Analyzing the Intel Itanium Memory Ordering Rules Using Logic Programming and SAT*

Yue Yang, Ganesh Gopalakrishnan, Gary Lindstrom, and Konrad Slind

School of Computing, University of Utah
{yyang,ganesh,gary,slind}@cs.utah.edu

Abstract. We present a non-operational approach to specifying and analyzing shared memory consistency models. The method uses higher order logic to capture a complete set of ordering constraints on execution traces, in an axiomatic style. A direct encoding of the semantics with a constraint logic programming language provides an interactive and incremental framework for exercising and verifying finite test programs. The framework has also been adapted to generate equivalent boolean satisfiability (SAT) problems. These techniques make a memory model specification executable, a powerful feature lacked in most non-operational methods. As an example, we provide a concise formalization of the Intel Itanium memory model and show how constraint solving and SAT solving can be effectively applied for computer aided analysis. Encouraging initial results demonstrate the scalability for complex industrial designs.

1 Introduction

Modern shared memory architectures rely on a rich set of memory access related instructions to provide the flexibility needed by software. For instance, the Intel ItaniumTM processor family [1] provides two varieties of loads and stores in addition to fence and semaphore instructions, each associated with different ordering restrictions. A memory model defines the underlying memory ordering semantics. Proper understanding of these ordering rules is essential for the correctness of shared memory consistency protocols that are aggressive in their ordering permissiveness, as well as for compiler transformations that rearrange multithreaded programs for higher performance. Due to the complexity of advanced computer architectures, however, practicing designers face a serious problem in reliably comprehending the memory model specification.

Consider, for example, the assembly code shown in Fig. 1 that is run concurrently on two Itanium processors (such code fragments are generally known as *litmus tests*): The first processor, P1, executes a *store* of datum 1 into ad-

* This work was supported by a grant from the Semiconductor Research Corporation for Task 1031.001, and Research Grants CCR-0081406 and CCR-0219805 of NSF.

P1	P2
<code>st a,1;</code>	<code>ld.acq r1,b;</code>
<code>st.rel b,1;</code>	<code>ld r2,a;</code>

Fig. 1. A litmus test showing the ordering properties of store-release and load-acquire. Initially, $a = b = 0$. Can it result in $r1 = 1$ and $r2 = 0$? The Itanium memory model does not permit this result.

dress **a**; it then performs a *store-release*¹ of datum 1 into address **b**. Processor P2 performs a *load-acquire* from **b**, loading the result into register **r1**. It is followed by an *ordinary load* from location **a** into register **r2**. The question arises: if all locations initially contain 0, can the final register values be $r1=1$ and $r2=0$? To determine the answer, the Itanium memory model must be consulted. The formal specification of the Itanium memory model is given in an Intel application note [2]. It comprises a complex set of ordering rules, 24 of which are expressed explicitly based on a large amount of special terminology. One can follow a pencil-and-pen approach to reason that the execution shown in Fig. 1 is not permitted by the rules specified in [2]. Based on this, one can conclude that even though the instructions in P2 pertain to different addresses, the underlying hardware is not allowed to carry out the ordinary load at the beginning, and by the same token, a shared memory consistency protocol or an optimizing compiler cannot reorder the instructions in P2. A further investigation shows that the above result would be permitted if the `st.rel` in P1 is changed to a `st`, or the `ld.acq` in P2 is changed to a `ld`. Therefore, `st.rel` and `ld.acq` must both be used in pairs to achieve the “barrier” effect in this scenario.

A litmus test like this can reveal critical information to help system designers make right decisions in code selection and optimization. But as bigger tests are used and more intricate rules are involved, trace properties quickly become non-intuitive and hand-proving program compliance can be very difficult. How can one be assured that there does not exist an interacting rule that might introduce unexpected implications? Also, a large scale design is often composed of simpler components. To avoid being overwhelmed by the overall complexity, a useful technique is to isolate the rules related to specific architectural features so that the model can be analyzed piece by piece. For example, if one can selectively enable/disable certain rules, one may quickly find out that the “program order” rules in [2] are critical to the scenario in Fig. 1 while many others are irrelevant.

These issues suggest that a series of useful features is needed from the specification framework to help people better understand the underlying model. Unfortunately, most non-operational specification methods leave these issues unresolved because they use notations that do not support analysis through execu-

¹ Briefly, a store-release instruction will, at its completion, ensure that all previous instructions are completed; a load-acquire instruction correspondingly ensures that all following instructions will complete only after it completes. These explanations are far from precise - what do “previous” and “completion” mean? A formal specification of a memory model is key to precisely capture these and all similar notions.

tion. Given that designers need lucid and reliable memory model specifications, and given that memory model specifications can live for decades, it is crucial that progress be made in this regard.

