

Integrating Transactions into the Data-Driven Multi-threading Model using the TFlux Platform

Andreas Diavastos, Pedro Trancoso

Dept. of Computer Science

University of Cyprus

Emails: {diavastos, pedro}@cs.ucy.ac.cy

Mikel Luján and Ian Watson

School of Computer Science

University of Manchester

Emails: {mikel, watson}@cs.manchester.ac.uk

Abstract—Multi-core processors have renewed interest in programming models which can efficiently exploit general purpose parallelism. *Data-Flow* is one such model which has demonstrated significant potential in the past. However, it is generally associated with functional styles of programming which do not deal well with shared mutable state. There have been a number of attempts to introduce state into Data-Flow models and functional languages but none have proved able to maintain the simplicity and efficiency of pure Data-Flow parallelism. Transactional memory is a concurrency control mechanism that simplifies sharing data when developing parallel applications while at the same time promises to deliver affordable performance. In this paper we report our experience of integrating Transactional Memory and Data-Flow. The ability of the Data-Flow model to expose large amounts of parallelism is maintained while Transactional Memory provides simplified sharing of mutable data in those circumstances where it is important to the expression of the program. The isolation property of transactions ensures that the exploitation of Data-Flow parallelism is not compromised.

In this study we extend the TFlux platform, a Data-Driven Multi-threading implementation, to support transactions. We achieve this by proposing new pragmas that allow the programmer to specify transactions. In addition we extend the runtime functionality by integrating a software transactional memory library with TFlux. To test the proposed system, we ported two applications that require transactional memory: Random Counter and Labyrinth an implementation of Lee's parallel routing algorithm. Our results show good opportunities for scaling when using the integration of the two models.

Keywords—Programming models, Future Multi-cores, Transactional Memory, Data-Flow model, Data-Driven Multi-threading

I. INTRODUCTION

As technology delivers higher integration of devices into processors, the multi-core design has become the de-facto standard for processor architecture. It promises to deliver high performance whilst maintaining an acceptable complexity and power budget. The trends show a continuous increase in the number of cores and it is expected that by 2020 processors will include 1000s of cores [1]. This will lead to new challenges, one of them being the programmability of such large-scale systems. If their power is to be harnessed on the solution of a wide range of problems it will be necessary

to develop new parallel programming models which are both efficient and easy to use.

There has recently been a re-surgence of interest in the Data-Flow model as a way to efficiently exploit large-scale parallelism. Even though the original implementations of Data-Flow were not efficient, more recent developments have overcome this [6], [9], [21]. However, the models are suited largely to the implementation of programming styles which are essentially purely functional. Indeed it is the absence of side effects in functional models which permits easy parallelisation. Unfortunately, there are many cases in real programs where the use of shared mutable state is either necessary for efficiency or is a fundamental part of the problem being solved. In these circumstances, functional approaches are unsuitable.

This limitation has long been recognised and there have been a number of attempts to integrate state into both Data-Flow models and functional languages. The early work on M-structures in Id [4] added implicit locking to data items to avoid the explicit manipulation of synchronisation. However, this merely hid the complexity of the synchronisation rather than removed it and, in many ways, made the writing of shared state programs more error prone. Functional languages such as SML [15] and F# [23] have introduced mutable variables in an attempt to extend their practicality. Unfortunately this state can rapidly destroy both the mathematical cleanliness of the language and the ability to exploit parallelism with Data-Flow like execution models. Haskell originally introduced state in a more disciplined way by the use of MONADS [14]. However, although this enables the isolation of state via the type system and hence preserves mathematical properties, the state manipulation is serialised and thus does not address the problems of writing parallel programs.

The Transactional Memory (TM) model facilitates sharing data in a manner which isolates individual sharers from the complexities of synchronisation. It was originally proposed as a way of simplifying parallel programming in conventional languages but has been shown to provide a clean and simple way to add the sharing of state to a functional language [19]. Transactional Haskell uses the MONAD

approach to allow the expression of explicit threads within Haskell programs by defining transactional variables which are manipulated serially within a thread but interact in parallel across threads. It has been shown that parallel state based programs can be specified while maintaining much of the purity of functional programming. It is our belief that a more general and more usable programming model can be produced by adding transactions to Data-Flow using more pragmatic programming approaches and avoiding the complexity of MONADS.

This paper reports our first experiences of integrating transactions into a thread based Data-Flow model. We use TFlux [22] as the Data Driven Multi-threading (DDM) implementation and TinySTM [18] for transactional support. This merged implementation will be referred to as DDM+TM, hereafter. We propose new TFlux directives for defining transactional threads and variables. We also show how it is possible to program applications with the combined model and present a preliminary performance evaluation study.

