

MIT Open Access Articles

Modular Deductive Verification of Multiprocessor Hardware Designs

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Vijayaraghavan, Muralidaran et al. "Modular Deductive Verification of Multiprocessor Hardware Designs." CAV 2015: Computer Aided Verification (July 2015): 109-127 © 2015 Springer International Publishing Switzerland

As Published: http://dx.doi.org/10.1007/978-3-319-21668-3_7

Publisher: Springer International Publishing

Persistent URL: <https://hdl.handle.net/1721.1/122607>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Modular Deductive Verification of Multiprocessor Hardware Designs

Muralidaran Vijayaraghavan¹, Adam Chlipala¹, Arvind¹, and Nirav Dave²

¹ MIT {vmurali, adamc, arvind}@csail.mit.edu
² SRI International ndave@csl.sri.com

Abstract. We present a new framework for modular verification of hardware designs in the style of the Bluespec language. That is, we formalize the idea of components in a hardware design, with well-defined input and output channels; and we show how to specify and verify components individually, with machine-checked proofs in the Coq proof assistant. As a demonstration, we verify a fairly realistic implementation of a multicore shared-memory system with two types of components: memory system and processor. Both components include nontrivial optimizations, with the memory system employing an arbitrary hierarchy of cache nodes that communicate with each other concurrently, and with the processor doing speculative execution of many concurrent read operations. Nonetheless, we prove that the combined system implements sequential consistency. To our knowledge, our memory-system proof is the first machine verification of a cache-coherence protocol parameterized over an arbitrary cache hierarchy, and our full-system proof is the first machine verification of sequential consistency for a multicore hardware design that includes caches and speculative processors.

1 Introduction

A modern high-performance, cache-coherent, distributed-memory hardware system is inherently complex. These systems by their nature are highly concurrent and nondeterministic. The goal of this work is to provide a framework for full verification of such complex hardware systems.

Modularity has long been understood as a key property for effective design and verification in this domain. We want to break a system into pieces that can be specified and verified independently. In our design, processors and memory systems both independently employ intricate optimizations that exploit opportunities for parallelism. We are able to prove that each of those two main components still provides strong enough guarantees to support SC, and then we compose those theorems into a result for the full system. Either component may be optimized further without requiring any changes to the implementation, specification, or proof of the other. Our concrete optimizations include speculation in processors and using a hierarchy of caches in memory.

We thus present **the first mechanized proof of correctness of a realistic multiprocessor, shared-memory hardware system**, including **the first mechanized correctness proof of a directory-based cache-coherence protocol for arbitrary cache hierarchies**, *i.e.*, the proof is parameterized over an unknown number of processors connected to an unknown number of caches in an unknown number of levels (*e.g.*,

L1, L2). Our proof has been carried out in the Coq proof assistant and is available at <http://github.com/vmurali/SeqConsistency>. Since our technique is based on proof assistants, the computational complexity remains constant for any choice of parameters. In the process, we introduce **a methodology for modular verification of hardware designs**, based on the theory of labeled transition systems (LTSes).

LTSes as hardware descriptions are an established idea [21, 22, 6], and there are compilers that convert LTSes into efficient hardware. Our work is based on the Bluespec language [7, 10], whose semantics match the formalism of this paper. Bluespec models hardware components as atomic rules of a transition system over state elements, and there is a commercial compiler to synthesize such codes into circuits (*i.e.*, Verilog code) with competitive performance. Our cache-coherent memory system is directly transliterated from a Bluespec implementation [17] used to implement a cycle-accurate simulator for a cache-coherent multiprocessor PowerPC system [27]. The hardware synthesized from that implementation is rather efficient: an 8-core system with a 2-level cache hierarchy can run 55 million instructions per second on the BEE FPGA board [14].

Our high-level agenda in this work is, first, to import to the hardware-verification domain good ideas from the worlds of programming-language semantics and formal software verification; and, second, to demonstrate some advantages of human-guided deductive techniques over model-checking techniques that less readily support modularity and generalization over infinite families of systems, and which may provide less insight to hardware designers (*e.g.*, by not yielding human-understandable invariants about systems).

Paper Organization: We begin with a discussion of related work in Section 2. In Section 3 we introduce our flavor of the labeled transition systems formalism, including a definition of trace refinement. In Section 4, we show a generic decomposition of any multiprocessor system independently of the memory model that it implements, and we discuss the store-atomicity property of the memory subcomponent. In Section 5 we give a simple formal model of sequential consistency. The following sections refine the two main subcomponents of our multiprocessor system. Section 7 discusses definition and verification of a speculative processor model, and Section 8 defines and proves our hierarchical cache-coherence protocol. Finally, in Section 9 we show the whole-system modular proof of our complex system, and we end with some conclusions in Section 10.