In this paper, we take a fresh look at the non-operational specification method and explore what verification techniques can be applied. We make the following contributions. First, we present a compositional method to axiomatically capture all aspects of the memory ordering requirements, resulting a comprehensive, constraint-based memory consistency model. Second, we propose a method to encode these specifications using FD-Prolog.² This enables one to perform interactive and incremental analysis. Third, we have harnessed a boolean satisfiability checker to solve the constraints. To the best of our knowledge, this is the first application of SAT methods for analyzing memory model compliance. As a case study of this approach, we have formalized a core subset of the Itanium memory model and used constraint programming and boolean satisfiability for program analysis.

Related Work. The area of memory model specification has been pursued under different approaches. Some researchers have employed *operational* style specifications [3] [4] [5] [6], in which the update of a global state is defined step-by-step with the execution of each instruction. For example, an operational model [4] for Sparc V9 [7] was developed in Murphi. With the model checking capability supported by Murphi, this executable model was used to examine many code sequences for Sparc V9. While the operational descriptions often mirror the decision process of an implementer and can be exploited by a model checker, they are not declarative. Hence they tend to emphasize the *how* aspects through their usage of specific data structures, not the *what* aspects that formal specifications are supposed to emphasize.

Other researchers have used *non-operational* (also known as *axiomatic*) specifications, in which the desired properties are directly defined. Non-operational styles have been widely used to describe conceptual memory models [8] [9]. One noticeable limitation of these specifications is the lack of a means for automatic execution. An axiomatic specification of the Alpha memory model was written by Yu [10] in Lisp. Litmus tests were written in S-expression. Verification conditions were generated for the litmus tests and fed to the Simplify [11] verifier of Compaq/SRC. In contrast, our specification is much closer to the actual industrial specification, thanks to the declarative nature of FD-Prolog. The FD constraint solver offers a more interactive and incremental environment. We have also applied SAT and demonstrated its effectiveness.

Lamport and colleagues have specified the Alpha and Itanium memory models in TLA+ [12] [13]. These specifications build visibility order inductively and support the execution of litmus tests. While their approach also precisely specifies the ordering requirement, the manner in which such inductive definitions

² FD-Prolog refers to Prolog with a finite domain (FD) constraint solver. For example, SICStus Prolog and GNU Prolog have this feature.

Fig. 2. The process of making an axiomatic memory model executable. Legality of a litmus test can be checked by either a constraint solver or a SAT solver.

are constructed will vary from memory model to memory model, making comparisons among them harder. Our method instead relies on primitive relations and directly describes the components to make up a full memory model. This makes our specification easier to understand, and more importantly, to compare against other memory models. This also means we can enable or disable some sub-rules quite reliably without affecting the other primitive ordering rules - a danger in a style which merges all the ordering concerns in a monolithic manner.

Roadmap. In the next section, we introduce our methodology. Section 3 describes the Itanium memory ordering rules. Section 4 demonstrates the analysis of the Itanium memory model through execution. We conclude and propose future works in Section 5. The concise specification of the Itanium ordering constraints is provided in the Appendix, with additional details appearing at our web site http://www.cs.utah.edu/formal_verification/itanium.

2 Overview of the Framework

A pictorial representation of our methodology is shown in Fig. 2. We use a collection of primitive ordering rules, each serving a clear purpose, to specify even the most challenging commercial memory models. This approach mirrors the style adopted in modern declarative specifications written by the industry, such as [2]. Moreover, by using pure logic programs supported by certain modern flavors of Prolog that also include finite domain constraints, one can directly capture these higher order logic specifications and also interactively execute the specifications to obtain execution results for litmus tests. Alternatively, we can obtain SAT instances of the boolean constraints representing the memory model through symbolic execution, in which case boolean satisfiability tools can be employed to quickly answer whether the tests are legal or not.

2.1 Specification Method

To define a memory model, we use predicate calculus to specify all constraints imposed on an ordering relation *order*. The constraints are almost completely

first-order. However, since *order* is a parameter to the specification, the constraints are most easily captured with higher order predicate calculus (we use the HOL logic [14]). Previous non-operational specifications often *implicitly* require general ordering properties, such as totality, transitivity, and circuit-freeness. This is the main reason why such specifications cannot readily be executed. In contrast, we are fully explicit about such properties, and so our constraints completely characterize the memory model.

The flexibility of our notation allows us to specify different memory models under the same framework. We have assembled a large collection of constraints for many conventional memory models, such as Sequential Consistency [8], Coherence, Causal Consistency [9], PRAM [15], and Processor Consistency [16]. Due to space limitation, this paper concentrates on demonstrating how to specify and analyze the Itanium memory ordering rules.

2.2 Executing Axiomatic Specifications

A straightforward transcription of the formal predicate calculus specification into a Prolog-style logic program makes it possible for interactive and incremental execution of litmus tests. This encourages exploration and experiment in the development and validation of complex coherence protocols. To make a specification executable, we instantiate it over a finite execution and convert the verification problem to a satisfiability problem.