This paper is organized as follows, Section II describes the concepts, the motivation and the issues of combining the two models. Section III describes the implementation we propose using the TFlux system and TinySTM model. In Section IV we analyze the applications we implemented and used for our experimental evaluation. Section V and Section VI describe the experimental setup and the performance scaling of two applications developed using DDM+TM. In Section VII we discuss the conclusions and future work.

II. COMBINING DATA-FLOW WITH TRANSACTIONS

There has recently been renewed interest in the Data-Flow approach to computation that was pioneered in late 70s and 1980s [10], [17], [5], [27], [13]. Those projects demonstrated that it was feasible to express sufficiently large amounts of parallelism that a significant problem was how to throttle and schedule it. Subsequent projects, for example TFlux [22] showed that a coarsening of the granularity (from instructions to tasks) could result in more controllable and more efficient parallel execution.

Numerical Linear Algebra (NLA) is one area where Data-Flow ideas have recently been adopted. This is apparent both in LAPACK and BLAS functionality (see, *e.g.* PLASMA project [24]) and for sparse matrices [12]. NLA constitutes one of the main kernels for scientific computing and their main next challenge is how to scale to petaflop-scale high-performance systems. The new generation of NLA algorithms are moving towards expressing parallelism but leaving the scheduling to the runtime trying to harness the available combination of resources (multi-cores, many-cores, clusters, GPUs). It has been demonstrated that a Data-Flow execution using Parallel Linear Algebra for Scalable Multi-core Architectures (PLASMA) can easily generate millions of tasks.

There is a clear relation between Data-Flow computations and parallelization of functional languages. Another highly prominent use of functional programming techniques can be observed in the MapReduce frameworks [7]. Provided the map and reduce operations are side-effect free, we can automatically parallelize their execution using a Data-Flow approach.

Although the benefits of Data-Flow can be demonstrated by the above applications, it is clear that there are cases where shared state is essential. This may be because it is a fundamental part of the program, it facilitates software development or a pure functional execution will be inefficient (*e.g.* due to memory management overheads). To overcome these limitations, it is desirable to introduce shared mutable state into the Data-Flow approach. However, this needs to be done in a way which does not require the specification of explicit synchronisation between parts of the program. This both introduces significant programming complexity and can often lead to unnecessary serialisation of the execution. Therefore approaches which introduce conventional locking are unlikely to lead to models which are either easy to use or efficient.

Transactional Memory (TM) [11] is a model for manipulating mutable shared data which attempts to reduce complexity by eliminating the need for explicit synchronisation. It works by allowing the programmer to specify that certain sections of a program must be executed atomically but without the need to consider any of the synchronised control that might be required. Execution of atomic sections takes place optimistically, that is with an assumption that any shared data within the section will not be changed by any other concurrent execution. If such a conflict does occur, the underlying runtime system ensures that only one execution succeeds while others are transparently re-executed. This leads to the important property of isolation. A thread always proceeds as though it has exclusive access to any shared data within an atomic section. All synchronisation complexity is removed and it is unnecessary to serialise accesses to achieve correct execution. Although, in practice, some serialisation may occur due to the resolution of conflicts, the optimistic nature of the model ensures that maximal parallelism is achieved. Although TM was originally proposed to reduce the complexity in the context of conventional threaded programming languages, the isolation property makes it an ideal way of introducing mutable state into Data-Flow or functional approaches.

A common example used for motivating TM illustrates the need for mutable shared state [19]. Consider a computation which is trying to perform concurrent credit/debits between bank accounts. Firstly, the state is fundamental to the problem. The account balances must be globally accessible variables which can be updated and persist. The credit/debit operations must be atomic to preserve the overall correctness of the balances. Assuming that we don't know the identity

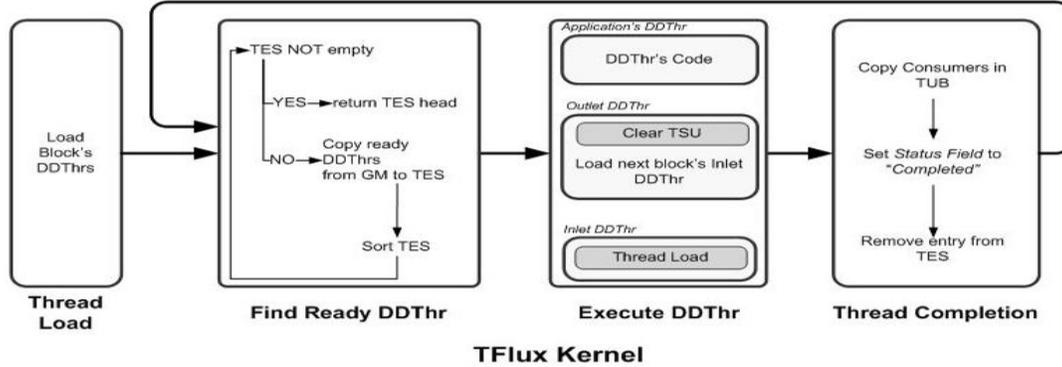


Figure 1. TFlux Runtime system.

