Dynamic Verification of Cache Coherence Protocols

Jason F. Cantin, Mikko H. Lipasti, James E. Smith

Department of Electrical and Computer Engineering
University of Wisconsin-Madision
Madison, WI 53706

{ jcantin, lipasti, jes }@ece.wisc.edu

Abstract

A method for improving the fault-tolerance of cache coherent multiprocessors is proposed. By dynamically verifying coherence operations in hardware, errors caused by manufacturing faults, soft errors, and design mistakes can be detected. Analogous to the DIVA concept for single-processor systems, a simple version of the protocol functions as a checker for the aggressive implementation. An example implementation is shown, and the overhead is estimated for a small SMP system.

1 Introduction

Cache coherence protocols are notoriously difficult to design and verify [1]. Though a protocol description may specify only a few states (e.g., MOESI), implementations quickly become very complicated as states are added to handle the non-atomicity of memory operations, preserve correctness, and implement protocol optimizations [2]. The complexity increases the possibility of subtle errors in the specification and/or low-level implementation. Furthermore, transient failures caused by non-ideal operating environments, or cosmic rays and alpha particles interacting with very small devices, are likely to pose major reliability problems [3, 4]. Thus, tolerance against design errors and transient faults will be important for ensuring the reliability and scalability of cache coherent multiprocessor systems.

Recently, Rotenberg observed that the result of a complex computation may be checked for correctness more efficiently than it was first computed provided the check is delayed in time [5]. Austin proposed a novel approach for runtime verification of complex superscalar processors based on this principle [6]. Because the verification hardware is simple and centralized, its correctness can be easily verified. We refer to this process as dynamic verification (DV).

We propose using DV techniques to improve the fault-tolerance of cache coherent multiprocessor systems. However, a centralized check processor approach as used for single processor systems exhibiting serial semantics [6] is probably inappropriate for distributed cache coherence hardware based on parallel multiprocessor semantics. Consequently, we propose a distributed version of DV for concurrently checking cache coherence protocols during execution. As an example, we demonstrate this concept with a symmetric multiprocessor system. In this paper, we concentrate on the error detection mechanisms. We leave the integration of DV with recovery techniques for future work.

1.1 Dynamic Verification

As mentioned above, a complex computation can be checked for correctness more efficiently than it was computed in the first place, provided the check is delayed in time. The key is that the checker can exploit parallelism exposed by the original computation, and need only verify results that update the architected machine state. This allows results to be recomputed in a simpler, more efficient way.

In the single-processor case, the primary execution core (e.g. a superscalar implementation) ultimately produces a sequence of state changes <PC, reg, data> or <PC, mem address, data> that capture the semantics of the

computation. As proposed by Austin, this sequence is held in the reorder buffer (ROB) and can be passed to a check processor after any speculation has been resolved [6].

The check processor lags behind and re-executes the program. However, because of the time lag, the check processor does not need to predict branches, disambiguate addresses, or handle pipeline hazards. These dependences were identified and resolved by the execution core. Instead, the checker sees a filtered execution stream, with effectively perfect branch and value prediction. The check processor can then recompute the result of each instruction in a simpler way. After the check processor produces a result, that value is compared with the corresponding value produced by the execution processor. Hence, each instruction is dynamically verified.

The benefits of dynamic verification are the following:

- It detects hardware faults; assuming faults in the complex implementation and the checker are not correlated.
- It detects design errors in the complex implementation, assuming that the checker is correct. The check processor is simple, so verification should be straightforward.
- The level of design verification for the complex implementation can be relaxed because the checker can be used for detection of design errors

1.2 Cache Coherence

Modern multiprocessor systems are typically constructed from commodity processors with on-chip caches or cache hierarchies. Despite the replication of data in caches, it must appear to the programmer that there is one coherent memory. Cache-coherence protocols are used to efficiently maintain this illusion.

Figure 1 shows a simple example of a coherence protocol, MSI, where the states for a cache line are "Modified", "Shared", and "Invalid". When the data is not present in the cache, its state is *Invalid* (I). When a read-only copy of the line is present, the state is *Shared* (S), indicating that copies may exist elsewhere. When a single, modifiable copy is present the state is *Modified* (M), indicating that this is the most up-to-date copy of the data.