The Algorithm. Given a finite execution *ops* with n operations, there are n^2 ordering pairs, constituting an ordering matrix $\mathcal{M}$, where the element $\mathcal{M}_{ij}$ indicates whether operations i and j should be ordered. We go through each ordering rule in the specification and impose the corresponding constraint regarding the elements of $\mathcal{M}$. Then we check the satisfiability of all the ordering requirements. If such a $\mathcal{M}$ exists, the trace *ops* is legal, and a valid interleaving can be derived from $\mathcal{M}$. Otherwise, *ops* is not a legal trace.

Applying Constraint Logic Programming. Logic programming differs from conventional programming in that it describes the logical structure of the problems rather than prescribing the detailed steps of solving them. This naturally reflects the philosophy of the axiomatic specification style. As a result, our formal specification can be easily encoded using Prolog. Memory ordering constraints can be solved through a conjunction of two mechanisms that FD-Prolog readily provides. One applies backtracking search for all constraints expressed by logical variables, and the other uses non-backtracking constraint solving based on *arc consistency* [17] for FD variables, which is potentially more efficient and certainly more complete (especially under the presence of negation) than with logical variables. This works by adding constraints in a monotonically increasing manner to a constraint store, with the in-built constraint propagation rules of FD-Prolog helping refine the variable ranges (or concluding that the constraints are not satisfiable) when constraints are asserted to the constraint store.

Applying Boolean Satisfiability Techniques. The goal of a boolean satisfiability problem is to determine a satisfying variable assignment for a boolean formula or to conclude that no such assignment exists. A slight variant of the Prolog code can let us benefit from SAT solving techniques, which have advanced tremendously in recent years. Instead of solving constraints using a FD solver, we can let Prolog emit SAT instances through symbolic execution. The resultant formula is true if and only if the litmus test is legal under the memory model. It is then sent to a SAT solver to find out the result.

3 Specifying the Itanium Memory Consistency Model

The original Itanium memory ordering specification is informally given in various places in the Itanium architecture manual [1]. Intel later provided an application note [2] to guide system developers. This document uses a combination of English and informal mathematics to specify a core subset of memory operations in a non-operational style. We demonstrate how the specification of [2] can be adapted to our framework to enable computer aided analysis. Virtually the entire Intel application note has been captured.³ We assume proper address alignment and common address size for all memory accesses, which would be the common case encountered by programmers (even these restrictions could be easily lifted). The detailed definition of the Itanium memory model is presented in the Appendix. This section explains each of the rules. The following definitions are used throughout this paper:

Instructions – Instructions with memory access or memory ordering semantics. Five instruction types are defined in this paper: load-acquire (**ld.acq**), store-release (**st.rel**), unordered load (**ld**), unordered store (**st**), and memory fence (**mf**). An instruction i may have *read semantics* (**isRd** $i = \text{true}$) or *write semantics* (**isWr** $i = \text{true}$). **Ld.acq** and **ld** have read semantics. **St.rel** and **st** have write semantics. **Mf** has neither read nor write semantics. Instructions are decomposed into *operations* to allow a finer specification of the ordering properties.

Execution – Also known as an *execution trace*, contains all memory operations generated by a program. Stores are annotated with the write data and loads are annotated with the return data. An execution is *legal* if there exists an order among the operations that satisfies all memory model constraints.

Address Attributes – Every memory location is associated with an address attribute, which can be write-back (WB), uncacheable (UC), or write-coalescing (WC). Memory ordering semantics may vary for different attributes. Predicate **attribute** is used to find the attribute of a location.

³ This paper formally captures 21 out of 24 rules from [2]. The remaining 3 rules deal with semaphore operations, which are straightforward to add using the same approach.

Operation Tuple – A tuple containing necessary attributes is used to mathematically describe memory operations. A memory operation i is represented by a tuple $\langle P, Pc, Op, Var, Data, WrId, WrType, WrProc, Reg, UseReg, Id \rangle$, where

p $i = P$:	issuing processor
pc $i = Pc$:	program counter
op $i = Op$:	instruction type
var $i = Var$:	shared memory location
data $i = Data$:	data value
wrID $i = WrId$:	identifier of a write operation
wrType $i = WrType$:	type of a write operation
wrProc $i = WrProc$:	target processor observing a write operation
reg $i = Reg$:	register
useReg $i = UesReg$:	flag of a write indicating if it uses a register
id $i = Id$:	global identifier of the operation

A read instruction or a fence instruction is decomposed into a single operation. A write instruction is decomposed into multiple operations, comprising a local write operation (**wrType** $i = Local$) and a set of remote write operations (**wrType** $i = Remote$) for each target processor (**wrProc** i), which also includes the issuing processor. Every write operation i that originates from a single write instruction shares the same program counter (**pc** i) and write ID (**WrID** i).

