with providing thread-safety.
 - ▶ By using these adaptors and thread pools, the STL algorithms can be replaced with multi-threaded versions.
- ▶ Amongst other influences, PPD was born out of discussions with Richard Harris and motivated by Kevlin Henney’s presentation to the [ACCU Conference '04] regarding threads.

PPD: More Details Regarding the Futures.

The use of futures, termed *execution contexts*, within PPD is crucial:

- ▶ This hides the data-related locks from the user, as the future wraps the retrieval of the data with a hidden lock, the resultant future-like object behaving like a proxy.
- ▶ This is an implementation of the *split-phase constraint*.
- ▶ The futures are also key in passing exceptions across thread boundaries.
- ▶ The void return-type optimization has been implemented.
- ▶ They may be used to delete that work from the thread-pool.
- ▶ They are only stack-based, cannot be heap-allocated and only `const` references may be taken.
 - ▶ This guarantees that there can be no aliasing of these objects. This provides extremely useful guarantees that will be used later.
 - ▶ They can only be created from the returned object when work is transferred into the pool.
 - ▶ There are no (Boost) promises!

PPD: The Thread Pools.

Primary method of controlling threads is thread pools, available in many different types:

- ▶ Work-distribution models: e.g. master-slave, work-stealing.
 - ▶ Size models: e.g. fixed size, unlimited.....
 - ▶ Threading implementations: e.g. *sequential*, PThreads, Win32:
 - ▶ The sequential trait means all threading within the library can be removed, but maintaining the interface. One simple recompilation and all the bugs left are *yours*, i.e. your business model code.
 - ▶ A separation of design constraints. Also side-steps the debugger issue!
 - ▶ Contains a special queue that contains the work in strict FIFO order or user-defined priority order.
- ▶ Threading models: e.g. the cost of performing threading, i.e. heavyweight like PThreads or Win32.
- ▶ Further assist in managing OS threads and exceptions across the thread boundaries.
- ▶ i.e. a PRAM programming model.

PPD: Exceptions.

Passing exceptions across the thread boundary is a thorny issue:

- ▶ Futures are used to receive any exceptions that may be thrown by the mutation when it is executed within the thread pool.
- ▶ If the data passed back from a thread-pool is not retrieved then the future may throw that exception from its destructor!
 - ▶ The design idea expressed is that the exception is important, and if something might throw, then the execution context **must** be examined to verify that the mutation executed in the pool succeeded.
 - ▶ This idea goes further:
 - ▶ All OS threads managed by the thread pool are destroyed when the pool terminates. The implication is that if there was no execution context to pass the exception into, then that thread will re-throw the exception from its destructor.
 - ▶ Ugly as hell, but otherwise what happens? Swallow the exception? That's *uglier*.

PPD: The Concurrency Assistance.

I'm really trying hard to not say “thread safe” as it isn't!

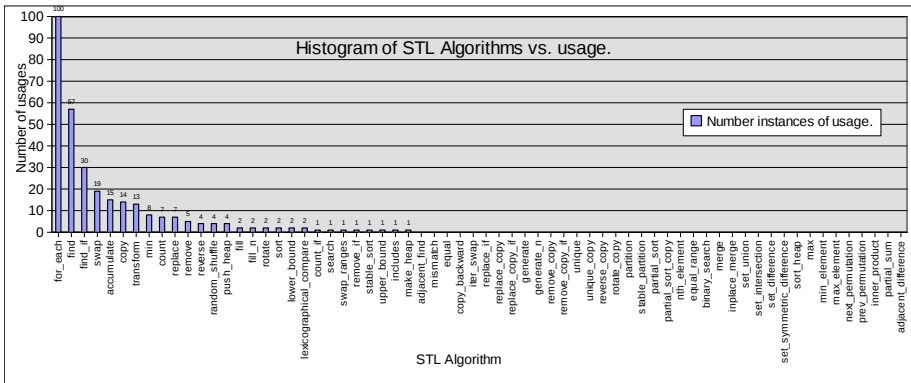
- ▶ The provision of adaptors for the STL collections assist with providing thread-safety.
 - ▶ The adaptors are used to provide the locks required by the thread pool library to maintain thread safety of the adapted collections.
 - ▶ In technical terms the type of lock supplied allows me to specify if the collection supports EREW or CREW operations.
- ▶ STL algorithms are neat ways to express lots of potential parallelism also the size of the collection is dynamic - known only at run-time.
 - ▶ Otherwise use HPFs auto-parallelizing loops!
 - ▶ These collections also assist in making PPD appear to be “drop in” to sequential code.
 - ▶ Naturally one would profile to determine which areas PPD might assist with!