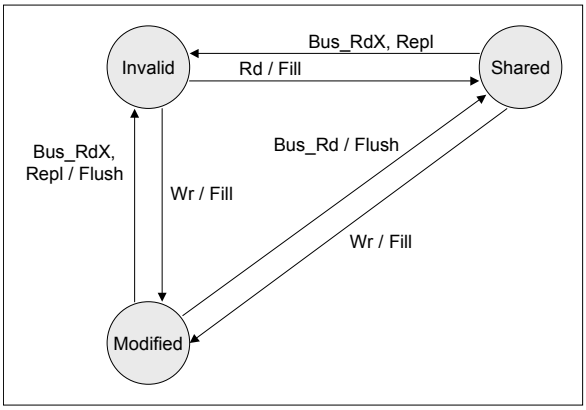

Figure 1: State Diagram of MSI Protocol (Adapted from version shown in [2])

Note that not all combinations of cache states are allowed. For example, two processors with a modifiable copy of a cache line lead to an erroneous system state. For MSI, the possible state combinations are shown in Fig. 2.

Processor A State	Processor B State		
	I	S	M
I	I	S	M
S	S	S	Error
M	M	Error	Error

Figure 2: Allowed State Combinations for MSI

This is a very simple cache coherence protocol by current standards. Real implementations require many additional states to handle non-atomic memory operations and optimizations [2]. For example, just accounting for pending write-backs complicates implementation of the protocol considerably (Figure 3).

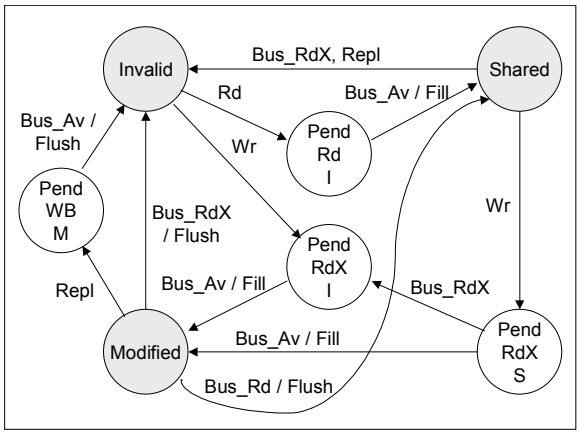

Figure 3: MSI Protocol with Transient States

These additional transient states do more than simply complicate the diagrams. They exponentially increase the state space for testing techniques, and make formal verification more cumbersome. For MSI, the addition of transient states to handle the non-atomicity of a real bus increases the verification time an order of magnitude for our experiments with the NuSMV symbolic model verification tools [7]. Random testing would also take longer, since more vectors would be necessary to achieve good coverage.

Before proceeding further, we define some useful terminology. We refer to the states used in Figure 1, as *stable* states. These states are defined in high-level descriptions of the protocol, and used to reason about interactions between processors and memory. Given atomic memory operations, the *stable* states are sufficient to correctly realize the protocol. The additional states added to implement the protocol with real hardware are referred to as *transient* states, following the convention in [2]. We refer to the state machine composed of stable states as the *simple protocol*, and the combined state machine (stable and transient states) as the *implementation protocol*. The current state of the implementation protocol is the *implementation state*, and the current state in the simple protocol is referred to as the *architected state*. Finally, we use the term *coherence transaction* to refer to the tuple consisting of the initial state, final state, input, action, and address of an event handled by the protocol.

2 Dynamic Verification of Cache Coherence

DV can be used for cache coherence. Unlike the centralized DIVA checker paradigm, a mechanism for dynamically verifying cache coherence should be logically distributed. Figure 4 shows a conceptual view. In the single-processor case, complex hardware does the computation initially, using some combination of implementation state (ROB, prediction tables, etc.) and architected state. The check processor maintains only architected state, and verifies the computations. For cache coherence, the implementation protocol initially computes the state, sends requests, and services external requests. The checker circuits maintain the architected tag state, and communicate with each other via additional messages to ensure coherence between nodes. Further, for cache coherence, the checker may be faster than the implementation, since fewer steps are required to make a transition.

Completed transitions between stable states are passed in completion order to checker circuits based on the simple protocol. This is an important detail, and warrants some elaboration. Depending on the memory model, completion order may not be program order, since the replacement of conflicting cache blocks creates a dependence between memory operations with regard to coherence. For example, a later program Load may replace a conflicting block that would be used by an earlier, not yet executed Store. If the resulting state transitions were sent to a checker in program order, an error would be falsely detected. Instead, coherence operations must be sent to the checker from the implementation protocol in the order they were completed and made visible to the program.

DV can detect subtle implementation mistakes, manufacturing faults, and transient faults in the control circuitry. Furthermore, if the checker is co-designed with the implementation protocol, it can be used to flush out specification errors early in the design phase.