3.1 The Itanium Memory Ordering Rules

As shown below, predicate **legal** is a top-level constraint that defines the legality of a trace ops by checking the existence of an $order$ among ops that satisfies all requirements. Each requirement is formally defined in the Appendix.

legal $ops \equiv \exists order.$

```

requireLinearOrder  $ops\ order \wedge$ 
requireWriteOperationOrder  $ops\ order \wedge$ 
requireProgramOrder  $ops\ order \wedge$ 
requireMemoryDataDependence  $ops\ order \wedge$ 
requireDataFlowDependence  $ops\ order \wedge$ 
requireCoherence  $ops\ order \wedge$ 
requireReadValue  $ops\ order \wedge$ 
requireAtomicWBRelease  $ops\ order \wedge$ 
requireSequentialUC  $ops\ order \wedge$ 
requireNoUCBypass  $ops\ order$ 

```

Table 1 illustrates the hierarchy of the Itanium memory model definition. Most constraints strictly follow the rules from [2]. We also explicitly add a predicate **requireLinearOrder** to capture the general ordering requirement since [2] has only English to convey this important ordering property.

Table 1. The specification hierarchy of the Itanium memory ordering rules.

<i>requireLinearOrder</i> - requireWeakTotal - requireTransitive - requireAsymmetric	<i>requireMemoryDataDependence</i> - MD:RAW - MD:WAR - MD:WAW	<i>requireReadValue</i> - validWr - validLocalWr - validRemoteWr - validDefaultWr - validRd
<i>requireWriteOperationOrder</i> - local/remote case - remote/remote case	<i>requireDataFlowDependence</i> - DF:RAR - DF:RAW - DF:WAR	<i>requireNoUCBypass</i>
<i>requireProgramOrder</i> - acquire case - release case - fence case	<i>requireCoherence</i> - local/local case - remote/remote case	<i>requireSequentialUC</i> - RAR case - RAW case - WAR case - WAW case
	<i>requireAtomicWBRelease</i>	

General Ordering Requirement (Appendix A.1). This requires order to be a weak total order which is also circuit-free.

Write Operation Order (Appendix A.2). This specifies the ordering among write operations that originate from a single write instruction. It guarantees that no write can become visible remotely before it becomes visible locally.

Program Order (Appendix A.3). This restricts reordering among instructions of the same processor with respect to the program order.

Memory-Data Dependence (Appendix A.4). This restricts reordering among instructions from the same processor when they access *common locations*.

Data-Flow Dependence (Appendix A.5). This is intended to specify how local *data dependency* and *control dependency* should be treated. However, this is an area that is not fully specified in [2]. Instead of pointing to an informal document as done in [2], we provide a formal specification covering most cases of data dependency, namely establishing data dependency between two memory operations by checking the conflict usages of local registers.⁴

Coherence (Appendix A.6). This constrains the order of writes to a *common location*. If two writes to the same location with the attribute of WB or UC become visible to a processor in some order, they must become visible to all processors in that order.

⁴ We do not cover branch instructions or indirect-mode instructions that also induce data dependency. We provide enough data dependency specification to let designers experiment with straight-line code that uses registers - this is an important requirement to support execution.

P1	P2
(1) <code>st_local(a,1);</code>	(7) <code>ld.acq(1,b);</code>
(2) <code>st_remote1(a,1);</code>	(8) <code>ld(0,a);</code>
(3) <code>st_remote2(a,1);</code>	
(4) <code>st.rel_local(b,1);</code>	
(5) <code>st.rel_remote1(b,1);</code>	
(6) <code>st.rel_remote2(b,1);</code>	

Fig. 3. An execution resulted from the program in Fig. 1. Stores are decomposed into local stores and remote stores. Loads are associated with return values.

Read Value (Appendix A.7). This defines what data can be observed by a read operation. There are three scenarios: a read can get the data from a local write (`validLocalWr`), a remote write (`validRemoteWr`), or the default value (`validDefaultWr`). Similar to shared memory read value rules, predicate `validRd` guarantees consistent assignments of registers - the value of a register is obtained from the most recent previous assignment of the same register.

Total Ordering of WB Releases (Appendix A.8). This specifies that store-releases to write-back (WB) memory must obey remote write atomicity, i.e., they become remotely visible atomically.

Sequentiality of UC Operations (Appendix A.9). This specifies that operations to `uncacheable(UC)` memory locations must have the property of *sequentiality*, i.e., they must become visible in program order.

No UC Bypassing (Appendix A.10). This specifies that `uncacheable(UC)` memory does not allow local bypassing from UC writes.

4 Making the Itanium Memory Model Executable

We have developed two methods to analyze the Itanium memory model. The first, as mentioned earlier, uses Prolog backtracking search, augmented with finite-domain constraint solving. The second approach targets the powerful SAT engines that have recently emerged.