PPD: Ugly Details.

- ▶ All the traits make the objects look quite complex.
 - ▶ typedefs are used so that the thread pools and execution contexts are aliased to more convenient names.
- ▶ It requires undefined behaviour or broken compilers, for example:
 - ▶ Exceptions and multiple threads are unspecified in the current standard. Fortunately most compilers “do the right thing” and implement an exception stack per thread, but this is not standardized.
 - ▶ Most optimizers in compilers are broken with regard to code hoisting. This affects the use of atomic counters and potentially intrinsic locks having code that they are guarding hoisted around them.
 - ▶ `volatile` does not solve this, because `volatile` has been too poorly specified to rely upon. It does not specify that operations on a set of objects must come after an operation on an unrelated object.

An Examination of Sample Code for STL Usage.

A large code-base that has been developed for over 6 years was examined for the distribution of STL algorithms used within it to guide implementation of algorithms within PPD. The target code-base consisting of over 180 projects was examined for usages of 66 of the STL algorithms. The distribution is:



PPD: The STL-style Parallel Algorithms.

That analysis, amongst other reasons lead me to implement:

1. `for_each`
2. `find_if` and `find` (which is trivially based on `find_if`)
3. `count_if` and `count` (which is trivially based on `count_if`)
4. `transform` (both overloads)
5. `copy` (trivially based upon `transform`)
6. `accumulate` (both overloads)

Which means I should cover over 75% of the cases in the code-base examined.

- ▶ Why no merge/sort/swap/min/replace?
 - ▶ merge/sort are harder because implementing a parallel sort (e.g. Batcher's bitonic mergesort or Cole's) well is tricky,
 - ▶ not got around to the rest.
- ▶ So I'm missing an interesting 5% of sorting-type cases.

PPD: Technical Details, Theorems.

Due to the restricted properties of the execution contexts and the thread pools a few important results arise:

1. The thread schedule created is only an acyclic, directed graph. In fact it is a tree.
2. From this property I have proved that the schedule PPD generates is *deadlock and race-condition free*.
3. Moreover in implementing the STL-style algorithms those implementations are optimal, in a mathematical sense, i.e. there are provable bounds on both the execution time and minimum number of processors required to achieve that time³.

³Much of this has come from [6], for which I am grateful to Martin Waplington and [7].

PPD: The Key Theorems: Deadlock & Race-Condition Free.

1. Deadlock Free:

Theorem

That if the user refrains from using any other threading-related items or atomic objects other than those provided by PPD, for example those in 26 and 28, then they can be guaranteed to have a schedule free of deadlocks.

2. Race-condition Free:

Theorem

That if the user refrains from using any other threading-related items or atomic objects other than those provided by PPD, for example those in 26 and 28, and that the work they wish to mutate may not be aliased by another object, then the user can be guaranteed to have a schedule free of race conditions.

PPD: The Key Theorems: Optimal Schedule with Futures.

Theorem

If the user only uses the threading-related items provided by the execution contexts, then the schedule of the work transferred to PPD will be executed in an algorithmic order that is between $O(\log(n))$ and $O(n)$. In those cases PPD should add at most a constant order to the execution time of the schedule.

- ▶ Note that the user might implement code that orders the transfers of work in a sub-optimal manner, PPD has very limited abilities to improve automatically that arrangement of work.
- ▶ This theorem states that once the work has been transferred, PPD will not add any further sub-optimal algorithmic delays, but might add some constant-order delay.

PPD: A Corollary: Optimal Schedule with STL-style algorithms.

Corollary

If the user uses the adaptors, for the STL collections, with CREW-PRAM semantics, and uses the STL-style algorithms, for example 28, then the scheduling algorithm implemented within PPD guarantees the following bounds will hold on that schedule:

- 1. That the optimal time in which n -items of work on p -processors will be executed is: $\left\lceil \frac{n}{p} \right\rceil - 1 + \log p$*
- 2. That the optimal number of processors with which to execute that work is: $\left\lceil \frac{n}{\log n} \right\rceil$*
- 3. That using at least the optimal number of processors, the optimal time is: $\lceil \log n \rceil$*

PPD: A Result of the Optimality Theorems.

If PPD is used for scheduling multi-threaded code then:

Corollary

Using any of the features of PPD will add at worst a $\lceil \log n \rceil$ time to the schedule, which is also the optimal time, then using any other additional threading schedule is likely to make that aggregate schedule worse.