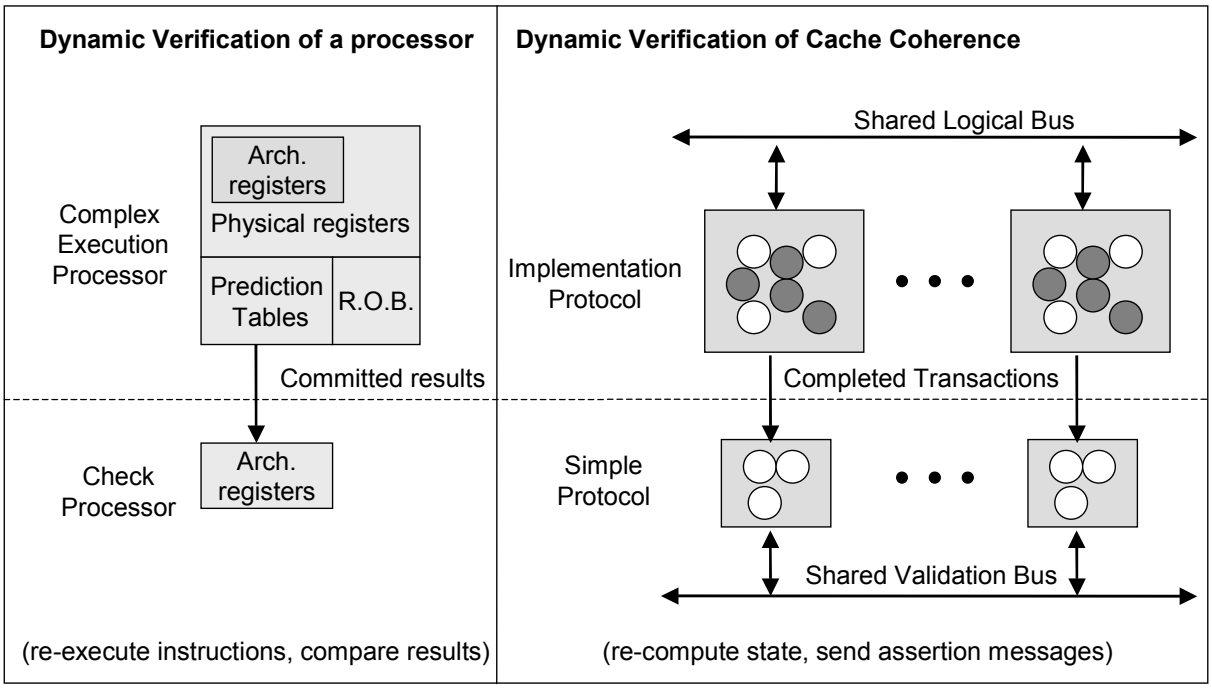

Figure 4: Conceptual View of Dynamic Verification for Cache Coherence

With this approach, the checker is implemented in a software model as part of the design verification effort, and used to check a model of the implementation protocol in simulations.

2.1 Symmetric Multiprocessor Example

To incorporate dynamic verification of coherence into a multiprocessor, a special checker circuit is placed in each node to verify the protocol actions locally (Figure 5). The checker implements a simplified version of the coherence protocol logic (i.e., no transient states), and maintains its own copy of the tags. A watchdog timer is included to detect omission failures [5].

A second logical network is used to check the protocol globally. We call this network the *validation network* to distinguish it from the main interconnection network. This network is used to broadcast final states in order to check for illegal combinations of states between nodes (e.g., two caches with modifiable copies). We refer to these additional messages as *assertions*, since the node is declaring that it has or has had certain access rights to the block. For the SMP case, this is just a second bus for addresses and final states.

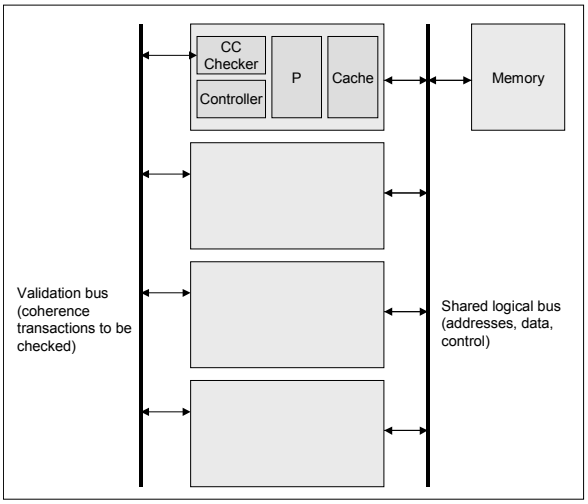

Figure 5: SMP with support for DV

2.2 SMP Coherence Checker Operation

Following a network transaction in the implementation protocol, the address, command, initial and final stable states are sent to the local checkers. Each checker re-computes the final state of the cache line and compares it to the implementation protocol's result. If the final states do not match, an error has occurred in the implementation.