of the accounts when specifying the problem, it is clear that the accounts might overlap and conflicts could occur. A conventional locking approach would need to deal with the cases where overlap might occur by taking explicit locks and dealing with complex interactions such as deadlock. The problem could, of course, be greatly simplified by serialising all the operations but this would defeat the desire to exploit parallelism. However, in many cases, there will be no overlap and an optimistic approach can proceed with maximal parallelism. We can envisage a Data-Flow solution where threads have been generated to perform calculations on each account using a purely functional approach and then invoking a transaction to perform a balance transfer. Any number of such threads can be generated to operate in parallel without any need to consider how they interact.

In the TERAFLUX project [1] we are investigating how to combine different variants of Data-Flow models (including synchronous Data-Flow) with Transactional memory. This paper reports our experiences with a first combination of the data-driven runtime system, TFlux, and software TM driven by the implementation of two applications.

III. PROPOSED DDM+TM IMPLEMENTATION

A. TFlux System

This study uses the Data-Driven Multi-threading model and in particular its TFlux implementation [22]. TFlux is a complete system that defines a set of ‘#pragma’ directives to program a DDM application. TFlux includes a source-to-source pre-processor that automatically adds calls to the TFlux library to implement the threads and their scheduling and a runtime system that implements the Thread Scheduling Unit (TSU). The TSU is the unit that handles the scheduling of threads in a Data-Flow manner. In Figure 2 we depict the different modules of TFlux.

In order to program a DDM application with TFlux, directives must be added to regular C code. The most relevant directives are the ones which enable a set of instructions to be defined as a thread (see Table I). In addition it is necessary

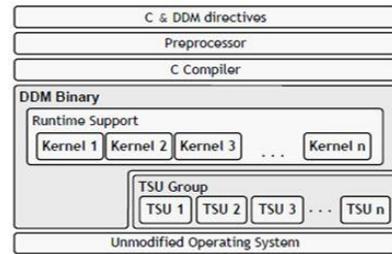


Figure 2. The layered design of the TFlux system [22].

to define the inputs and outputs of a thread as well as the producer and consumer relationships between threads. Using this information the system is able to form the code for the threads as well as the thread dependency graph, which is the structure that needs to be loaded into the TSU for the scheduling to be executed in a Data-Flow manner. The task of the scheduling unit is to manage the counters that control the firing of threads. Each time a producer thread terminates its execution, the consumer’s counter is decremented by one. When the counter reaches zero, all needed results have been produced and thus the thread is ready for execution. All these operations are part of the TFlux runtime system, which includes all data structures required to manage the thread scheduling as well as the scheduling code itself. The operations of the runtime system are depicted in Figure 1. There are hardware and software versions of the TFlux platform. In this study we use the software version, also called TFluxSoft, without losing generality of the system. In TFluxSoft the TSU’s execution is handled by one core of the multi-core system we are using.

B. TinySTM Software Library

In order to support the transactional execution, *i.e.* the monitoring of the updates to variables, the conflict detection and the restarting of the execution in case of abort, we need to extend the TFlux runtime. Rather than developing

Table I
TFLUX DDM PRAGMA DIRECTIVES.

#pragma ddm startprogram	Define the start and the end of a DDM program
#pragma ddm endprogram	
#pragma ddm block <i>ID</i>	Define the start and the end of a block of threads with identifier <i>ID</i>
#pragma ddm endblock	
#pragma ddm thread <i>ID</i> kernel <i>NUMBER</i>	Define the boundaries of a DDM thread with identifier <i>ID</i> and the kernel <i>NUMBER</i> to execute on
#pragma ddm endthread	
#pragma ddm for thread <i>ID</i>	Define the boundaries of a DDM loop thread with identifier <i>ID</i>
#pragma ddm endfor	
#pragma ddm kernel <i>NUMBER</i>	Declare the number of kernels to be used
#pragma ddm var <i>TYPE NAME</i>	Declare a shared variable with <i>NAME</i> and <i>TYPE</i>
#pragma ddm private var <i>TYPE NAME</i>	Declare a private variable with <i>NAME</i> and <i>TYPE</i>

from scratch we preferred to extend the TFlux runtime system with an existing software TM implementation. We chose the TinySTM [18] as it appeared to provide a simple approach to the integration. However, we could have used other Software TM systems such as TL2 [8] or RSTM [20]. Within TERAFLUX, we intend to investigate other TM implementations, including hardware support, for the scalability of the execution required for large-scale Data-Flow applications.