The Constraint Logic Programming Approach

Our formal Itanium specification is implemented in SICStus Prolog [18]. Litmus tests are contained in a separate test file. When a test number is selected, the FD constraint solver examines all constraints automatically and answers whether the selected execution is legal. By running the litmus tests we can learn the degree to which executions are constrained, i.e., we can obtain a general view of the global ordering relation between pairs of instructions.

	1	2	3	4	5	6	7	8
1	0	1	1	0	0	0	0	0
2	0	0	1	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	1	1	1	0	1	1	1	0
5	1	1	1	0	0	1	1	0
6	1	1	1	0	0	0	1	0
7	1	1	1	0	0	0	0	0
8	1	1	1	1	1	1	1	0

Fig. 4. A legal ordering matrix for the execution shown in Fig. 3 when `requireProgramOrder` is disabled. A value 1 indicates that the two operations are ordered. A possible interleaving 8 4 5 6 7 1 2 3 is also automatically derived from this matrix.

Consider, for example, the program discussed earlier in Fig. 1. Its instructions are decomposed into operations as shown in Fig. 3. After taking this trace as input, the Prolog tool attempts all possible orders until it can find an instantiation that satisfies all constraints. For this particular example, it returns “illegal trace” as the result. If one comments out the `requireProgramOrder` rule and examines the trace again, the tool quickly finds a legal ordering matrix and a corresponding interleaving as shown in Fig. 4. Many other experiments can be conveniently performed in a similar way. Therefore, not only does this approach give people the notation to write rigorous as well as readable specifications, it also allows users to play with the model, asking “what if” queries after selectively enabling/disabling the ordering rules that are crucial to their work. We can also use the built-in predicate `setof` provided by Prolog to collect all legal return values for read operations. This is achieved by repeatedly backtracking and gradually building up a list of the solutions.

Although translating the formal specification to Prolog is fairly straightforward, there does exist some “logic gap” between predicate calculus and Prolog. Most Prolog systems do not directly support quantifiers. Therefore, we need to implement the effect of a universal quantifier by enumerating the related finite domain. The existential quantifier is realized by the backtracking mechanism of Prolog when proper predicate conditions are set.

The SAT Approach

As an alternative method, we use our Prolog program as a driver to emit propositional formulae asserting the solvability of the constraints. After being converted to the DIMACS format, the final formula is sent to a SAT solver, such as **ZChaff** [19] or **berkmin** [20]. Although the clause generation phase can be detached from the logic programming approach, the ability to have it coexist with FD-Prolog might be advantageous since it allows the two methods to share the same specification base.

Performance Results

Performance statistics from some litmus tests is shown below. These tests are chosen from [2] and represented by their original table numbers. Results are

measured on a Pentium III 900 MHz machine with 256 MB of RAM running Windows 2000. SICStus Prolog is run under compiled mode. The SAT solver used is ZChaff.

Test	Result	FD Solver(sec)	Vars	Clauses	SAT(sec)	CNF Gen (sec)
[2, Table 5]	illegal	0.49	64	679	0.01	3.67
[2, Table 10]	legal	3.00	100	1280	0.01	8.23
[2, Table 15]	illegal	22.29	576	15706	0.01	211.76
[2, Table 18]	illegal	2.40	144	2125	0.01	15.75
[2, Table 19]	legal	4.81	144	2044	0.01	15.68

Although satisfiability problems are NP-complete, the performance in practice has been acceptable. For the method using SAT solvers, the clause generation time is noticeably larger than the actual SAT solving time, since the entire formula is encoded at once through symbolic execution and is recursively simplified afterwards. Alternative boolean formula encoding techniques, such as the one discussed in [21], may help speed up this process.

5 Conclusions

The setting in which contemporary memory models are expressed and analyzed needs to be improved. Towards this, we present a framework based on axiomatic specifications (expressed in higher order logic) of memory ordering requirements. It is straightforward to encode these requirements as constraint logic programs or, by an extra level of translation, as boolean satisfiability problems. Our techniques are demonstrated through the adaptation and analysis of the Itanium memory model. Being able to tackle such a complex design also attests to the scalability of our framework for cutting-edge commercial architectures.

Our methodology provides several benefits. First, the ability to execute the underlying model is a powerful feature that promotes understanding. Second, the compositional specification style provides modularity, reusability, and scalability. It also allows one to change constraints incrementally for investigation purposes. Third, the expressive power of the underlying logic allows one to define a wide range of requirements using the same notation, providing a rich taxonomy for memory consistency models. Finally, the method of converting axiomatic rules to a propositional formula allows one to perform property checking through boolean reasoning, thus opening up a new means to conduct memory model verification.

