

Dynamic Verification for Hybrid Concurrent Programming Models

Erdal Mutlu
Koc University
ermutlu@ku.edu.tr

Vladimir Gajinov
Barcelona Supercomputing
Center
vladimir.gajinov@bsc.es

Adrián Cristal
Barcelona Supercomputing
Center
adrian.cristal@bsc.es

Serdar Tasiran
Koc University
stasiran@ku.edu.tr

Osman S. Unsal
Barcelona Supercomputing
Center
osman.unsal@bsc.es

1. INTRODUCTION

Most modern computation platforms feature multiple CPU and GPU cores. For many large applications, it is more convenient for programmers to make use of combination of different programming models to coordinate different kinds of concurrency and communication in the program. In this paper, we explore hybrid concurrent programming models that combine shared memory with dataflow abstractions. Shared memory multi-threading is ubiquitous in concurrent programs. By contrast, in the dataflow programming model, the execution of an operation is constrained only by the availability of its input data – a feature that makes dataflow programming convenient and safe when it fits the problem at hand.

Using the dataflow programming model in conjunction with shared memory mechanisms can make it convenient and natural for programmers to express the parallelism inherent in a problem as evidenced by recent proposals [4, 8] and adoptions [5, 6, 7]. The proposed hybrid programming models [4, 8] provide programmers with dataflow abstractions for defining tasks as the main execution unit with corresponding data dependencies. Contrary to the pure dataflow model which assumes side-effect free execution of the tasks, these models allow tasks to share the data using some form of thread synchronization, such as locks or transactional memory (TM). In this way, they facilitate implementation of complex algorithms for which shared state is the fundamental part of how the computational problem at hand is naturally expressed.

Enabling a combination of different programming models provides a user with a wide choice of parallel programming abstractions that can support a straightforward implementation of a wider range of problems. However, it also increases the likelihood of introducing concurrency bugs, not only those specific to a given well-studied programming model, but also those that are the result of unexpected program behavior caused by an incorrect use of different programming abstractions within the same program. Since the hybrid dataflow models we consider in this paper are quite novel, many of the bugs that belong to the latter category may not have been studied. The goal of this work is to identify these bugs and design a verification tool that can facilitate automated behavior exploration targeting their detection.

We present a dynamic verification tool for characterizing and exploring behaviors of programs written using hybrid dataflow programming models. We focus in particular on the Atomic DataFlow (ADF) programming model [4] as a representative of this class of programming models. In the ADF model, a program is based on tasks for which data dependencies are explicitly defined by a programmer and used by the runtime system to coordinate the task execution, while the memory shared between potentially concurrent tasks is managed using transactional memory (TM). While ideally these two domains should be well separated within a program, concurrency bugs can lead to an unexpected interleaving between these domains, leading to incorrect program behavior.

We devised a randomized scheduling method for exploring programs written using ADF. The key challenge in our work was precisely characterizing and exploring the concurrency visible and meaningful to the programmer, as opposed to the concurrency present in the dataflow runtime or TM implementations. For exploration of different interleavings, we adapted the dynamic exploration technique “Probabilistic Concurrency Testing (PCT)” [3] to ADF programs in order to amplify the randomness of observed schedules [2].

2. MOTIVATION

In this section, we illustrate the effectiveness of ADF programming model on a simple example and describe an unexpected execution scenario for motivating our dynamic verification method. In this example, two input streams of integers, x and y , are given and the goal is to provide, for each pair of tokens consumed from these streams, the larger of the two values, and also to maintain the global maximum of all values received on both streams. Figure 1 shows how this program is implemented in pure dataflow, shared memory and the ADF programming models. In pure dataflow, maintaining the global maximum is inefficient (the max of two tokens needs to be fed back to the task for the next step) whereas in the shared memory model, the global maximum is treated as the mutable state which can easily be updated in-place. However, condition variables are needed to detect the availability of new input tokens. ADF combines the best of two worlds. The atomicity of the code in the “atomic” block in Figure 1-b and c can be ensured using

locks or transactional memory.




Figure 1: Computing maxima using: a) dataflow, b) shared memory and c) ADF

Due to the asynchronous concurrent execution of tasks in the ADF model, users can face unexpected execution orders causing atomicity violations between dataflow tasks. To illustrate such a behavior, consider two ADF tasks, *max_min* (Figure 2) that compute the maximum and minimum values from two input streams while updating a global minimum and maximum, and *comp_avg* that uses the output streams provided by *max_min* for comparing the average values of *g_max* and *g_min* with the input values and returning the bigger one. As seen in Figure 2-c, the dependencies between these tasks can be expressed with ADF programming model naturally as shown in Figure 2-a and b. However, while these particular implementations appear correct separately, when combined, they may result in unexpected behavior in an ADF execution. As the updates on the global variables, *g_max* and *g_min*, are performed in separate atomic blocks, concurrently running tasks can read incorrect values of global variables. Imagine an execution where first pair from the input streams are processed by *max_min* and then passed to *comp_avg*. During the execution of *comp_avg*, *max_min* can start to process the second pair and update *g_max* value causing *comp_avg* to read the new *g_max* value from the second iteration while reading *g_min* value from the first one. Such concurrency scenarios that arise due to an interaction between dataflow and shared memory may be difficult to foresee for a programmer and are not addressed properly by verification methods for pure dataflow or pure shared memory model.