C. DDM+TM

When adding support for transactions to TFlux an important decision concerned the granularity of transactions. The simplest approach would be to declare a whole thread as a transaction. With this option we would enhance the system by providing the programmer with two types of threads: pure Data-Flow threads or transactional threads. However, a thread may contain code that needs to be transactional but combined with non-transactional code. Furthermore, it may be appropriate to specify several atomic regions within a thread. This could lead to potentially wasteful aborts when either a transaction is only a small portion of the thread or multiple atomic regions need to be aborted together. Therefore, we opted for providing new pragma directives to define the beginning and end of transactional sections within threads. These new TFlux directives are presented in Table II.

Another design issue is how we identify variables which are transactional. These variables will require that their read and write operations are observed to form the read-set and write-set during a speculative execution of the transaction. These sets are used to detect conflicts. For all these transactional variables we also need to version the results to allow a clean restart of the transaction if necessary. One option is to monitor every memory access that is performed within a transaction. However, this is not necessary for unshared variables, for example those which are thread local. Therefore, in common with other TM approaches, we explicitly declare which variables are transactional. The directive

```
#pragma ddm atomic tvar(NAME : READ/WRITE)
```

offers such functionality. Note that each transactional variable is associated within a thread with a READ, WRITE or READ_WRITE qualifier. This qualifier provides information on the use of the variable within the thread which can be used by the TM implementation to optimize the execution.

For DDM+TM, we decided to have a complete separation between transactional and non-transactional variables. Transactional variables must always be accessed within a transaction. Non-transactional variables are normally private to a thread during execution and thus cannot generate conflicts. Other non-transactional threads will only be allowed to access a non-transactional variable if the scheduling can guarantee independence. With this decision we avoid the problems of weak isolation. Note that we do not modify DDM by imposing this decision. DDM+TM could be implemented without speculation by performing a scheduling where the transactional variables are treated as inputs and outputs of the threads that are read and written. This will result in the sequential execution of the code in case of threads accessing shared data where we cannot determine the dependencies before runtime.

For transactional variables, we also provide the programmer with pragmas to define them within the declaration of a transactional thread. This is required to support the monitoring of variables that may have more than one alias (e.g. parameter variables inside the code of a function). The monitoring of these variables is specified as a parameter in the thread declaration (see Table II).

As TFlux has two pragmas for declaring threads, the table contains

```
#pragma ddm atomic thread ID and
#pragma ddm atomic for thread ID
```

declaring a transactional thread and a transactional loop thread, respectively. The tvar(NAME : READ/WRITE) extension defines the thread variables that are transactional.

The last proposed directive

```
#pragma ddm atomic transaction
```

allows the declaration of a transaction as a portion of a thread. For certain applications such as those considered in this work, this offers better performance (see Section VI).

These directives are a subset of the possible ones which

Table II
TFLUX DDM+TM PRAGMA DIRECTIVES.

#pragma ddm atomic thread <i>ID</i> tvar(<i>NAME</i> : <i>READ/WRITE/READ_WRITE</i>)	DDM+TM thread boundaries with identifier <i>ID</i> and the atomic variables to monitor for either <i>READ</i> or <i>WRITE</i>
#pragma ddm atomic endthread	
#pragma ddm atomic for thread <i>ID</i> tvar(<i>NAME</i> : <i>READ/WRITE/READ_WRITE</i>)	DDM+TM loop thread boundaries with identifier <i>ID</i> and the atomic variables to monitor for either <i>READ</i> or <i>WRITE</i>
#pragma ddm atomic endfor	
#pragma ddm atomic transaction tvar(<i>NAME</i> : <i>READ/WRITE/READ_WRITE</i>)	DDM+TM boundaries of a transaction that is smaller than a thread and the atomic variables to monitor for either <i>READ</i> or <i>WRITE</i>
#pragma ddm atomic endtransaction	
#pragma ddm atomic tvar(<i>NAME</i> : <i>READ/WRITE/READ_WRITE</i>)	Declare an atomic variable to monitor either for <i>READ</i> or <i>WRITE</i>
#pragma ddm atomic abort	Manually abort a transaction

Table III
TFLUX DDM+TM PRAGMA DIRECTIVES CORRESPONDENCE WITH TINYSTM RUNTIME CALLS.