Future Work. One possible enhancement is to develop the capability of exercising *symbolic (non-ground)* litmus tests. Such a tool may be used to automatically synthesize critical instructions of concurrent code fragments comprising compiler idioms or other synchronization primitives. For example, one could imagine using a symbolic store instruction in a program and asking a tool to solve whether it should be an “ordinary” or a “release” store to help generate aggressive code. Another area of improvement is in reducing the logic gap between the formal specification and the tools that execute the specification. One possibility is to

apply a *quantified boolean formulae* (QBF) solver that directly accepts quantifiers. The research of QBF solvers is still at a preliminary stage compared to propositional SAT. We hope our work can help accelerate its development by providing industrially motivated benchmarks.

References

1. Intel Itanium Architecture Software Developer's Manual, <http://developer.intel.com/design/itanium/manuals.htm>
2. A Formal Specification of Intel Itanium Processor Family Memory Ordering, Application Note, Document Number: 251429-001 (October, 2002)
3. K. Gharachorloo: Memory consistency models for shared-memory multiprocessors. Technical Report CSL-TR-95-685, Stanford University (December 1995)
4. D. Dill, S. Park, A. Nowatzky: Formal Specification of Abstract Memory Models. Research on Integrated Systems: Proceedings of the 1993 Symposium
5. Prosenjit Chatterjee, Ganesh Gopalakrishnan: Towards a Formal Model of Shared Memory Consistency for Intel Itanium. ICCD 2001, Austin, TX (Sept 2001)
6. Yue Yang, Ganesh Gopalakrishnan, Gary Lindstrom: Specifying Java Thread Semantics Using a Uniform Memory Model. Joint ACM Java Grande - ISCOPE Conference (2002)
7. The SPARC Architecture Manual, Version 9, Prentice Hall (1993)
8. Leslie Lamport: How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs. IEEE Transactions on Computers, 28(9): 690–691 (1979)
9. M. Ahamad, G. Neiger, James Burns, Prince Kohli, Philip Hutto: Causal Memory: Definitions, Implementation and Programming. Technical Report: GIT-CC-93/95
10. Yuan Yu, through personal communication
11. Simplify, <http://research.compaq.com/SRC/esc/Simplify.html>
12. TLA+, <http://research.microsoft.com/users/lamport/tla/tla.html>
13. Rajeev Joshi, Leslie Lamport, John Matthews, Serdar Tasiran, Mark Tuttle, Yuan Yu: Checking Cache-Coherence Protocols with TLA+ Formal Methods in System Design. Formal Methods in System Design, 22(2): 125–131 (Mar 2003)
14. M. J. C. Gordon, T. F. Melham: Introduction to HOL: A theorem proving environment for higher order logic. Cambridge University Press (1993)
15. R. J. Lipton, J. S. Sandberg: PRAM: A Scalable Shared Memory. CS-TR-180-88
16. Mustaque Ahamad, Rida Bazzi, Ranjit John, Prince Kohli, Gil Neiger: The Power of Processor Consistency. 5th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 93) (1993)
17. J. Jaffar, J-L. Lassez: Constraint Logic Programming. Principles Of Programming Languages, Munich, Germany (January 1987)
18. SICStus Prolog, <http://www.sics.se/sicstus>
19. M. Moskewicz, C. Madigan, Y. Zhao, L. Zhang, S. Malik: Chaff: Engineering an Efficient SAT Solver. 39th Design Automation Conference, Las Vegas (June 2001)
20. E. Goldberg, Y. Novikov: BerkMin: a Fast and Robust Sat-Solver. Design, Automation and Test in Europe Conference and Exhibition Paris, France (2002)
21. Sanjit A. Seshia, Shuvendu K. Lahiri, Randal E. Bryant: A Hybrid SAT-Based Decision Procedure for Separation Logic with Uninterpreted Functions. 40th Design Automation Conference (DAC), pages 425–430 (June 2003)

Appendix: Formal Itanium Memory Ordering Specification

$\text{legal ops} \equiv \exists \text{ order.}$

$\text{requireLinearOrder ops order} \wedge \text{requireWriteOperationOrder ops order} \wedge$
 $\text{requireProgramOrder ops order} \wedge$
 $\text{requireMemoryDataDependence ops order} \wedge$
 $\text{requireDataFlowDependence ops order} \wedge \text{requireCoherence ops order} \wedge$
 $\text{requireReadValue ops order} \wedge \text{requireAtomicWBRelease ops order} \wedge$
 $\text{requireSequentialUC ops order} \wedge \text{requireNoUCBypass ops order}$

A.1 General Ordering Requirement

$\text{requireLinearOrder ops order} \equiv$

$\text{requireWeakTotal ops order} \wedge \text{requireTransitive ops order} \wedge$
 $\text{requireAsymmetric ops order}$

$\text{requireWeakTotal ops order} \equiv \forall i, j \in \text{ops. id } i \neq \text{id } j \Rightarrow (\text{order } i j \vee \text{order } j i)$