2 Related Work

Hardware verification is dominated by model checking; for instance, processor verification [12, 32] and more recently, Intel’s execution cluster verification [26]. Many abstraction techniques are used to reduce designs to finite state spaces, which can be explored exhaustively. There are limits to the construction of sound abstractions, so verifications of, *e.g.*, cache-coherence protocols have mostly treated systems with concrete topologies, involving particular finite numbers of caches and processors. For instance, explicit-state model checking tools like Murphi [19] or TLC [29, 25] are only able to handle single-level cache hierarchies with fewer than ten addresses and ten CPUs, as opposed to the billions of addresses in a real system, or the ever-growing number of CPUs. Symbolic model-checking techniques have fared better: McMillan *et al.* have

verified a 2-level MSI protocol based on the Gigamax distributed multiprocessor using SMV [34]. Optimizations on these techniques (*e.g.*, partial order reduction [8], symmetry reduction [9, 23, 16, 15, 20, 43], compositional reasoning [33, 31, 24], extended-FSM [18]) also scale the approach, verifying up to 2 levels of cache hierarchy, but they are still unable to handle multi-level hierarchical protocols. In fact, related work by Zhang et al. [43] insists that parameterization should be restricted to single dimensions for the state-of-the-art tools to scale practically. In all these cases, finding invariants automatically is actually hard. Chou et al. [16] require manual insertion of extra invariants, called “non-interference lemmas”, to eliminate counter-examples that violate the required property. Flow-based methodology [41] gives yet another way of manually specifying invariants. In general, we believe that the level of complexity of the manually specified invariants between those approaches and ours is similar. Moreover, we might hope to achieve higher assurance and understanding of design ideas by verifying *infinite families* of hardware designs, which resist reduction to finite-state models. Past work by Zhang et al [43] has involved model-checking hierarchical cache-coherence protocols [44], with a restriction to *binary* trees of caches only, relying on paper-and-pencil proofs about the behavior of fractal-like systems. The authors agree that, as a result, the protocol suffers from a serious performance handicap. Our cache protocol in this paper is chosen to support more realistic performance scaling.

Theorem provers have also been used to verify microprocessors, *e.g.*, HOL to verify an academic microprocessor AVI-1 [42]. Cache-coherence proofs have also been done with mechanized theorem provers, though all previous work has verified only single-level hierarchies. Examples include using ACL2 for verifying a bus-based snoop protocol [35], using a combination of model-checking and PVS [36] to verify the FLASH protocol [28], and using PVS to mechanize some portions of the paper-and-pencil proof verifying that the Cachet cache-coherence protocol [40] does not violate the CRF memory model. The first two of these works do not provide insights that can be used to design and verify other protocols. The last falls short of proving a “full functional correctness” property of a memory system. In this paper, we verify that property for a complex cache protocol, based on human-meaningful invariants that generalize to related protocols.

Before verifying that a hardware design is correct, it is necessary to pin down the intended semantics. While our case study in this paper hews to the uncontroversial sequential-consistency model, much recent work has studied popular hardware platforms empirically to derive their memory models, and to provide tools to verify a multi-threaded program’s behavior in the memory model [1–5, 30, 37–39].

3 Labeled Transition Systems

We make extensive use of the general theory of labeled transition systems, a semantics approach especially relevant to communicating concurrent systems. As we are formalizing processors for Turing-complete machine languages, it is challenging to prove that a system preserves almost any aspect of processor behavior from a model like SC. To focus our theorems, we pick the time-honored property of *termination*. An optimized system should terminate or diverge iff the reference system could also terminate or di-

verge, respectively. All sorts of other interesting program properties are reducible to this one, in the style of computability theory. Our basic definitions of transition systems build in special treatment of halting, so that we need not mention it explicitly in most contexts to come.

Definition 1 A *labeled transition system (LTS)* is a ternary relation, over $\mathcal{S}^H \times \mathcal{L}^\epsilon \times \mathcal{S}^H$, for some sets \mathcal{S} of states and \mathcal{L} of labels. We usually do not mention these sets explicitly, as they tend to be clear from context. We write X^ϵ for lifting of a set X to have an extra “empty” element ϵ (like an `option` type in ML). We write X^H for lifting of a set X to have an extra “halt” element H . We also implicitly consider each LTS to be associated with an initial state in \mathcal{S} .

For LTS A , we write $(s) \xrightarrow[A]{\ell} (s')$ as shorthand for $(s, \ell, s') \in A$, and we write A_0 for A 's initial state. The intuition is that A is one process within a concurrent system. The label ℓ from set \mathcal{L} of labels is produced when A participates in some IO exchange with another process; otherwise it is an empty or “silent” label ϵ . As a shorthand, we sometimes omit labels for ϵ steps.

3.1 Basic Constructions on LTSes

From an LTS representing single-step system evolution, we can build an LTS capturing arbitrary-length evolutions.