DDM+TM Pragma Directives	TinySTM runtime support
#pragma ddm atomic thread <i>ID</i> tvar(<i>NAME</i> : <i>READ/WRITE/READ_WRITE</i>)	stm_init_thread()
#pragma ddm atomic endthread	stm_exit_thread()
#pragma ddm atomic for thread <i>ID</i> tvar(<i>NAME</i> : <i>READ/WRITE/READ_WRITE</i>)	stm_init_thread()
#pragma ddm atomic endfor	stm_exit_thread()
#pragma ddm atomic transaction tvar(<i>NAME</i> : <i>READ/WRITE/READ_WRITE</i>)	sigjmp_buf *_e = stm_start(&_a) (*_e != NULL) sigsetjmp(*_e, 0)
#pragma ddm atomic endtransaction	stm_commit()
#pragma ddm atomic tvar(<i>NAME</i> : <i>READ/WRITE/READ_WRITE</i>)	int temp = (int) stm_load((stm_word_t *)&NAME) OR stm_store((stm_word_t *)&NAME, (stm_word_t *)temp)
#pragma ddm atomic abort	stm_abort()

have been defined for TM [11]. However, they are enough to implement the applications described in this paper and we consider them to be the core directives. Extra transactional functionality can be added by declaring

```
#pragma ddm atomic abort
```

in the case the programmer wants to manually abort a transaction. In Table III we show the correspondence of the proposed pragmas with the calls to the TinySTM runtime that will automatically be generated by the preprocessor.

Finally we would like to mention the issue of scheduling. The TFlux runtime system takes care of scheduling the threads. The software TM also takes decisions about whether to abort and re-execute a transaction. Previous work on TM [3], [2] has shown that controlling the scheduling using information about transactions can improve the performance and reduce wasted work due to aborts. In this first attempt at adding TM to TFlux, we have opted not to make changes in the TSU to support transactional behavior. However, we are investigating the possibility of offloading the re-scheduling of an aborted transactional thread to the TSU instead of the TinySTM system.

IV. WORKLOADS

For this proposal of transaction integration with the Data-Flow model we have tested two applications. The Random Counters and the Labyrinth implementation of Lee's algorithm from the STAMP Benchmark suite [16]. These

were originally implemented in a conventional language using TM and are not naturally Data-Flow applications. However, for this study we are focussing mainly on the functionalities and overheads of the implementation rather than the added benefits of the integration such as exploiting implicit parallelism.

A. Random Counters

In the Random Counters application we use multiple threads to increment the values of random array positions. We create and initialize all elements of the array to zero. We then spawn multiple threads to execute loop thread code. Each loop iteration is executed in parallel and each thread will do the same work: generate a random sequence of numbers - that represent index positions in the array and increment by 1 the values located in those positions. In the code shown in Figure 3 we present the implementation of this algorithm, where the

```
#pragma ddm atomic for and
#pragma ddm atomic endfor
```

directives define the DDM+TM thread boundaries, showing that all the code of the thread is considered to be transactional. The purpose of this algorithm is to create concurrent accesses to memory. This will in turn create memory conflicts between threads trying to access the same memory locations. In an unsynchronised parallel implementation this execution would create a false result. In the DDM model

```

#pragma ddm atomic for thread 1
    tvar(counters : READ_WRITE)
for (cv00 = 0; cv00 < THREADS; cv00++)
{
    for (j = 0; j < numOfCounters; j++)
        indexTM[j] = rand() % arraySize;
    for (j = 0; j < numOfCounters; j++)
        counters[indexTM[j]]++;
}
#pragma ddm atomic endfor

```

Figure 3. Random Counters implementation with TFlux DDM+TM pragma directives.

this code would have to be executed sequentially since the dependencies cannot be defined prior to the execution (due to the random memory accesses). By using TM we protect the values of the shared variables and ensure that, at the end of the execution, we will get the correct result while running the threads in parallel.

B. Labyrinth implementation of Lee's Algorithm

Lee's Algorithm is used in the process of producing an automated interconnection of electronic components. It guarantees to find a shortest interconnection between two points using the Expansion-Backtracking technique shown in Figure 4.

Starting from the source point S , the grid points are numbered by expanding a wavefront until the destination is reached (Figure 4(a)-(e)). At each phase during the expansion stage each grid point in the wavefront marks its unnumbered neighbors with an increment of its value. Once the destination D is reached, a route is traced back to the source by following any decreasing sequence of numbered grid points that are not used by others. Once the shortest route has been determined, the grid points reserved for this route cannot be used by others [26].

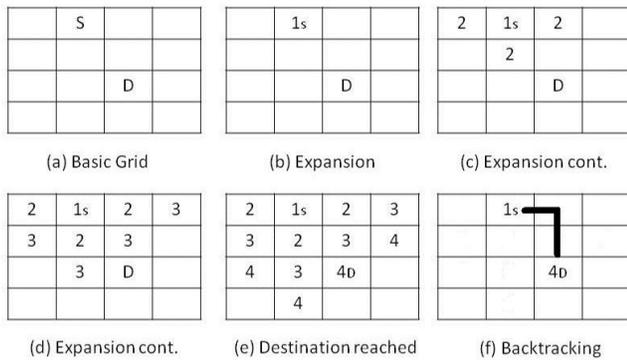


Figure 4. Lee's algorithm example.