$\text{requireTransitive ops order} \equiv \forall i, j, k \in \text{ops. } (\text{order } i j \wedge \text{order } j k) \Rightarrow \text{order } i k$

$\text{requireAsymmetric ops order} \equiv \forall i, j \in \text{ops. } \text{order } i j \Rightarrow \neg(\text{order } j i)$

A.2 Write Operation Order

$\text{requireWriteOperationOrder ops order} \equiv \forall i, j \in \text{ops.}$
 $\text{orderedByWriteOperation } i j \Rightarrow \text{order } i j$

$\text{orderedByWriteOperation } i j \equiv \text{isWr } i \wedge \text{isWr } j \wedge \text{wrID } i = \text{wrID } j \wedge$
 $(\text{wrType } i = \text{Local} \wedge \text{wrType } j = \text{Remote} \wedge \text{wrProc } j = \text{p } i \vee$
 $\text{wrType } i = \text{Remote} \wedge \text{wrType } j = \text{Remote} \wedge$
 $\text{wrProc } i = \text{p } i \wedge \text{wrProc } j \neq \text{p } i)$

A.3 Program Order

$\text{requireProgramOrder ops order} \equiv \forall i, j \in \text{ops.}$
 $(\text{orderedByAcquire } i j \vee \text{orderedByRelease } i j \vee \text{orderedByFence } i j) \Rightarrow$
 $\text{order } i j$

$\text{orderedByProgram } i j \equiv \text{p } i = \text{p } j \wedge \text{pc } i < \text{pc } j$

$\text{orderedByAcquire } i j \equiv \text{orderedByProgram } i j \wedge \text{op } i = \text{ld.acq}$

$\text{orderedByRelease } i j \equiv \text{orderedByProgram } i j \wedge \text{op } j = \text{st.rel} \wedge$
 $(\text{isWr } i \Rightarrow (\text{wrType } i = \text{Local} \wedge \text{wrType } j = \text{Local} \vee$
 $\text{wrType } i = \text{Remote} \wedge \text{wrType } j = \text{Remote} \wedge \text{wrProc } i = \text{wrProc } j))$

$\text{orderedByFence } i j \equiv \text{orderedByProgram } i j \wedge (\text{op } i = \text{mf} \vee \text{op } j = \text{mf})$

A.4 Memory-Data Dependence

$\text{requireMemoryDataDependence } ops \text{ order} \equiv \forall i, j \in ops.$

$(\text{orderedByRAW } i \ j \vee \text{orderedByWAR } i \ j \vee \text{orderedByWAW } i \ j) \Rightarrow \text{order } i \ j$

$\text{orderedByMemoryData } i \ j \equiv \text{orderedByProgram } i \ j \wedge \text{var } i = \text{var } j$

$\text{orderedByRAW } i \ j \equiv$

$\text{orderedByMemoryData } i \ j \wedge \text{isWr } i \wedge \text{wrType } i = \text{Local} \wedge \text{isRd } j$

$\text{orderedByWAR } i \ j \equiv$

$\text{orderedByMemoryData } i \ j \wedge \text{isRd } i \wedge \text{isWr } j \wedge \text{wrType } j = \text{Local}$

$\text{orderedByWAW } i \ j \equiv \text{orderedByMemoryData } i \ j \wedge \text{isWr } i \wedge \text{isWr } j \wedge$

$(\text{wrType } i = \text{Local} \wedge \text{wrType } j = \text{Local} \vee$

$\text{wrType } i = \text{Remote} \wedge \text{wrType } j = \text{Remote} \wedge$

$\text{wrProc } i = p \ i \wedge \text{wrProc } j = p \ i)$

A.5 Data-Flow Dependence

$\text{requireDataFlowDependence } ops \text{ order} \equiv \forall i, j \in ops.$

$\text{orderedByLocalDependence } i \ j \Rightarrow \text{order } i \ j$

$\text{orderedByLocalDependence } i \ j \equiv \text{orderedByProgram } i \ j \wedge \text{reg } i = \text{reg } j \wedge$

$(\text{isRd } i \wedge \text{isRd } j \vee$

$\text{isWr } i \wedge \text{wrType } i = \text{Local} \wedge \text{useReg } i \wedge \text{isRd } j \vee$

$\text{isRd } i \wedge \text{isWr } j \wedge \text{wrType } j = \text{Local} \wedge \text{useReg } j)$

A.6 Coherence

$\text{requireCoherence } ops \text{ order} \equiv \forall i, j \in ops.$

$(\text{isWr } i \wedge \text{isWr } j \wedge \text{var } i = \text{var } j \wedge$

$(\text{attribute } (\text{var } i) = \text{WB} \vee \text{attribute } (\text{var } i) = \text{UC}) \wedge$

$(\text{wrType } i = \text{Local} \wedge \text{wrType } j = \text{Local} \wedge p \ i = p \ j \vee$