The checker also performs a tag-lookup with the address to get the architected state of the cache line. The architected state stored in the checker must match the initial state reported by the implementation protocol. Disagreement signals an error. See Figures 6 and 7 for a simplified checker datapath.

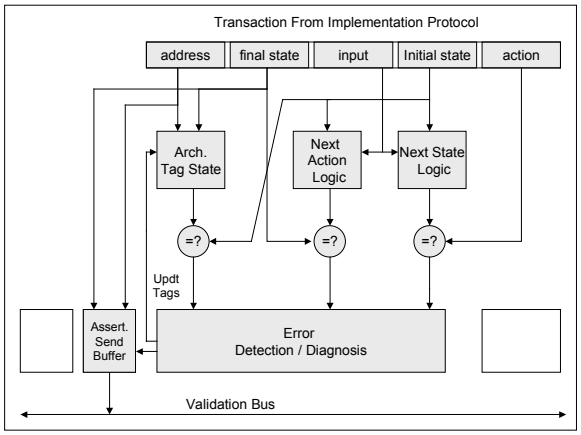

Figure 6: Coherence Checker Logic (checking a transaction).

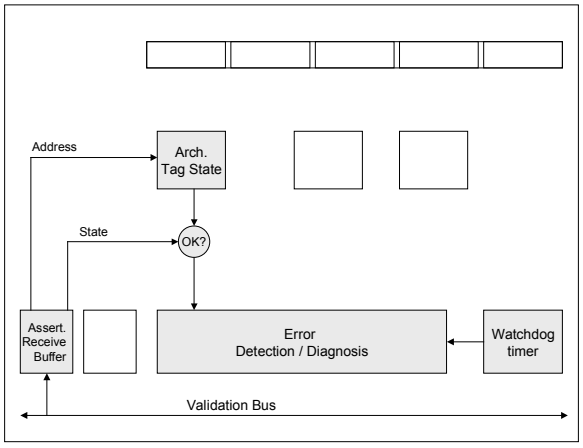

Figure 7: Coherence Checker Logic (checking an assertion).

Once the transaction has been verified locally, the cache states must be checked globally. The node that initiated the transaction (via a request) broadcasts the final state and address of the cache block over the validation network. The other nodes snoop the network and determine if the broadcasted cache state conflicts with the state of a cached copy they hold. For example, if a node acquires a modifiable copy of a memory block, it sends an assertion message indicating that it has the block in "M" state. The other checkers must determine if they have any shared or modified copies still in their caches. If an illegal combination of cache states is detected for an address (Figure 2), an error is signaled.

Once the node that provided data for the transaction (the receiver) sees the assertion, its checker knows that the transaction has completed. The transaction may then be retired and removed from any queues. Note: we do not allow further updates to the architected state of the receiving node until a corresponding check message is received. Depending on whether or not recoverability is desired, the initiator can retire the transaction after sending the assertion, or wait to make sure that no errors are signaled.

3 Evaluation of SMP Coherence Checker Implementation

Design of the coherence checker is in progress, however we can reason about its effectiveness and performance. Ideally, a checker implementation should have full fault coverage. By full coverage, we mean complete detection of faults that have propagated to the point of being visible to the coherence checker (e.g., a stable state transition that violates coherence is made by the implementation protocol). For example, a design error may cause the omission of an invalidate message in the implementation protocol, however this will not be detected by our scheme until it results in an improper stable state transition.

In addition, the coherence checking hardware should not slow down the system by introducing too much overhead, signaling too many false positives, or lagging too far behind the implementation for efficient recovery.

3.1 Coherence Checker Coverage and Specificity

Symbolic model verification (model checking) is a powerful technique for verifying finite state machines and protocols [1, 7, 8, 9]. It has been successfully used to verify cache coherence protocols [8, 9]. Given a model of a system and a set of logical properties, a tool can automatically determine if the modeled system satisfies the

properties in all cases. We can use symbolic model checking techniques to verify the correctness of the simple protocol, determine the checker implementation's coverage, and determine if the checker implementation incorrectly signals an error (a false positive).

Determining if the simple protocol is correct is straightforward. A model of the simplified state machine can be written in a language such as SMV (top part of Figure 8). A set of necessary conditions for maintaining coherence is then specified formally in temporal logic (e.g., CTL). The model checking software determines if the model always meets the conditions, or produces a counter-example. Without the complexity of the full implementation, the state space will be relatively small and quickly searched by the model verifier such as NuSMV [7].

To determine if the coherence checker implementation detects errors, we write a detailed model for the checker implementation (middle part of Figure 8). Next, we formally define (in CTL) what errors the coherence checker implementation should detect. The checker implementation achieves full coverage if, for every defined error condition, the coherence checker signals an error.

