

# Extended abstract: Type oriented programming for task based parallelism

Nick Brown  
EPCC, James Clerk Maxwell  
Building, Peter Guthrie Tait Road,  
Edinburgh  
nick.brown@ed.ac.uk

Ludovic Capelli  
EPCC, James Clerk Maxwell  
Building, Peter Guthrie Tait Road,  
Edinburgh

James Mark Bull  
EPCC, James Clerk Maxwell  
Building, Peter Guthrie Tait Road,  
Edinburgh

## 1 Challenge

Writing parallel codes is difficult [1] and exhibits a fundamental trade-off between abstraction and performance. The high level language abstractions designed to simplify the complexities of parallelism make certain assumptions that impacts performance and scalability. On the other hand lower level languages, providing many opportunities for optimisation, require in-depth knowledge and the programmer to consider tricky details of parallelism. An approach is required which can bridge the gap and provide both the ease of programming and opportunities for control and optimisation.

## 2 Type oriented programming

By optionally decorating their codes with additional type information, programmers can either direct the compiler to make certain decisions or rely on sensible default choices. In [2] we introduced a research based programming language, Mesham, which explores these ideas in relation to data parallelism. This is illustrated in listing 1 with three declarations of variables.

```
1 var a: Int ;  
2 var b: Int :: allocated [ single [ on [ 0 ] ] ] ;  
3 var c: Int :: allocated [ single [ on [ 0 ] ] ] ::  
   channel [ 0 , 1 ] ;
```

**Listing 1.** Type examples in the Mesham language

The first declaration determines that  $a$  is an integer and, in the absence of further type information, the compiler generates code that allocates this variable on every process. The behaviour of an assignment  $a:=22$  is to perform a local assignment on every process. In the second declaration we have added extra type information to further direct the compiler, in this case explicitly specifying that  $b$  will be allocated only on process 0. Based upon this additional information the same assignment  $b:=22$  generates a local assignment on process 0 and remote communications on every other process. By default the behaviour of these communications is simple and safe, however might not be particularly performant. In declaration three the programmer has added extra type information to guide the compiler to handle remote data access as a point to point communication (a

channel) rather than the default RMA. In our approach the programmer, after experimentation and profiling, can further directed parallelism in a high level manner without having to consider the low level implementation details or significantly rewrite their code. The use of types in listing 1 is different from simply specifying variable storage because, for instance, the assignment  $a:=b$  invokes a broadcast from process 0 of the value held in  $b$  to all other processes, writing this into their local  $a$ .

We denote this combination of types, which then determines the behaviour of all variable usage, using the  $::$  operator. This is known as a type chain. Precedence is from right to left, so certain types can override the behaviour of other types based upon their order in the chain. Types can be arbitrarily chained together and any potentially conflicting combinations are handled by this precedence rule. The specific types themselves are separate from the core semantics of the language which means that, from a language perspective, it is trivial to add or remove types for specific domains. This approach has been applied to areas from traditional HPC [3] to graph based codes [4].

Our use of types is different from annotation approaches, such as OpenMP [5], because types are part of, and integrate fully with, the language rather than a bolt on. Therefore the programmer has flexibility to create new types in their code and reasoning about type information using existing language constructs. Through constructing type chains we provide a mechanism for building up complex type information in a structured, hierarchical, manner and it is this type chain that provides the behaviour of operations performed on the variable throughout its life.

## 3 Task based parallelism

Many traditional HPC codes have been oriented around data parallelism, where a data is split up and distributed amongst processes. However associated techniques such as halo swapping often result in a bulk synchronous style of parallelism, where processes proceed in computation and then communication/synchronisation steps. This synchronisation, which is often global, is inefficient and to reach the scale of millions of processes must be avoided.

In task based parallelism computational tasks themselves are decomposed amongst the processes. Typically tasks are scheduled based upon a number of dependencies and will execute once these dependencies are met. Task based parallelism forces the programmer to break away from their bulk synchronous approach and promotes the asynchronous nature of codes. Whilst there is active research into task based programming models, the level of abstraction is again a major challenge. Many existing technologies rely on the runtime to make sensible decisions, such as scheduling. This can have a significant impact on performance without the user being able to explicitly control or tune important parameters.