- ▶ i.e. PPD will not make your chosen parallel schedule sub-optimal.
- ▶ If you can express your parallel schedule in PPD, then those portions expressed using PPD will use an optimal schedule.
 - ▶ But that doesn't mean that just because you use PPD, your chosen parallel schedule *will be* optimal.

That's enough theory, I'll stop waving my hands! What about using PPD...

PPD: Basic Example using `execution_contexts`.

Listing 1: General-Purpose use of a Thread Pool and Future.

```
struct res_t {
    int i;
};
struct work_type {
    void process(res_t &) {}
};
pool_type pool(2);
execution_context(context(pool<<joinable()<<time_critical()<<work_type()));
pool.erase(context);
context->i;
```

- ▶ To assist in making the code appear to be readable, a number of typedefs have been omitted.
- ▶ The `execution_context` is created from adding the wrapped work to the pool.
 - ▶ `process(res_t &)` is the only invasive artifact of the library for this use-case.
 - ▶ Alternative member functions may be specified, if a more complex form is used.
 - ▶ Compile-time check: the type of the transferred work is the same as that declared in the `execution_context`.

PPD: The typedefs used.

Listing 2: typedefs in the execution context example.

```
typedef ppd::thread_pool<
    pool_traits::worker_threads_get_work, pool_traits::fixed_size,
    pool_adaptor<
        generic_traits::joinable, platform_api, heavyweight_threading,
        pool_traits::normal_fifo, std::less
    >
> pool_type;
typedef pool_type::create1<work_type>::creator_t;
typedef creator_t::execution_context execution_context;
typedef creator_t::joinable joinable;
typedef priority<pool_type, api_params_type::time_critical> time_critical;
```

- ▶ Express flexibility & optimizations available: also the apparent complexity of the typedefs for the `thread_pool`.
 - ▶ `pool_traits::prioritised_queue`: specifies work in the `thread_pool` should be ordered by the trait `std::less`.
 - ▶ Could be used to order work by an estimate of time taken to complete - very interesting, but not implemented.
- ▶ `execution_context` checks on the argument-type of the `process(...)` function at compile-time,
 - ▶ also if it is `const`, i.e. pure, which could be used to implement memoizing - also very interesting, but not implemented.

PPD: Example using `for_each()`.

Listing 3: For Each with a Thread Pool and Future.

```
typedef ppd::safe_colln<
    vector<int>, lock_traits::recursive_critical_section_lock_type
> vtr_colln_t;
typedef pool_type::void_exec_cxt execution_context;
struct accumulate {
    static int last; // for_each needs some kind of side-effect...
    void operator()(int t) { last+=t; };
};

vtr_colln_t v;
v.push_back(1); v.push_back(2);
execution_context context(pool<<joinable()>>pool.for_each(v, accumulate()));
*context;
```

- ▶ The `safe_colln` adaptor provides the lock for the collection,
 - ▶ in this case a simple EREW lock,
 - ▶ only locks the collection, not the elements in it.
- ▶ The `for_each()` returns an `execution_context`:
 - ▶ The input collection has a read-lock taken on it.
 - ▶ Released when all of the asynchronous applications of the functor have completed, when the read-lock has been dropped.
 - ▶ It makes no sense to return a random copy of the input functor.

PPD: Example using transform().

Listing 4: Transform with a Thread Pool and Future.

```
typedef ppd::safe_colln<
    vector<long>, lock::rw<lock_traits>::decaying_write,
    no_signalling, lock::rw<lock_traits>::read
> vtr_out_colln_t;

vtr_colln_t v; vtr_out_colln_t v_out;
v.push_back(1); v.push_back(2);
pool_type::void_exec_cxt context(
    pool<<joinable()><<pool.transform(
        v, v_out, std::negate<vtr_colln_t::value_type>()
    )
);
*context;
```

- ▶ `vtr_out_colln_t` provides a CREW lock on `v_out`: it is re-sized, then the write-lock atomically decays to a read-lock.
 - ▶ Note how the collection is locked, but not the data elements themselves.
- ▶ The `transform()` returns an `execution_context`:
 - ▶ Released when all of the asynchronous applications of the unary operation have completed & the read-locks have been dropped.
 - ▶ Similarly, it makes no sense to return a random iterator.

PPD: Example using accumulate().