To determine specificity, we can also use model checking (bottom part of Figure 8). We can combine the two models mentioned above such that the simple protocol is checked by the implementation of the coherence checker. Since the simple protocol has been proven correct, the coherence checker implementation should never detect an error in this configuration. Any errors signaled by the coherence checker in this configuration are considered false positives.

This is analogous to proving the correctness of a DIVA checker in the single processor case, however the burden of defining the necessary conditions for correctness is placed on the designer. It is also necessary for the designer to validate the model. This is true in general for verifying concurrent systems with model checking (a design may still be incorrect, but satisfy the designers specifications). Further, model checking can only tell us if the checker implementation will detect errors specified by the designer, such as disagreement between the implementation protocol and the simple protocol.

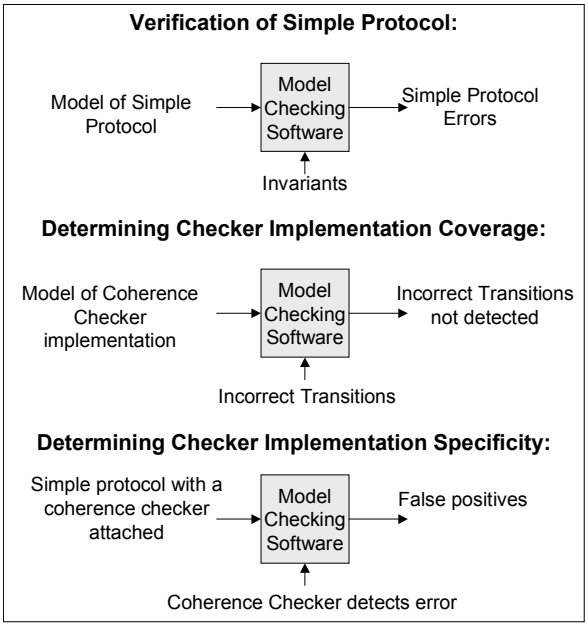

Figure 8: Determining Checker Coverage and Specificity

3.2 Coherence Checker Overhead

Conceptually, the coherence checker implementation proposed should not become a bottleneck for the base system, or increase system cost excessively. The absence of transient states keeps the checker logic simple and fast. Further, the second bus proposed for the SMP configuration produces no more transactions than the main address bus, so duplicating the address portion of the bus may be sufficient to support the extra messages in that case. However, if a second physical network is infeasible, the main network may be used to send extra messages (with low priority). If this is the case, bandwidth overhead is incurred by the extra messages used for DV.

To estimate overhead, we collected data from a 4-processor SMP system with 1MB 4-way set associative L2 caches and 64-byte lines. For simplicity, an MSI protocol is used, though the results do not change significantly for a MOESI protocol. From this data, we can determine (relatively) how often the checker must verify a

coherence operation. Figure 9 shows (for five benchmarks) that between 0.54% and 7.17% of memory references (loads and stores) result in a change of architected state for a cache line. We infer from this that the checking of local state transitions is infrequent and need not be fast.

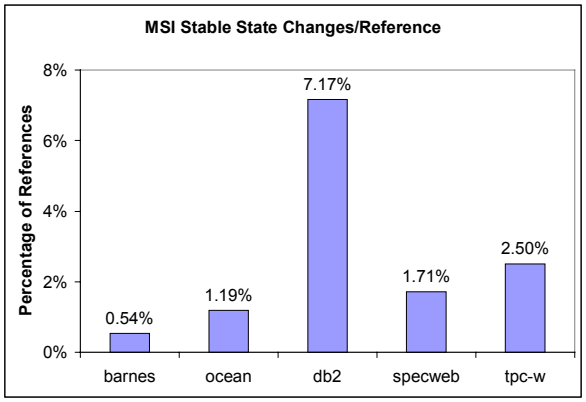

Figure 9: Percentage of Memory References that Result in Stable State Transitions

The estimated overhead incurred by extra messages is calculated for several strategies below (See Figure 10). For each of these strategies, the transactions they check are shown in Figure 11.

The first approach is to have the node's checker send a message each time that a modifiable copy of a cache line is acquired. This ensures that no other shared or modified copies exist. This simple approach requires 31% or fewer extra messages for the benchmarks simulated, but cannot detect certain types of errors. Cache lines brought into a node in the shared state are not checked to see if modified copies still exist, and replacements are not checked.

A second approach is to have check messages sent for all transactions resulting in a data transfer. This checks that data is brought into a node in the correct state, but does not check upgrades (S→M). This approach requires 45% or fewer additional messages.