Definition 2 The *transitive-reflexive closure* of A , written A^* , is a derived LTS. Where A 's states and labels are \mathcal{S} and \mathcal{L} , the states of A^* are \mathcal{S} , and the labels are \mathcal{L}^* , or sequences of labels from the original system. A^* steps from s to s' when there exist zero or more transitions in A that move from s to s' . The label of this transition is the concatenation of all labels generated in A , where the empty or “silent” label ϵ is treated as an identity element for concatenation.

We also want to compose n copies of an LTS together, with no explicit communication between them. We apply this construction later to lift a single-CPU system to a multi-CPU system.

Definition 3 The *n -repetition* of A , written A^n , is a derived LTS. Where A 's states and labels are \mathcal{S} and \mathcal{L} , the states of A^n are \mathcal{S}^n , and the labels are $[1, n] \times \mathcal{L}$, or pairs that tag labels with which component system generated them. These labels are generated only when the component system generates a label. The whole system halts whenever one of the components halts.

Eventually, we need processes to be able to communicate with each other, which we formalize via the $+$ composition operator. It connects same-label transitions in the two systems, treating the label as a cooperative communication event that may now be hidden from the outside world, as an ϵ label.

Definition 4 Where A and B are two LTSes sharing labels set \mathcal{L} , and with state sets \mathcal{S}_A and \mathcal{S}_B respectively, the *communicating composition* $A + B$ is a new LTS with states $\mathcal{S}_A \times \mathcal{S}_B$ and an empty label set, defined as follows:

$$\begin{array}{c}
A \frac{(a) \xrightarrow{A} (a') \quad a' \neq H}{(a, b) \xrightarrow{A+B} (a', b)} \quad B \frac{(b) \xrightarrow{B} (b') \quad b' \neq H}{(a, b) \xrightarrow{A+B} (a, b')} \quad H_A \frac{(a) \xrightarrow{A} (H)}{(a, b) \xrightarrow{A+B} (H)} \\
H_B \frac{(b) \xrightarrow{B} (H)}{(a, b) \xrightarrow{A+B} (H)} \quad \text{Join} \frac{(a) \xrightarrow{A} (a') \quad (b) \xrightarrow{B} (b') \quad a', b' \neq H}{(a, b) \xrightarrow{A+B} (a', b')}
\end{array}$$

3.2 Refinement Between LTSes

We need a notion of when one LTS “implements” another. Intuitively, transition labels and halting are all that the outside world can observe. Two systems that produce identical labels and termination behavior under all circumstances can be considered as safe substitutes for one another. We will only need an asymmetrical notion of compatibility:

Definition 5 For some label domain \mathcal{L} , let $f : \mathcal{L} \rightarrow \mathcal{L}^\epsilon$ be a function that is able to replace labels with alternative labels, or erase them altogether. Let LTSes A and B have the same label set \mathcal{L} . We say that A **trace-refines** B **w.r.t.** f , or $A \sqsubseteq_f B$, if:

$$\forall s_A, \eta. (A_0) \xrightarrow{A^*} (s_A) \Rightarrow \exists s_B. (B_0) \xrightarrow{B^*} (s_B) \wedge (s_A = H \Leftrightarrow s_B = H)$$

Each label in the trace is replaced by the mapping of f on it, and labels mapped to ϵ by f are dropped. f is overloaded to denote the multilabel version when applied to η .

As a shorthand, we write $A \sqsubseteq B$ for $A \sqsubseteq_{\text{id}} B$, for id an identity function, forcing traces in the two systems to match exactly. Under this notion of identical traces, we say that A is sound w.r.t. B .

3.3 A Few Useful Lemmas

We need the following theorems in our proof.

Theorem 1. \sqsubseteq is reflexive and transitive.

Theorem 2. If $A \sqsubseteq_f B$, then $A^n \sqsubseteq_{f^n} B^n$, where f^n is f lifted appropriately to deal with indices ($f^n(i, \ell) = (i, \ell')$ when $f(\ell) = \ell'$, and $f^n(i, \ell) = \epsilon$ when $f(\ell) = \epsilon$).

Theorem 3. If $A \sqsubseteq_f A'$ and $B \sqsubseteq_f B'$, then $A + B \sqsubseteq A' + B'$. In other words, if systems A and B are individually simulated by A' and B' on identical alteration of the traces, then the composed system $A + B$ will be sound with respect to $A' + B'$.

4 Decomposing a Shared-Memory Multiprocessor System

Any conventional multiprocessor system can be divided logically into 3 components as shown in Figure 1. The top-level system design is shown at top right, while the details of its components, the memory system and the processor (P_i), are shown in the magnified boxes. The processor component P_i can be implemented in a variety of ways, from one executing instructions one-by-one in program order, to a complex one speculatively executing many instructions concurrently to exploit parallelism. The memory component is normally implemented using a hierarchy of caches, in order to




Fig. 1: Components of a multiprocessor system

increase the performance of the overall system, since the latency of accessing memory directly is large compared to that of accessing a much smaller cache. The component between each processor and the global memory subsystem contains local buffers, LB_i , each specific to a processor P_i . We use local buffers to stand for any structure between the core and the memory subsystem.