Listing 5: Accumulate with a Thread Pool and Future.

```
typedef pool_type::accumulate_t<
    vtr_colln_t
>::execution_context execution_context;

vtr_colln_t v;
v.push_back(1); v.push_back(2);
execution_context context(
    pool.accumulate(v,1, std::plus<vtr_colln_t::value_type>())
);
*context;
```

- ▶ The `accumulate()` returns an `execution_context`:
 - ▶ Released when all of the asynchronous applications of the binary operation have completed & the read-lock is dropped.
- ▶ Note the use of the `accumulate_t` type: it contains a specialized counter that atomically accrues the result using suitable locking according to the API and type.
 - ▶ This is effectively a map-reduce operation.
 - ▶ `find()`, `find_if()`, `count()` and `count_if()` are also implemented, but omitted, for brevity.

PPD: Example using `logical_or()`.

Listing 6: Logical Or with a Thread Pool and a Future.

```
struct bool_work_type {  
    void process(bool &) const {}  
};  
pool_type pool(2);  
pool_type::bool_exec_cxt context(  
    pool.logical_or(bool_work_type(), bool_work_type())  
);  
*context;
```

- ▶ The `logical_or()` returns an `execution_context`:
 - ▶ Each argument is asynchronously computed, independent of each other:
 - ▶ the compiler attempts to verify that they are pure, with no side-effects (they could be memoized),
 - ▶ no short-circuiting, so breaks the C++ Standard.
 - ▶ Once the results are available, the `execution_context` is released.
- ▶ `logical_and()` has been equivalently implemented.

PPD: Example using boost::bind.

Listing 7: Boost Bind with a Thread Pool and a Future.

```
struct res_t {
    int i;
};
struct work_type {
    res_t exec() {return res_t();}
};
pool_type pool(2);
execution_context context(
    pool<<joinable()<<wrap(boost::bind(&work_type::exec, work_type()))
);
context->i;
```

- ▶ The `wrap(boost::bind())` returns an `execution_context`:
 - ▶ There should be no unbound arguments in the call to `boost::bind()`, no placeholders.
 - ▶ The function passed should return a copy-constructible type, which is the result.
 - ▶ Contrast with the reference argument result-type in the earlier example.
 - ▶ The only cost is the extra stack-space required to store the function-pointer & bound arguments.
 - ▶ Once the results are available, the `execution_context` is released.

PPD: Some Thoughts Arising from the Examples.

- ▶ The read-locks on the collections assist the user in avoiding operations that might invalidate iterators on those collections, when they should use EREW locks.
 - ▶ Internally PPD creates tasks that recursively divide the collections into p equal-length sub-sequences, where p is the number of threads in the pool, at most $O(\log(\min(p, n)))$.
- ▶ The operations on the elements should be thread-safe & must not affect the validity of iterators on those collections.
- ▶ The user may call many STL-style algorithms in succession, each one placing work into the thread pool to be operated upon, and later recall the results or note the completion of those operations via the returned `execution_contexts`.
 - ▶ This allows the user to fill the thread pool with work and maintain high throughput.
 - ▶ But this implies that the work that is done on each iterator is more costly than the synchronization and thread-management operations that are incurred. This is unlikely for such simple algorithms as `count()` and `copy()`.

PPD: Sequential Operation.

Recall the `thread_pool` typedef that contained the API and threading model:

Listing 8: Definitions of the typedef for the thread pool.

```
typedef ppd::thread_pool<
    pool_traits::worker_threads_get_work, pool_traits::fixed_size,
    generic_traits::joinable, platform_api, heavyweight_threading,
    pool_traits::normal_fifo, std::less
> pool_type;
```

To make the `thread_pool` and therefore all operations on it sequential, and remove all threading operations, one merely needs to replace the trait `heavyweight_threading` with the trait `sequential`.

- ▶ It is that simple, no other changes need be made to the client code, so all threading-related bugs have been removed.
- ▶ This makes debugging significantly easier by separating the business logic from the threading algorithms.

PPD: Other Features.

Other parallel algorithms implemented:

- ▶ `unary_fun()` and `binary_fun()` which would be used to execute the operands in parallel, in a similar manner to the similarly-named STL binary functions.
 - ▶ Examples omitted for brevity.

Able to provide estimates of:

- ▶ The optimal number of threads required to complete the current work-load.
- ▶ The minimum time to complete the current work-load given the current number of threads.

This could provide:

- ▶ An indication of soft-real-time performance.
 - ▶ Furthermore this could be used to implement dynamic tuning of the thread pools.

PPD: Limitations.