Third, assertions may be issued for all bus read or upgrade operations. This approach is a combination of the first two, and has improved coverage. Each time a block is brought into the cache or upgraded, a message is sent by the initiator of the transaction. However, this is even more costly than the previous methods (as much as 64%).

Finally, an assertion message may be sent for all bus transactions. In addition to all cases checked by the third approach, writebacks resulting from replacements are also checked. In effect, the second bus proposed for checking purposes mirrors the main bus, ensuring that all transactions were completed correctly. This incurs 100% address bandwidth overhead (same overhead as replicating the main address bus, but with the ability to catch design errors).

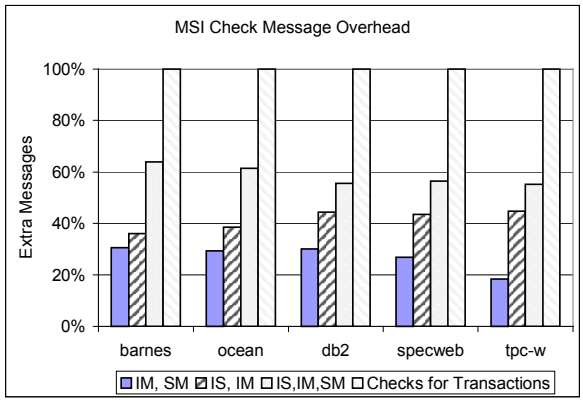

Figure 10: Calculated Message Overhead for Checker Implementation Strategies

Initiator Trans.	Remote State			Bus	Check Message Coverage			
	I	S	M	Trans.	Scheme 1	Scheme 2	Scheme 3	Scheme 4
I --> S	Ok	Ok	Error	Yes		Checked	Checked	Checked
I --> M	Ok	Error	Error	Yes	Checked	Checked	Checked	Checked
S --> M	Ok	Error	--	Yes	Checked		Checked	Checked
S --> I (Repl)	Ok	Ok	--	No				
M --> I (Repl)	Ok	--	--	Yes				Checked

Figure 11: Coverage Provided By Check Messages

None of the strategies mentioned checks the transition from *Shared* to *Invalid* that occurs when a cache line is replaced. This is a silent transition that does not involve updating memory, or sending data to other nodes. The data is simply discarded, and we rely on the checker internal to the node to make sure that the state transition takes place.

4 Related Work

Testing-based approaches are conventionally used for detecting both design errors and fabrication defects [10]. For design verification, the most commonly used technique is to develop a set of test vectors and use them to drive logic simulation [11, 12, 13]. The designer or a verification engineer may devise such test vectors. Parts of this process may be automated, but its effectiveness depends on the insight and skill of the engineers involved. This method is extremely time-consuming, and typically missed some bugs.

For hardware failures, test patterns are generated to exercise the hardware, often using a fault model such as the stuck-at model [10]. The process can be largely automated [14, 15], and test coverage can be quantitatively estimated. The test patterns are then applied to verify correctness –up to the level of test coverage. In the field, this technique works better for permanent faults than for transient ones, as the fault must be present at the time the test is applied. Also, this method may require some downtime when the test is applied.

There have been a number of proposals for using a simplified "watchdog" processor [16] to check a main processor. Watchdog processors have many of the advantages we envision, but watchdog processors do not duplicate the entire computation. They check only certain aspects of the computation, for example the control flow [17], the memory access behavior [18], and "reasonableness" based on programmed-in assertions [19]. They do not check the entire computation, and therefore do not detect all faults, nor be used for complete state recovery.

Rotenberg proposed a multithreaded processor that implements a form of time redundancy where a computation thread is re-executed later in time and the results of the two thread executions are compared [5]. His approach focused on transient hardware faults in the multithreaded processor's datapath, and also built on previous approaches using time-shifted redundant execution [20, 21]. He referred to this new technique as Active-stream / Redundant-stream Simultaneous Multithreading (AR-SMT).

Reinhardt and Mukherjee further explored the use of multithreading for transient fault detection in [22]. They introduced an important abstraction for simultaneous and redundantly multithreaded (SRT) processors, identified some key implementation challenges, and suggested some microarchitectural solutions.

As described earlier, Austin proposed dynamic checking with a separate check processor for the second computation [6].

Conventional forms of dynamic checking have been proposed and implemented for many years. Probably the oldest is replication with comparison checking as protection against hardware failures [23, 24, 25, 26, 27]. This method can be effective against both permanent and transient hardware errors, but it does not catch design errors. Furthermore, it is likely to be more expensive than the dynamic inductive checking method, because the check processor is a complete replica and is not simplified. Many systems have used replication for failure protection. The IBM G5 [28] is a recent version where both processors are on the same chip.