For the purpose of this work we used the Labyrinth TM implementation of Lee's algorithm from STAMP [16] and

```

#pragma ddm atomic for thread 1
for (cv00 = 0; cv00 < THREADS; cv00++)
{
    #pragma ddm atomic transaction
        tvar(queue : READ_WRITE)
        [1] Get a (S, D) pair from the
            work queue
    #pragma ddm atomic endtransaction

    #pragma ddm atomic transaction
        tvar(global_grid : READ)
        [2] Copy global grid to threads
            local grids
    #pragma ddm atomic endtransaction

        [3] Expansion stage
        [4] Backtracking stage

    #pragma ddm atomic transaction
        tvar(global_grid : WRITE)
        [5] Write selected route back to
            global grid
    #pragma ddm atomic endtransaction
}
#pragma ddm atomic endfor

```

Figure 5. Lee's Algorithm pseudo-code with TFlux DDM+TM pragma directives.

ported it to DDM+TM. In Figure 5 we present the integration of TFlux DDM+TM pragmas into the application. Using

```

#pragma ddm atomic for and
#pragma ddm atomic endfor

```

directives we declare the boundaries of a DDM+TM thread. In step 1 each thread will take a (S, D) pair from the global work queue as an atomic action. In step 2 each thread will copy the global grid to its local memory space, and in steps 3-4 each thread will execute the algorithm's expansion and backtracking phase locally. Finally in step 5 each thread will write the selected route back to the global grid as an atomic action. In order for steps 1, 2 and 5 to be executed correctly we must ensure that no data conflicts occur on the work queue or global grid. Declaring steps 1, 2 and 5 as transactional code with the

```

#pragma ddm atomic transaction and
#pragma ddm atomic endtransaction

```

we provide an atomic execution for those steps ensuring that the final result will be correct. As for steps 3 and 4 the execution is done locally using only data local to each thread. This does not require this part of the code to be transactional and lets it execute in parallel within the boundaries of the thread, thus avoiding the need to re-execute non-transactional code in the event of a conflict. This is therefore an example where not all the code of a thread is transactional and hence need not be declared with the transactional version of the pragma.

Table IV
RANDOM COUNTER STATISTICS FOR TM AND DDM+TM IMPLEMENTATIONS.

Number of Threads	Commits	Aborts	
		TM	DDM+TM
2	200	38	47
4	400	157	479
8	800	920	1283

V. EXPERIMENTAL SETUP

We have evaluated the integration of TM with the Data-Flow model using the Random Counters and Lee’s algorithm described in Section IV. For Random Counters we used a conventional multi-threaded TM implementation as a baseline for our implementations of the DDM+TM version of the application. In this preliminary study we are interested in the overheads of the DDM+TM model over the TM implementation rather than detailed performance. For Lee’s algorithm we used the Labyrinth implementation from [16] that uses TM calls from TinySTM and integrated the TFlux directives to create the DDM+TM version. The results are for the parallel DDM+TM implementation of the applications on a 12-core machine with 2 6-core AMD Opteron(tm) Processors 2427 running at 2200MHz. The available memory of the system is 31GB of main memory and 512KB of L2 cache. The system is running an Ubuntu SMP x86-64 operating system.

The porting of the applications for the Data-Flow model was performed using the pragma directives from [25] of the TFlux Soft system [22] version 1.5.0. To integrate the transactions in the applications we used the TinySTM library. The execution time measurements were collected using the *gettimeofday* system call to measure the execution time of the section of code that has been parallelized. The input data sizes used for each application are depicted in Table V. All the results collected are for the Data-Flow model implementing transactions over the baseline execution. The baseline execution was considered to be the sequential execution of the applications implementing transactions with the TinySTM library. To calculate the results while avoiding any statistical errors we used the average of 10 executions removing the largest and the smallest execution time.

VI. EXPERIMENTAL RESULTS

In Table IV we present the statistics of the Random Counter application both for TM and DDM+TM implementations. It shows the number of commits for executions with a different number of threads and the average number of aborts for 10 executions. For each execution the statistics are different due to the use of *random* in the application. Since each execution of the application is random the statistics will follow a uniform random distribution. Consequently, the results demonstrate that the integration of TM with DDM

Table VI
LABYRINTH STATISTICS FOR TM AND DDM+TM IMPLEMENTATIONS.

Number of Threads	Commits	Aborts	
		Small	Large
2	1028	18	18
4	1032	47	46
8	1040	101	99
16	1056	182	175

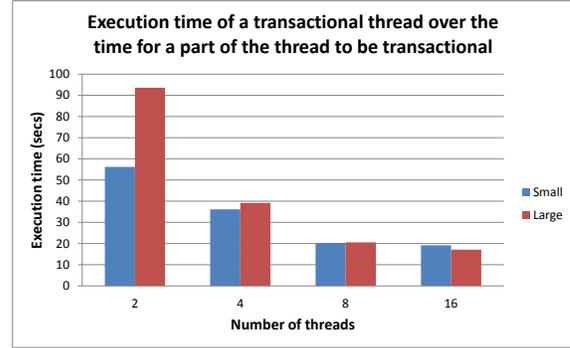


Figure 6. Comparison of complete and partially transactional threads in the Labyrinth implementation.

behaves similarly to a normal TM implementation but cannot be used to extract conclusions about detailed performance.