1. PPD assumes an SMP machine-model in which the inter-thread communication time is only roughly modeled via types of thread pool.
 - ▶ The implementation of all of the various types of thread pool are not complete.
2. The test cases are reasonable, uses boost.test:
 - ▶ They have reasonable coverage, but needs much more sophisticated test cases - library may not be bug-free.
3. The missing STL algorithms:
 - ▶ The lack of `merge()`, `sort()`, `replace()` etc., algorithms means that PPD misses covering some 5% of instances of usage of these interesting algorithms in the example code-base.
4. Linking to a good-quality multi-threaded memory manager would be useful, such as [Hoard, 8], but this is optional.
5. It breaks the C++ ISO Standard in certain subtle ways, for example the logical operations do not short-circuit.

PPD: Cross-Platform Issues.

Not actually cross platform!

- ▶ Different qualities of C++ compilers (G++ vs MSVC++) conspire to defeat this ability (compiler bugs, repeated re-compilation against alternate platforms, etc, etc).
 - ▶ Always a moving target, upgrading sometimes means regressions in the compilers.
 - ▶ This has been a major stumbling block, requiring almost continual retesting.
 - ▶ I should not have attempted to design the library with a large code-base of cross-platform code, and a thin layer to normalise the APIs. I should have reimplemented the library for each API.
 - ▶ This is a quality of implementation issue relating to the compilers, not an issue relating to the C++ language.

Conclusions.

Recollecting my original problem statement, what has been achieved?

- ▶ A library that assists with creating deadlock and race-condition free schedules.
- ▶ Given the typedefs, a procedural way automatically to convert STL algorithms into optimal threaded ones.
- ▶ The ability to add to the business logic parallel algorithms that are sufficiently separated to allow relatively simple debugging and testing of the business logic independently of the parallel algorithms.
- ▶ All this in a traditionally hard-to-thread imperative language, such as C++, indicating that it is possible to re-implement such an example library in many other programming languages.

These are, in my opinion, rather useful features of the library!

Future Directions.

Given the issues and limitations of PPD there are further areas that it would be useful to investigate:

- ▶ Most of the library is concerned with CREW-PRAM, but what about EREW-PRAM constraints?
 - ▶ Literature indicates that the algorithmic complexity is increased by $O(\log(n))$.
- ▶ GSS(k), or baker's scheduling, would be useful to reduce the synchronisation costs on thread-limited systems or pools in which the synchronisation costs are high.
- ▶ A performance analysis of using PPD vs. other threading models using the [SPEC2006, 9] benchmark suite has just been started.
- ▶ Some investigation of how PPD might support asynchronous I/O threading operations.
- ▶ How does or even should PPD interact with:
 - ▶ The C++ 0x standardisation efforts regarding threading.
 - ▶ The excellent work in Boost.Threads?

References 1



S. B. Baden, R. S. Schriebe, K. S. Gatlin, and S. J. Fink, "A preliminary evaluation of hpf," in *In Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing*, 1997.



T. A. El-ghazawi, W. W. Carlson, and J. M. Draper, "Upc language specifications v1.1.1," tech. rep., 2003.



C. Pheatt, "Intel®threading building blocks," *J. Comput. Small Coll.*, vol. 23, no. 4, pp. 298–298, 2008.



D. Abrahams, "Boost." www.boost.org, July 2009.



X. Tang, *Compiling for Multithreaded Architectures*.
PhD thesis, University of Delaware, Delaware, USA, Fall 1999.



A. Gibbons and W. Rytter, *Efficient parallel algorithms*.
New York, NY, USA: Cambridge University Press, 1988.



H. Casanova, A. Legrand, and Y. Robert, *Parallel Algorithms*.
Chapman & Hall/CRC Press, 2008.



E. Berger, K. McKinley, R. Blumofe, and P. Wilson, "Hoard: A scalable memory allocator for multithreaded applications," in *ASPLOS-IX: Ninth International Conference on Architectural Support for Programming Languages and Operating Systems*, (Cambridge, MA), pp. 117–128, 2000.



J. Reilly, "Evolve or die: Making spec's cpu suite relevant today and tomorrow," in *IISWC*, p. 119, 2006.

References 2



<http://www.specbench.org/cpu2006/Docs/>



<http://www.hoard.org/>



<http://accu.org/index.php/conferences/2004/speakers2004>



<http://threadingbuildingblocks.org/>



http://www.boost.org/doc/libs/1_45_0/doc/html/thread.html



<http://www.stdthread.co.uk/>



<http://software.intel.com/en-us/articles/intel-concurrent-collections-for-cc/>



<http://www.picochip.com/>



<http://www.capsl.udel.edu/~venetis/Cyclops-64.html>