Popular ISAs, such as Intel x86, ARM, and PowerPC, do not guarantee sequential consistency. However, we want to emphasize that, in every weak-memory system we are aware of, *the main memory still exposes atomic loads and stores!* Weaker semantics in a core P_i arise only because of (1) reordering of memory instructions by the core and/or (2) the properties of the local buffers LB_i connected to P_i .




Fig. 2: LTS for a simple memory

So, we focus on this opportunity to simplify proof decomposition. We will prove that our main memory component satisfies an intuitive *store atomicity* property which is an appropriate specification of the memory component even for implementations of weaker memory models. Store atomicity can be understood via the operational semantics of Figure 2, describing an LTS that receives load and store requests (Ld and St) from processors and sends back load responses (LdRp). The transfer happens via input buffers $ins(p)$ from processor p and output buffers $outs(p)$ to processor p . Note that this model allows the memory system to answer pending memory requests in any order (as indicated by the bag union operator \uplus), even potentially reordering requests from a single processor, so long as, whenever it does process a request, that action appears to take place *atomically*.

The memory component is composed of a *hierarchy of caches*, with cache nodes labeled like “L1,” “L2,” etc., to avoid the latency of round trips with main memory. Therefore, it is the responsibility of the hierarchy of caches, which forms the memory subcomponent, to implement the store atomicity property. In fact, as we will prove in Section 8, the purpose of the cache-coherence protocol is to establish this invariant for the memory subcomponent. Concretely, we have verified a *directory-based* proto-

col for coordinating an arbitrary tree of caches, where each node stores a conservative approximation of its children’s states.

As an instance of the above decomposition, we will prove that a multiprocessor system with no local buffering in between the processor and the memory components indeed implements SC. We implement a highly speculative processor that executes instructions and issues loads out of order, but commits instructions (once some “verification” is done) in order.

The processor itself can be decomposed into several components. In the zoomed-in version of Figure 1, we show a highly speculative, out-of-order issue processor. We have the normal architectural state, such as values of registers. Our proofs are generic over *a family of instruction set architectures*, with parameters for opcode sets and functions for executing opcodes and decoding them from memory. Other key components are a *branch predictor*, which guesses at the control-flow path that a processor will follow, to facilitate speculation; and a *reorder buffer (ROB)*, which decides which instructions along that path to try executing ahead of schedule. Our proofs apply to an arbitrary branch predictor and any reorder buffer satisfying a simple semantic condition.

Our framework establishes theorems of the form “if system A has a run with some particular observable behavior, then system B also has a run with the same behavior.” In this sense, we say that A correctly implements B . Other important properties, such as *deadlock freedom* for A (which might get stuck without producing any useful behavior), we leave for future work.

5 Specifying Sequential Consistency

Halt	$\frac{\theta(i) = (s, pc) \quad \text{dec}(s, pc) = H}{(\theta, m) \xrightarrow{\text{SC}} (H)}$	NonMem	$\frac{\theta(i) = (s, pc) \quad \text{dec}(s, pc) = (\text{Nm}, x) \quad \text{exec}(s, pc, (\text{Nm}, x)) = (s', pc')}{(\theta, m) \xrightarrow{\text{SC}} (\theta[i := (s', pc')], m)}$
Load	$\frac{\theta(i) = (s, pc) \quad \text{dec}(s, pc) = (\text{Ld}, x, a)}{(\theta, m) \xrightarrow{\text{SC}} (\theta[i := (s', pc')], m)}$		$\frac{\text{exec}(s, pc, (\text{Ld}, x, m(a))) = (s', pc')}{(\theta, m) \xrightarrow{\text{SC}} (\theta[i := (s', pc')], m)}$
Store	$\frac{\theta(i) = (s, pc) \quad \text{dec}(s, pc) = (\text{St}, a, v)}{(\theta, m) \xrightarrow{\text{SC}} (\theta[i := (s', pc')], m[a := v])}$		$\frac{\text{exec}(s, pc, (\text{St})) = (s', pc')}{(\theta, m) \xrightarrow{\text{SC}} (\theta[i := (s', pc')], m[a := v])}$

Fig. 3: LTS for SC with n simple processors

Our final theorem in this paper establishes that a particular complex hardware system implements sequential consistency (SC) properly. We state the theorem in terms of the trace refinement relation \sqsubseteq developed in Section 3. Therefore, we need to define an LTS embodying SC. The simpler this system, the better. We do not need to worry about its performance properties, since we will prove that an optimized system remains faithful to it.