In Figure 6 we present the difference in the execution time of the Labyrinth benchmark when reducing the size of the transactional code inside a thread. From now on we reference the execution of a whole transactional thread as *large* and the execution of a part of the thread as transactional as *small*. We observe that on average the execution time is smaller when we declare only a part of the thread as transactional. As we increase the number of threads we see this difference decreasing. Correlating these results with Table VI, where we present the numerical results of the previous executions, we see that when the number of aborts is the same for the two executions the *small* implementation is faster and when the number of aborts for the *small* implementation exceeds the number of aborts of the *large* implementation the execution times are almost the same. From this correlation we conclude that the overhead of rescheduling a whole transactional thread is higher than rescheduling a small part of the thread. This is because the rescheduled code to be executed will be more in the former case.

Finally we present the experimental results that indicate the speedup for the DDM+TM implementation of the Labyrinth application over the TM sequential implementation; these are depicted in Figure 7. From these results it is easy to observe that for the two smaller input files (256x256x3-n256 and 256x256x5-n256) the application scales well up to 10 threads and then starts degrading. For the largest input file it scales beyond 10 threads with a max-

Table V
EXPERIMENTAL WORKLOAD DESCRIPTION AND PROBLEM SIZES.

Benchmark	Description	Problem size		
Random Counters	Randomly increment shared values	10 updates	100 updates	1000 updates
Labyrinth	Find connecting routes in a shared grid	256x256x3-n256	256x256x5-n256	512x512x7-n512

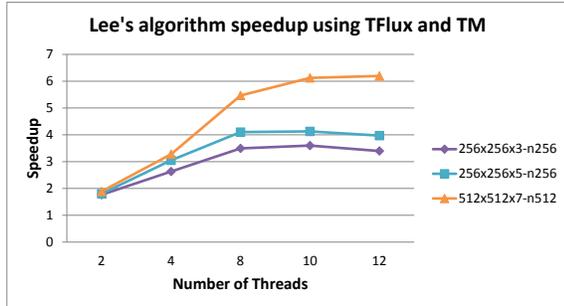


Figure 7. Lee's TM routing algorithm - scalability of TFlux extended with TM.

imum speedup of 6.2x over the sequential implementation with TM at 12 threads. The reason for this degradation in the speedup for more than 10 threads is that although we use a 12-core machine we also have the TSU and the OS simultaneously running with the application and occupying one core each [22]. As seen from Figure 7 the scalability of the DDM+TM model is good as it seems that at this early stage of the study there are no obvious extra overheads.

VII. CONCLUSIONS

We have reported our lessons from integrating, and performance experiments of, Data Driven Multi-threading (DDM) with Transactional Memory (TM). DDM offers the benefit of discouraging sharing of mutable data while expressing large amounts of parallelism. However, it is not always possible to avoid sharing mutable state and that reduces parallelism in the DDM model. By adding TM to DDM, we allow further parallelism to be exploited while introducing mutable shared state in a composable manner.

We have extended the TFlux directives to develop Data-Flow applications with new pragmas for TM. We have also described the extended TFlux runtime system with functionality provided by a software TM. We have tested the new runtime system by developing two applications that require TM: Random Counter and an implementation of Lee's parallel routing algorithm.

The experiments and our reported experience indicates that, from the runtime integration point of view, existing software implementations can be integrated without major problems. The biggest interaction and point for further exploration comes with the scheduling. When a complete DDM thread is a single transaction, the TSU could take ownership of the restart mechanism and reschedule the thread.

In future work we will address this and other potential optimizations as well as explore further parallelism in Data-Flow applications by using transactions.

ACKNOWLEDGMENTS

The authors would like to thank HiPEAC Network of Excellence for the collaboration grant and TERAFLUX project funded by the EC with Grant Agreement number 249013. Dr. Luján is supported by a Royal Society University Research Fellowship.