For detecting design errors, formal methods [29, 30] provide an alternative to conventional simulation-based testing. Formal methods typically use an architecture specification and an implementation specification, and then show the two are equivalent. This equivalence is essentially proven for all possible computations, either via model checking [1, 8, 9], theorem proving [31], or a combination [32, 33]. As high-performance implementations of coherence protocols become more complex, the computational complexity of formal methods becomes an issue.

5 Future Work

In future work, we will refine our implementation with data obtained from model checking and detailed timing simulations. Detailed simulations will determine the actual overhead of

the check processor implementation for commercial workloads.

Though a bus-based SMP system with a simple protocol was used here to illustrate the concepts, more complex protocols (such as MOESI) and scalable directory-based coherence schemes will be explored. We intend to develop a framework for checker design, verification, and performance evaluation to facilitate the process of incorporating DV into parallel systems.

In the DIVA approach for single-processor systems, the check hardware had a well-defined sequence of operations to check via the reorder buffer. Unfortunately, such a serialization is not present for coherence operations in a multiprocessor that is not sequentially consistent, since program loads can (correctly) become visible to the system before earlier stores. In future work we will investigate DV for the memory model itself, and define what constraints must be placed on the implementation protocol in order to provide a simple interface to a checker.

Finally, we intend to combine DV with hardware and software recovery techniques. Once an error has been detected and diagnosed, it may be possible to restart from a checkpoint or use some form of forward error recovery.

6 Conclusions

With dynamic verification, errors in a cache coherence protocol caused by manufacturing faults, soft errors, and design mistakes can be detected at run-time. Since most memory operations do not cause a change in cache state, a simple checker can check the coherence protocol of an aggressive processor. Further, with a second network for assertions, globally verifying cache coherence does not place pressure on the data network or memory. This approach can be combined with recovery techniques, and methods of dynamically verifying program execution [6] to produce fault-tolerant multiprocessor systems.

7 Acknowledgements

We thank Timothy Heil, Ashutosh Dhodapkar, Sebastien Nussbaum, and Tejas Karkhanis for comments on drafts of this paper. This work is supported by NSF grant CCR-0083126 and by IBM. Jason Cantin is supported by a Wisconsin Distinguished Graduate Fellowship.

8 References

[1] D. L. Dill, A. J. Drexler, A. J. Hu, and C. H. Yang. “Protocol Verification as a Hardware Design Aid.” *International Conference on Computer Design, VLSI in Computers and Processors*, Oct. 1992: 522-525.

[2] D. E. Culler and J. P. Singh, *Parallel Computer Architecture: A Hardware / Software Approach*, San Francisco, CA: Morgan Kaufmann Publishers Inc., 1999.

[3] T. C. May and M. H. Woods. “Alpha-Particle-Induced Soft Errors in Dynamic Memories.” *IEEE Transactions on Electronic Devices*, 26(2), 1979.

[4] T. J. O’Gorman, J. M. Ross, A.H. Taber, J.F. Ziegler, H.P. Muhlfeld, C.J. Montrose, H.W. Curtis, J.L. Walsh, “Field Testing for Cosmic Ray Soft Errors in Semiconductor Memories.” *IBM Journal of Research and Development*, Jan. 1996: 41-49.

[5] E. Rotenberg. “AR-SMT: A Microarchitectural Approach to Fault-Tolerance in Microprocessors.” *Proceedings of the 29th International Symposium on Fault-Tolerant Computing*, June 1999: 84-91.

[6] T. Austin. “DIVA: A Reliable Substrate for Deep-Submicron Processor Design.” *Proceedings of the 32nd Annual ACM/IEEE International Symposium on Microarchitecture*, Dec. 1999: 196-207.

[7] A. Cimatti, E. Clarke, F. Giunchiglia, and M. Roveri, "NuSMV: A New Symbolic Model Verifier." N. Halbwachs and D. Peled, eds. *Proceedings of the 11th International Conference on Computer-Aided Verification*, Lecture Notes in Computer Science 1633, Springer Verlag, 1999: 495-499.

[8] E. Clarke, O. Grumberg, H. Hiraishi, S. Jha, D. E. Long, K. L. McMillan, and A. L. Ness. "Verification of the Futurebus+ Cache Coherence Protocol." *Proceedings of the 11th International Symposium on Computer Hardware Description Languages and their Applications*. Apr. 1993.

[9] E. Clarke, O. Grumberg, and D. Peled, *Model Checking*. Cambridge, MA: MIT Press, 1999.

[10] M. Abramovici, M. A. Breuer, and A. D. Friedman. *Digital Systems Testing and Testable Design*. New York: IEEE Press, 1992.