Figure 3 defines an n -processor, sequentially consistent system as an LTS, parameterized over details of the ISA. In particular, the ISA gives us some domains of architectural states s (e.g., register files) and of program counters pc . A function $\text{dec}(s, pc)$ figures out which instruction pc references in the current state, returning the instruction’s “decoded” form. A companion function $\text{exec}(s, pc, d)$ actually executes the instruction, returning a new state s' and the next program counter pc' .

The legal instruction forms, which are outputs of dec , are (Nm, x) , for an operation not accessing memory; (Ld, x, a) , for a memory load from address a ; (St, a, v) , for a memory store of value v to address a ; and H , for a “halt” instruction that moves the LTS to state H . The parameter x above represents the rest of the instruction, including the opcode, registers, constants, *etc.*

The legal inputs to exec encode both a decoded instruction and any relevant responses from the memory system. These inputs are (Nm, x) and St , which need no extra input from the memory; and (Ld, x, v) , where v gives the contents of the requested memory cell.

We define the initial state of SC as (θ_0, m_0) , where m_0 is some initial memory fixed throughout our development, mapping every address to value v_0 ; and θ_0 maps every processor ID to (s_0, pc_0) , using architecture-specific default values s_0 and pc_0 .

This LTS encodes Lamport’s notion of SC, where processors take turns executing nondeterministically in a simple interleaving. Note that, in this setting, an operational specification like the LTS for SC is precisely the right characterization of *full functional correctness* for a hardware design, much as a precondition-postcondition pair does that job in a partial-correctness Hoare logic. Our SC LTS fully constrains observable behavior of a system to remain consistent with simple interleaving. Similar operational models are possible as top-level specifications for systems following weaker memory models, by giving the LTS for the *local buffer* component and composing the three components simultaneously.

Our final, optimized system is parameterized over an ISA in the same way as SC is. In the course of the rest of this paper, we will define an optimized system O and prove $O \sqsubseteq \text{SC}$. To support a modular proof decomposition, however, we need to introduce a few intermediate systems first.

$\text{Halt} \frac{\text{dec}(s, pc) = H}{(s, pc, \perp) \xrightarrow{\text{P}_{\text{ref}}} (H)}$	$\text{NM} \frac{\text{dec}(s, pc) = (\text{Nm}, x) \quad \text{exec}(s, pc, (\text{Nm}, x)) = (s', pc')}{(s, pc, \perp) \xrightarrow{\text{P}_{\text{ref}}} (s', pc', \perp)}$
$\text{LdRq} \frac{\text{dec}(s, pc) = (\text{Ld}, x, a)}{(s, pc, \perp) \xrightarrow{\text{P}_{\text{ref}}} (s, pc, \top)}$	$\text{StRq} \frac{\text{dec}(s, pc) = (\text{St}, a, v)}{(s, pc, \perp) \xrightarrow{\text{P}_{\text{ref}}} (s, pc, \top)}$
$\text{LdRp} \frac{\text{dec}(s, pc) = (\text{Ld}, x, a) \quad \text{exec}(s, pc, (\text{Ld}, x, v)) = (s', pc')}{(s, pc, \top) \xrightarrow{\text{P}_{\text{ref}}} (s', pc', \perp)}$	$\text{StRp} \frac{\text{dec}(s, pc) = (\text{St}, a, v) \quad \text{exec}(s, pc, (\text{St})) = (s', pc')}{(s, pc, \top) \xrightarrow{\text{P}_{\text{ref}}} (s', pc', \perp)}$

Fig. 4: LTS for a simple decoupled processor (P_{ref})

6 Respecifying Sequential Consistency with Communication

Realistic hardware systems do not implement the monolithic SC model of Figure 3 directly. Instead, there is usually a split between processors and memory. Here we formalize that split using LTSes that compose to produce a system refining the simple SC model.

Figure 4 defines an LTS for a simple *decoupled* processor (P_{ref}). That is, memory does not appear within a processor’s state, but instead, to load from or store to an address, we send *requests* to the memory system and receive its *responses*. Both kinds of

messages are encoded as labels, ToM for requests to memory and ToP for responses from memory back to the processor.

A state of P_{ref} is a triple $(s, pc, wait)$, giving the current architectural state s and program counter pc , as well as a Boolean flag $wait$ indicating whether the processor is blocked waiting for a response from the memory system. As in the SC model, we change the state of the processor to H whenever dec returns H .

As initial state for system P_{ref} , we use (s_0, pc_0, \perp) .

The simple memory defined earlier in Figure 2 is meant to be composed with P_{ref} processors. A request to memory like (t, Ld, a) asks for the value of memory cell a , associating a *tag* t that the processor can use to match responses to requests. Those responses take the form (t, Ld, v) , giving the value v of the requested memory address.

A memory state is a triple $(ins, outs, m)$, giving not just the memory m itself, but also buffers *ins* and *outs* for receiving processor requests and batching up responses to processors, respectively. We define the initial state of the M_m LTS as $(\emptyset, \emptyset, m_0)$, with empty queues.