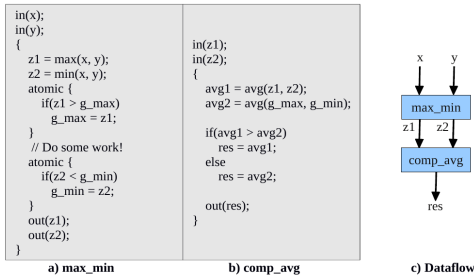


Figure 2: Motivating example

3. OUR METHOD

The ADF programming model has an inherently asynchronous concurrent execution model where tasks can be enabled and executed multiple times. In addition, programmers are allowed to provide their own relaxed synchronizations using transactional memory in ADF tasks, which can possibly influence the dataflow execution. In order to fully investigate

behaviors of programs written using a hybrid model such as ADF, the dynamic exploration technique has to be aware of both the dataflow structure and the specifics of the shared memory synchronization mechanism. Furthermore, the dynamic verification tool should not simply instrument the platform implementations for transactional memory, atomic blocks and dataflow. This would not only be very inefficient, but it would not provide value to the programmer. The user of a hybrid concurrent programming model is not interested in the concurrency internal to the platform implementing the model, which should be transparent to the programmer, but only in the non-determinism made visible at the programming model level.

For our initial study, we investigated the PCT algorithm which defines a *bug depth* parameter as the minimum number of ordering constraints that are sufficient to find a bug and uses a randomized scheduling method, with provably good probabilistic guarantees, to find all bugs of low depth. PCT makes use of a priority based scheduler that maintains randomly assigned priorities for each thread and a list of randomly assigned *priority change points* for simulating the ordering constraints.

We build upon the PCT algorithm but redefine priority assignment points, making use of TM transaction boundaries for priority change point assignment. Rather than using the original ADF work-stealing scheduler based on a pool of worker threads, we have devised a new scheduler that creates a thread with randomly assigned priority for each enabled task and sequentially schedules the threads by honoring their priorities. Likewise, instead of using the original priority change point assignment from the PCT method, we narrowed possible priority change point location to the beginning and the end of atomic regions only. Further, we devised a monitoring mechanism for checking globally defined invariants during an execution. Differently from pure shared memory programming models, writing global properties as assertion is not practical within ADF tasks that are executing in a dataflow fashion. We provide the users with the capability to write global invariants on shared variables. These can be checked at every step by our tool, or at randomly assigned points in the execution.

4. ONGOING WORK

We have started by investigating ADF implementations of DWARF [1] benchmark applications. These applications are mostly numerical computations that have a structured dataflow with little shared memory accesses. We believe these to be a good initial set of benchmarks for discovering possibly missed cases in dataflow-heavy implementations. In later experimental work, we plan to investigate the dynamic verification of the ADF implementation of a parallel game engine. In this complex application, the game map is divided between different tasks that process the objects moving between map regions. Dataflow is used to coordinate the execution of tasks that correspond to different game regions, whereas the TM synchronization is used to protect lists of objects, associated with each game region, that hold all the objects physically located within a region. By using the game engine application, we wish to evaluate how well our exploration method behaves with performance-critical applications characterized with highly-irregular behavior.

5. REFERENCES

- [1] K. Asanovic, R. Bodik, J. Demmel, T. Keaveny, K. Keutzer, J. Kubiawicz, N. Morgan, D. Patterson, K. Sen, J. Wawrzynek, D. Wessel, and K. Yelick. A view of the parallel computing landscape. *Commun. ACM*, 52(10):56–67, Oct. 2009.
- [2] Y. Ben-Asher, Y. Eytani, E. Farchi, and S. Ur. Producing scheduling that causes concurrent programs to fail. PADTAD '06, pages 37–40, New York, NY, USA, 2006. ACM.
- [3] S. Burckhardt, P. Kothari, M. Musuvathi, and S. Nagarakatte. A randomized scheduler with probabilistic guarantees of finding bugs. ASPLOS XV, pages 167–178, New York, NY, USA, 2010. ACM.
- [4] V. Gajinov, S. Stipic, O. Unsal, T. Harris, E. Ayguade, and A. Cristal. Integrating dataflow abstractions into the shared memory model. SBAC-PAD, pages 243–251, 2012.
- [5] Intel. Intel threading building blocks - flow graph. http://www.threadingbuildingblocks.org/docs/help/reference/flow_graph.htm.
- [6] Microsoft. Task parallel library - dataflow. <http://msdn.microsoft.com/en-us/library/hh228603.aspx>.
- [7] OpenMP. Openmp 4.0 specification. <http://www.openmp.org/mp-documents/OpenMP4.0.0.pdf>.
- [8] C. Seaton, D. Goodman, M. Luján, and I. Watson. Applying dataflow and transactions to Lee routing. In *Workshop on Programmability Issues for Heterogeneous Multicores*, 2012.