[11] T. Sasaki, A. Yamada, and T. Aoyama. "Hierarchical Design Verification for Large Digital Systems." *Proceedings of the 18th Design Automation Conference*, June 1981: 105-112

[12] M. Monachino. "Design Verification System for Large-scale LSI Designs." *Proceedings of the 19th Design Automation Conference*, June 1982: 83-90.

[13] A. Aharon. "Test Program Generation for Functional Verification of PowerPC Processors in IBM." *Proceedings of the 32nd Design Automation Conference*, June 1995: 279-285.

[14] J. P. Roth, W. G. Bouricius, and P. R. Schneider. "Programmed Algorithms to Compute Tests and Detect and Distinguish Between Failures in Logic Circuits." *IEEE Transactions on Electronic Computers*, Oct. 1967, EC-16(10):567-579.

[15] P. Goel. "An Implicit Enumeration Algorithm to Generate Tests for Combinational Logic Circuits." *IEEE Transactions on Computers*, Mar. 1981, C-30(3):215-222.

[16] A. Mahmood and E. J. McCluskey. "Concurrent Error Detection Using Watchdog Processors –A Survey." *IEEE Transactions on Computers*, Feb. 1982, C-37(2):160-173.

[17] D. J. Liu. "Watchdog Processor and Structural Integrity Checking." *IEEE Transactions on Computers*, July 1982, C-31: 681-685.

[18] M. Namjoo and E. J. McCluskey. "Watchdog Processors and Capability Checking." *Proceedings of the 12th International Symposium on Fault-Tolerant Computing*, June 1982: 245-248.

[19] A. Mahmood, D. J. Liu, and E. J. McCluskey. "Concurrent Detection Using a Watchdog Processor and Assertions." *Proceedings of the 1983 International Test Conference*, Oct. 1983: 622-628.

[20] J. Patel and L. Fung. "Concurrent Error Detection in ALUs by Recomputing with Shifted Operands." *IEEE Transactions on Computers*, July 1982, C-31(7): 589-595.

[21] G. Sohi, M. Franklin, and K. Saluja. "A Study of Time-Redundant Fault-Tolerance Techniques in High-Performance Pipelined Computers." *Proceedings of 19th Fault-Tolerant Computing Symposium*, June 1989: 436-443.

[22] S. Reinhardt and S. Mukherjee. "Transient Fault Detection via Simultaneous Multithreading." *Proceedings of The 27th International Symposium on Computer Architecture*, June 2000: 25-36.

[23] S. Webber. "The Stratus Architecture." D. Siewiorek and R. Swarz, eds. *Reliable Computer Systems: Design and Evaluation*. Bedford, MA: Digital Press, 1992.

[24] O. Serlin. "Fault-Tolerant Systems in Commercial Applications." *IEEE Computer*, Aug. 1984: 19-30.

[25] Katzman, J. A., "A Fault-Tolerant Computing System." Tandem Computers, Inc., Cupertino, CA, 1977.

[26] J.P. Eckert Jr., J. R. Weiner, H. D. Welsh, and H. F. Mitchell. "The UNIVAC system", *Proceedings of the Joint AIEE-IRE Computer Conference*, Dec. 1951: 6-16.

[27] A. W. Burks, H. H. Goldstein, and J. von Neumann, "Preliminary Discussion of the Logical Design of an Electronic Computing Instrument." *Papers of John von Neumann*. Cambridge, MA, MIT Press, 1987. 97-146.

[28] L. Spainhower and T. A. Gregg. "IBM S/390 Parallel Enterprise Server G5 Fault-Tolerance: A Historical Perspective." *IBM Journal of Research and Development*, May 1999, 43(5/6): 863.

[29] E. M. Clarke and J. M. Wing. "Formal Methods: State of the Art and Future Directions." *ACM Computing Surveys*, Dec. 1996, 28(4): 626-643.

[30] E. Clarke and R. Kurshian. "Computer-Aided Verification." *IEEE Spectrum*, June 1996, 33(6): 61-67.

[31] A. Kuehlman, A. Srinivasan, and D. LaPotin. "Verity --A Formal Program for Custom CMOS Circuits." *IBM Journal of Research and Development*, 1995, 39(1/2): 149-165.

[32] R. Kurshan and L. Lamport. "Verification of a Multiplier, 64 Bits and Beyond." *Proceedings of the 5th International Conference on Computer-Aided Verification*, Lecture Notes in Computer Science, Springer Verlag, (1993): 166-179.

[33] S. Rajan, N. Shankar, and M. Srivas. "An Integration of Model Checking with Automated Proof Checking." *Proceedings of the 7th International Conference on Computer-Aided Verification*, Lecture Notes in Computer Science 939, Springer Verlag, June 1995: 84-97.