Now we can compose these LTSes to produce an implementation of SC.

For a system of n processors, our decoupled SC implementation is $P_{\text{ref}}^n + M_m$.

Theorem 4. $P_{\text{ref}}^n + M_m \sqsubseteq SC$

Proof. By induction on traces of the decoupled system. We need to choose an abstraction function f from states of the complex system to states of the simple system. This function must be inductive in the appropriate sense: a step from s to s' on the left of the simulation relation must be matched by sequences of steps on the right from $f(s)$ to $f(s')$. We choose f that just preserves state components in states with no pending memory-to-processor responses. When such responses exist, f first executes them on the appropriate processors. \square

7 Speculative Out-Of-Order Processor

We implement a *speculative* processor, which may create many simultaneous outstanding requests to the memory, as an optimization to increase parallelism. Our processor proof is in some sense generic over correct speculation strategies. We parameterize over two key components of a processor design: a *branch predictor*, which makes guesses about processor-local control flow in advance of resolving conditional jumps; and a *re-order buffer*, which decides what speculative instructions (like memory loads) are worth issuing at which moments (in effect *reordering* later instructions to happen before earlier instructions have finished).

The branch predictor is the simpler of the two components, whose state is indicated with metavariable bp . The operations on such state are $\text{curPpc}(bp)$, to extract the current program-counter prediction; $\text{nextPpc}(bp)$, to advance to predicting the next instruction; and $\text{setNextPpc}(bp, pc)$, to reset prediction to begin at a known-accurate position pc . We need to impose no explicit correctness criterion on branch predictors; the processor uses predictions only as hints, and it always resets the predictor using setNextPpc after detecting an inaccurate hint.

Fetch	$\frac{}{(s, pc, wait, rob, bp) \xrightarrow{P_{so}} (s, pc, wait, insert(curPpc(bp), rob), nextPpc(bp))}$
Comp	$\frac{compute(rob) = (rob', \epsilon)}{(s, pc, wait, rob, bp) \xrightarrow{P_{so}} (s, pc, wait, rob', bp)} \quad SpLdRq \frac{compute(rob) = (rob', (SpecLd, t, a))}{(s, pc, wait, rob, bp) \xrightarrow{P_{so}} (s, pc, wait, rob', bp)}$
SpLdRp	$\frac{t \neq \epsilon}{(s, pc, wait, rob, bp) \xrightarrow{P_{so}} (s, pc, wait, updLd(rob, t, v), bp)} \quad \frac{ToP(t, Ld, v)}{P_{so}}$
Abort	$\frac{commit(rob) = (pc', -, -) \quad pc' \neq pc}{(s, pc, wait, rob, bp) \xrightarrow{P_{so}} (s, pc, wait, \phi, setNextPpc(bp, pc))}$
Halt	$\frac{commit(rob) = H}{(s, pc, \perp, rob, bp) \xrightarrow{P_{so}} (H)} \quad Nm \frac{commit(rob) = (pc, pc', (Nm, s'))}{(s, pc, \perp, rob, bp) \xrightarrow{P_{so}} (s', pc', \perp, retire(rob), bp)}$
StRq	$\frac{commit(rob) = (pc, pc', (St, a, v, s'))}{(s, pc, \perp, rob, bp) \xrightarrow{P_{so}} (s, pc, \top, rob, bp)} \quad LdRq \frac{commit(rob) = (pc, pc', (Ld, x, a, v, s'))}{(s, pc, \perp, rob, bp) \xrightarrow{P_{so}} (s, pc, \top, rob, bp)}$
StRp	$\frac{commit(rob) = (pc, pc', (St, a, v, s'))}{(s, pc, \top, rob, bp) \xrightarrow{P_{so}} (s', pc', \perp, retire(rob), bp)} \quad LdRpGd \frac{commit(rob) = (pc, pc', (Ld, x, a, v, s'))}{(s, pc, \top, rob, bp) \xrightarrow{P_{so}} (s', pc', \perp, retire(rob), bp)}$
LdRpBad	$\frac{commit(rob) = (pc, pc', (Ld, x, a, v', s')) \quad v' \neq v \quad exec(s, pc, (Ld, x, v)) = (s'', pc'')}{(s, pc, \top, rob, bp) \xrightarrow{P_{so}} (s'', pc'', \perp, \phi, setNextPpc(bp, pc''))}$

Fig. 5: Speculating, out-of-order issue processor

The interface and formal contract of a reorder buffer are more involved. We write rob as a metavariable for reorder-buffer state. ϕ denotes the state of an empty buffer. The operations associated with rob are:

- $insert(pc, rob)$, which appends the program instruction at location pc to the list of instructions that the buffer is allowed to consider executing.
- $compute(rob)$, which models a step of computation inside the buffer, returning both an updated state and an optional speculative load to issue. For instance, it invokes the dec and $exec$ functions (as defined for SC) internally to obtain the next program counter, state, *etc.* (but the actual states are not updated).
- $updLd(rob, t, v)$, which informs the buffer that the memory has returned result value v for the speculative load with tag $t \neq \epsilon$.
- $commit(rob)$, which returns the next instruction in serial program order, if we have accumulated enough memory responses to execute it accurately, or returns ϵ otherwise. When $commit$ returns an instruction, it also returns the associated program counter plus the next program counter we would advance to afterward. Furthermore, the instruction is extended with any relevant response from memory (used only for load instructions, obtained through $updLd$) and with the new architectural state (*e.g.*, register file) after execution.
- $retire(rob)$, which informs the buffer that its $commit$ instruction was executed successfully, so it is time to move on to the next instruction.

Figure 5 defines the speculative processor LTS P_{so} . This processor may issue arbitrary speculative loads, but it only ever *commits* the next instruction in serial program order. We will see the processor issue two kinds of loads, a speculative load (whose tag

is not ϵ) and a commit or real load (whose tag is ϵ). To maintain SC, every speculative load must have a matching verification load later on, and we maintain the illusion that the program only depends on the results of verification loads, which, along with stores, *must be issued in serial program order*.

When committing a previously issued speculative load instruction, the associated speculative memory load response is verified against the new commit load response. If the resulting values do not match, the processor terminates all past uncommitted speculation, by emptying the reorder buffer and resetting the next predicted program counter in the branch predictor to the correct next value. In common cases, performance of executing loads twice is good, because it is likely that the verification load finds the address already in a local cache, thanks to the recent processing of the speculative load. Moreover, 60% to 90% of verification loads can be avoided by tracking speculative loads [13]; in the future we plan to extend our proofs to include optimizations along these lines.

A full processor state is $(s, pc, wait, rob, bp)$, comprising architectural state, the program counter, a Boolean flag indicating whether the processor is waiting for a memory response about an instruction being committed, and the reorder-buffer and branch-predictor states. Its initial state is given by $(s_0, pc_0, \perp, \phi, bp_0)$. The interface of this processor with memory (*i.e.*, communication labels with ToM, ToP) is identical to that of the reference processor.

Finally, we impose a general correctness condition on reorder buffers (Figure 6). Intuitively, whenever the buffer claims in a commit output that a particular instruction is next to execute, causing certain state changes, that instruction must really be next in line according to how the program runs in the SC system, and its execution must really cause those state changes.

<p>ROB-invariant: If P_{so} reaches a state $(s, pc, wait, rob, bp)$,</p> $\begin{cases} \text{commit}(rob) = (pc, pc', (Nm, s')) & \Rightarrow \exists x. \text{dec}(s, pc) = (Nm, x) \wedge \text{exec}(s, pc, (Nm, x)) = (s', pc') \\ \text{commit}(rob) = (pc, pc', (Ld, x, a, v, s')) & \Rightarrow \text{dec}(s, pc) = (Ld, x, a) \wedge \text{exec}(s, pc, (Ld, x, v)) = (s', pc') \\ \text{commit}(rob) = (pc, pc', (St, a, v, s')) & \Rightarrow \text{dec}(s, pc) = (St, a, v) \wedge \text{exec}(s, pc, (St)) = (s', pc') \\ \text{commit}(rob) = H & \Rightarrow \text{dec}(s, pc) = H \end{cases}$
--

Fig. 6: Correctness of reorder buffer

When this condition holds, we may conclude the correctness theorem for out-of-order processors. We use a trace-transformation function `noSpec` that drops all speculative-load requests and responses (*i.e.*, those load requests and responses whose tags are not ϵ). See Definition 5 for a review of how we use such functions in framing trace refinement. Intuitively, we prove that any behavior by the speculating processor can be matched by the simple processor, with speculative messages erased.

Theorem 5. $P_{so} \sqsubseteq_{\text{noSpec}} P_{ref}$

Proof. By induction on P_{so} traces, using an abstraction function that drops the speculative messages and the *rob* and *bp* states to relate the two systems. The reorder-buffer correctness condition is crucial to relate its behavior with the simple in-order execution of P_{ref} . \square

Corollary 1. $P_{so}^n \sqsubseteq_{\text{noSpec}^n} P_{ref}^n$

Proof. Direct consequence of Theorems 5 and 2 (the latter is about n -repetition). \square