REFERENCES

- [1] Teraflux project. <http://www.teraflux.eu>.
- [2] M. Ansari, C. Kotselidis, K. Jarvis, M. Luján, C. C. Kirkham, and I. Watson. Advanced concurrency control for transactional memory using transaction commit rate. In *Euro-Par*, volume 5168 of *LNCS*, pages 719–728. Springer, 2008.
- [3] M. Ansari, M. Luján, C. Kotselidis, K. Jarvis, C. C. Kirkham, and I. Watson. Steal-on-abort: Improving transactional memory performance through dynamic transaction reordering. In *HiPEAC*, volume 5409 of *LNCS*, pages 4–18. Springer, 2009.
- [4] P. S. Barth, R. S. Nikhil, and Arvind. M-structures: Extending a parallel, non-strict, functional language with state. In *Proceedings of the 5th ACM Conference on Functional Programming Languages and Computer Architecture*, pages 538–568, London, UK, 1991. Springer-Verlag.
- [5] D. Cann. Retire Fortran?: a debate rekindled. *Commun. ACM*, 35:81–89, 1992.
- [6] K. Costas, P. Evripidou, and P. Trancoso. Data-driven multithreading using conventional microprocessors. *IEEE Transactions on Parallel and Distributed Systems*, pages 1176–1188, October 2006.
- [7] J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. In *Proceedings of the 6th conference on Symposium on Operating Systems Design & Implementation*, 2004.
- [8] D. Dice, O. Shalev, and N. Shavit. Transactional locking ii. *Proceedings of the 20th Intl. Symposium on Distributed Computing (DISC)*, 2006.
- [9] R. Giorgi, Z. Popovic, and N. Puzovic. Dta-c: A decoupled multi-threaded architecture for CMP systems. *19th International Symposium on Computer Architecture and High Performance Computing, 2007. SBAC-PAD 2007*, pages 263 – 270, October 2007.
- [10] J. R. Gurd, C. C. Kirkham, and I. Watson. The manchester prototype dataflow computer. *Commun. ACM*, 28:34–52, 1985.

- [11] T. Harris, J. R. Larus, and R. Rajwar. *Transactional Memory, 2nd edition*. Morgan & Claypool Publishers.
- [12] J. D. Hogg, J. K. Reid, and J. A. Scott. Design of a multicore sparse cholesky factorization using dags. *SIAM J. Scientific Computing*, 32(6):3627–3649, 2010.
- [13] W. M. Johnston, J. R. P. Hanna, and R. J. Millar. Advances in dataflow programming languages. *ACM Comput. Surv.*, 36:1–34, 2004.
- [14] S. P. Jones, A. Gordon, and S. Finne. Concurrent haskell. In *Annual Symposium on Principles of Programming Languages*, pages 295–308. ACM, 1996.
- [15] R. Milner, M. Tofte, and R. Harper. *The Definition of Standard ML*. MIT Press, Cambridge, MA, USA, 1990.
- [16] C. C. Minh, J. Chung, C. Kozyrakis, and K. Olukotun. Stamp: Stanford transactional applications for multi-processing. In *IISWC '08: Proceedings of The IEEE International Symposium on Workload Characterization*,, September 2008.
- [17] G. M. Papadopoulos and D. E. Culler. Monsoon: an explicit token-store architecture. In *Proceedings of the 17th annual international symposium on Computer Architecture*, ISCA '90, pages 82–91, 1990.
- [18] F. Pascal, F. Christof, and R. Torvald. Dynamic performance tuning of word-based software transactional memory. *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, 2008.
- [19] S. L. Peyton Jones. *Beautiful Concurrency*. 2007.
- [20] J. Sreeram, R. Cledat, T. Kumar, and S. Pande. Rstm : A relaxed consistency software transactional memory for multicores. *Parallel Architectures and Compilation Techniques, International Conference on*, 0:428, 2007.
- [21] K. Stavrou, P. Evripidou, and P. Trancoso. DDM-CMP: Data-driven multithreading on a chip multiprocessor. In *Embedded Computer Systems: Architectures, Modeling, and Simulation*, 3553:205–224, 2005.
- [22] K. Stavrou, M. Nikolaides, D. Pavlou, S. Arandi, P. Evripidou, and P. Trancoso. Tflux: A portable platform for data-driven multithreading on commodity multicore systems. *37th International Conference on Parallel Processing*, pages 25–34, 2008.
- [23] D. Syme, A. Granicz, and A. Cisternino. *Expert F# (Expert's Voice in .Net)*.
- [24] S. Tomov, R. Nath, H. Ltaief, and J. Dongarra. Dense linear algebra solvers for multicore with gpu accelerators. In *International Symposium on Parallel and Distributed Processing*, 2010.
- [25] P. Trancoso, K. Stavrou, and P. Evripidou. DDMCPP: The data-driven multithreading C pre-processor. In *Proceedings of the 11th Interact-11*, pages 32–39.
- [26] I. Watson, C. Kirkham, and M. Luján. A study of a transactional parallel routing algorithm. *PACT '07 Proceedings of the 16th International Conference on Parallel Architecture and Compilation Techniques*, pages 388–398, 2007.
- [27] I. Watson, V. Woods, P. Watson, R. Banach, M. Greenberg, and J. Sargeant. Flagship: A parallel architecture for declarative programming. In *ISCA*, pages 124–130, 1988.