## 4 Types for task based parallelism

Our use of types so far followed the data parallelism approach but we believe that, by decorating functions, types can also address the abstraction challenge of task based parallelism. In Mesham the declaration of a function, *myFunction* which takes in and returns an integer follows *function Int myFunction(var a:Int)*. When called, the behaviour is to execute the function immediately and return the value. It is possible to apply addition type information to the function declaration. The addition of the *spawnable* type such as *function Int myFunction(var a:Int):spawnable* overrides this default behaviour. Because of the additional type, calls to the function will schedule it for execution on a thread rather than executing it directly. The *spawnable* type effectively transforms functions into tasks, one per thread run concurrently.

An important aspect of scheduling these tasks is to have some way of referencing them. The semantics of the *spawnable* type is that function calls will return a variable of the type *Future[X]* (where *X* is the actual return type, in this case *Future[Int]*.) This future can be used as a handle to test and wait for completion. Listing 2 illustrates a simple Fibonacci example where the *fib* function is marked as *spawnable* (a concurrent task.) Based upon the function calls at lines 4 and 5, upon execution the task schedules two further recursive *fib* tasks and the calling of these functions immediately return futures to these as the variables *f1* and *f2*. Lines 6 and 7 synchronise on the futures (waits for their corresponding tasks to complete), before adding the integer values together and returning the result.

---

```

1 function Int fib(var val:Int) : spawnable {
2   if (val == 0 || val == 1) return val;
3   var f1, f2 : Future[Int];
4   f1 := fib (val - 1);
5   f2 := fib (val - 2);
6   synchronise (f1);
7   synchronise (f2);
8   return f1.val + f2.val;
9 }
```

---

**Listing 2.** Fibonacci task parallelism with explicit synchronisation on futures

However the code in listing 2 is naive as the explicit synchronisations block the calling thread which is wasteful. To avoid this we allow the programmer, via type information, to encode the dependencies of tasks within their code. Now the scheduler will not execute the task until these dependencies are met and hence there is no explicit synchronisation or blocking of threads. Listing 3 is the same Fibonacci task parallel code but using task dependencies, via the *dependencies* type, on the *add* function instead of explicit synchronisation. The behaviour of the *dependencies* type is that the decorated function will accept both normal valued arguments (in this case integers for variables *a* and *b*) as well as futures. If futures are provided (as is the case in listing 3) then the scheduler will wait until the tasks that they depend upon have completed before executing the scheduled, decorated, function. In this manner execution of the *fib* task returns a future on the *add* task which itself is dependent on futures of recursive calls to the *fib* task.

---

```

1 function Int fib(var val:Int) : spawnable {
2   if (val == 0 || val == 1) return val;
3   var f1, f2 : Future[Int];
4   f1 := fib (val - 1);
5   f2 := fib (val - 2);
6   return add(f1, f2);
7 }
8
9 function Int add(var a:Int, var b:Int) :
10  spawnable :: dependencies {
11  return a + b;
12 }
```

---

**Listing 3.** Fibonacci task parallelism with dependencies

It is possible to omit the *spawnable* type whilst keeping the *dependencies* type. In this case the same dependencies behaviour is present, with function execution being immediate and blocking rather than a task. This is how we implement the *synchronise* call of listing 2.

The idea of programmers decorating functions with type information in order to guide the compiler to generate the correct code for task based parallelism is of main interest here. In many cases existing sequential functions can be decorated and, with minimal modifications to the code, programmers can direct how these functions will execute and any dependencies at a high level. The types themselves, whilst interesting in their own right, are mainly used by us as a vehicle for illustrating the benefits of driving parallelism through types. We are developing additional types to control task placement, runtime scheduling priorities and resilience.

## References

- [1] Skillicorn, David B., and Domenico Talia. Models and languages for parallel computation. *Acm Computing Surveys (Csur)* 30.2 (1998): 123-169.
- [2] Brown, Nick. Applying Type Oriented Programming to the PGAS Memory Model. *7th International Conference on PGAS Programming Models* 2013.
- [3] Brown, Nick. Type oriented parallel programming for Exascale *Advances in Engineering Software*, Available online 26 April 2017, to appear in print.
- [4] Brown, Nick. A Type-Oriented Graph500 Benchmark *Proceedings of the 29th International Conference on Supercomputing - Volume 8488*, 460-469
- [5] OpenMP Architecture Review Board: OpenMP Application Program Interface Version 3.0 <http://www.openmp.org/mp-documents/spec30.pdf> , 2008