$\text{wrType } i = \text{Remote} \wedge \text{wrType } j = \text{Remote} \wedge \text{wrProc } i = \text{wrProc } j) \wedge$

$\text{order } i \ j)$

$\Rightarrow$

$(\forall p, q \in ops.$

$(\text{isWr } p \wedge \text{isWr } q \wedge \text{wrID } p = \text{wrID } i \wedge \text{wrID } q = \text{wrID } j \wedge$

$\text{wrType } p = \text{Remote} \wedge \text{wrType } q = \text{Remote} \wedge \text{wrProc } p = \text{wrProc } q) \Rightarrow$

$\text{order } p \ q)$

A.7 Read Value

$\text{requireReadValue } ops \text{ order} \equiv \forall j \in ops.$

$(\text{isRd } j \Rightarrow (\text{validLocalWr } ops \text{ order } j \vee \text{validRemoteWr } ops \text{ order } j \vee \text{validDefaultWr } ops \text{ order } j)) \wedge ((\text{isWr } j \wedge \text{useReg } j) \Rightarrow \text{validRd } ops \text{ order } j)$

validLocalWr $ops\ order\ j \equiv \exists i \in ops.$

$(isWr\ i \wedge wrType\ i = Local \wedge var\ i = var\ j \wedge p\ i = p\ j \wedge$
 $data\ i = data\ j \wedge order\ i\ j) \wedge$
 $(\neg \exists k \in ops. isWr\ k \wedge wrType\ k = Local \wedge var\ k = var\ j \wedge p\ k = p\ j \wedge$
 $order\ i\ k \wedge order\ k\ j)$

validRemoteWr $ops\ order\ j \equiv \exists i \in ops.$

$(isWr\ i \wedge wrType\ i = Remote \wedge wrProc\ i = p\ j \wedge var\ i = var\ j \wedge$
 $data\ j = data\ i \wedge \neg(order\ j\ i)) \wedge$
 $(\neg \exists k \in ops. isWr\ k \wedge wrType\ k = Remote \wedge var\ k = var\ j \wedge wrProc\ k = p\ j \wedge$
 $order\ i\ k \wedge order\ k\ j)$

validDefaultWr $ops\ order\ j \equiv$

$(\neg \exists i \in ops. isWr\ i \wedge var\ i = var\ j \wedge order\ i\ j \wedge$
 $(wrType\ i = Local \wedge p\ i = p\ j \vee wrType\ i = Remote \wedge wrProc\ i = p\ j)) \Rightarrow$
 $data\ j = default\ (var\ j)$

validRd $ops\ order\ j \equiv \exists i \in ops.$

$(isRd\ i \wedge reg\ i = reg\ j \wedge orderedByProgram\ i\ j \wedge data\ j = data\ i) \wedge$
 $(\neg \exists k \in ops. isRd\ k \wedge reg\ k = reg\ j \wedge$
 $orderedByProgram\ i\ k \wedge orderedByProgram\ k\ j)$

A.8 Total Ordering of WB Releases

requireAtomicWBRelease $ops\ order \equiv \forall i, j, k \in ops.$

$(op\ i = st.rel \wedge wrType\ i = Remote \wedge op\ k = st.rel \wedge wrType\ k = Remote \wedge$
 $wrID\ i = wrID\ k \wedge attribute\ (var\ i) = WB \wedge order\ i\ j \wedge order\ j\ k) \Rightarrow$
 $(op\ j = st.rel \wedge wrType\ j = Remote \wedge wrID\ j = wrID\ i)$

A.9 Sequentiality of UC Operations

requireSequentialUC $ops\ order \equiv \forall i, j \in ops. orderedByUC\ i\ j \Rightarrow order\ i\ j$

orderedByUC $i\ j \equiv$

$orderedByProgram\ i\ j \wedge attribute\ (var\ i) = UC \wedge attribute\ (var\ j) = UC \wedge$
 $(isRd\ i \wedge isRd\ j \vee$
 $isRd\ i \wedge isWr\ j \wedge wrType\ j = Local \vee$
 $isWr\ i \wedge wrType\ i = Local \wedge isRd\ j \vee$
 $isWr\ i \wedge wrType\ i = Local \wedge isWr\ j \wedge wrType\ j = Local)$

A.10 No UC Bypassing

requireNoUCBypass $ops\ order \equiv \forall i, j, k \in ops.$

$(isWr\ i \wedge wrType\ i = Local \wedge attribute\ (var\ i) = UC \wedge isRd\ j \wedge$
 $isWr\ k \wedge wrType\ k = Remote \wedge wrProc\ k = p\ k \wedge wrID\ k = wrID\ i \wedge$
 $order\ i\ j \wedge order\ j\ k) \Rightarrow$
 $(wrProc\ k \neq p\ j \vee var\ i \neq var\ j)$