8 Cache-Based Memory System

Processor/Memory Interface	
Ins $\frac{\left(\begin{array}{l} d, ch, cs, \\ dir, w, dirw, \\ ins, outs \end{array} \right)}{M_c} \xrightarrow{i, \text{ToM}(q)} \left(\begin{array}{l} d, ch, cs, dir, w, \\ dirw, ins[i := \{q\}] \\ \cup ins(i), outs \end{array} \right)$	Rm $\frac{\left(\begin{array}{l} d, ch, cs, \\ dir, w, dirw, \\ ins, outs \end{array} \right)}{M_c} \xrightarrow{i, \text{ToP}(r)} \left(\begin{array}{l} d, ch, cs, dir, \\ w, dirw, ins, \\ outs[i := rs] \end{array} \right)$
$outs(i) = rs \uplus \{r\}$	
Ld $\frac{\left(\begin{array}{l} d, ch, cs, dir, w, \\ dirw, ins, outs \end{array} \right)}{M_c} \xrightarrow{ins(c) = \{(t, \text{Ld}, a)\} \uplus rs \quad cs(c, a) \geq S} \left(\begin{array}{l} d, ch, cs, dir, w, dirw, ins[c := rs], \\ outs[c := outs(c) \uplus \{(t, \text{Ld}, d(c, a))\}] \end{array} \right)$	
St $\frac{\left(\begin{array}{l} d, ch, cs, dir, w, \\ dirw, ins, outs \end{array} \right)}{M_c} \xrightarrow{ins(c) = \{(St, a, v)\} \uplus rs \quad cs(c, a) \geq M} \left(\begin{array}{l} d[(c, a) := v], ch, cs, dir, w, dirw, ins[c := rs], \\ outs[c := outs(c) \uplus \{(St)\}] \end{array} \right)$	
Child Upgrade	
ChildSendReq $\frac{\left(\begin{array}{l} d, ch, cs, dir, w, \\ dirw, ins, outs \end{array} \right)}{M_c} \xrightarrow{parent(c, p) \quad cs(c, a) < x \quad w(c, a) = \epsilon} \left(\begin{array}{l} d, ch[(c, p, \text{Rq}) := (a, cs(c, a), x) \uplus ch(c, p, \text{Rq})], \\ cs, dir, w[(c, a) := x], dirw, ins, outs \end{array} \right)$	
ParentRecvReq $\frac{\left(\begin{array}{l} d, ch, cs, \\ dir, w, dirw, \\ ins, outs \end{array} \right)}{M_c} \xrightarrow{parent(c, p) \quad ch(c, p, \text{Rq}) = \{(a, y, x)\} \uplus rs \quad cs(p, a) \geq x \\ dirCompat(p, c, x, a) \quad dirw(p, c, a) = \epsilon \quad dir(p, c, a) \leq y} \left(\begin{array}{l} d, ch[(c, p, \text{Rq}) := rs][(p, c, \text{RR}) := (\text{Rp}, (a, y, x, v), x, \\ \text{if}(dir(p, c, a) = I) \text{ then } d(p, a) \text{ else } _) :: ch(p, c, \text{RR})], \\ cs, dir[(p, c, a) := x], w, dirw, ins, outs \end{array} \right)$	
ChildRecvRsp $\frac{\left(\begin{array}{l} d, ch, cs, \\ dir, w, dirw, \\ ins, outs \end{array} \right)}{M_c} \xrightarrow{parent(c, p) \quad ch(p, c, \text{RR}) = rs :: (\text{Rp}, (a, y, x, v))} \left(\begin{array}{l} d[(c, a) := \text{if}(y = I) \text{ then } v \text{ else } d(c, a)], ch[(p, c, \text{RR}) := rs], \\ cs[(c, a) := x], dir, w[(c, a) := \text{if}(w(c, a) \leq x) \text{ then } \epsilon \\ \text{else } w(c, a)], dirw, ins, outs \end{array} \right)$	
Parent Downgrade	
ParentSendReq $\frac{\left(\begin{array}{l} d, ch, cs, dir, w, \\ dirw, ins, outs \end{array} \right)}{M_c} \xrightarrow{parent(c, p) \quad dir(p, c, a) > x \quad dirw(p, c, a) = \epsilon} \left(\begin{array}{l} d, ch[(p, c, \text{RR}) := (\text{Rq}, (a, dir(p, c, a), x)) :: ch(p, c, \text{RR})], \\ cs, dir, w, dirw[(p, c, a) := x], ins, outs \end{array} \right)$	
ChildRecvReq $\frac{\left(\begin{array}{l} d, ch, cs, \\ dir, w, dirw, \\ ins, outs \end{array} \right)}{M_c} \xrightarrow{parent(c, p) \quad ch(p, c, \text{RR}) = rs :: (\text{Rq}, (a, y, x)) \\ (\forall i. parent(i, c) \Rightarrow dir(c, i, a) \leq x) \quad cs(c, a) > x} \left(\begin{array}{l} d, ch[(p, c, \text{RR}) := rs][(c, p, \text{Rp}) := (a, cs(c, a), x, \\ \text{if}(dir(c, a) = M) \text{ then } d(c, a) \text{ else } _) :: ch(c, p, \text{Rp})], \\ cs[(c, a) := x], dir, w, dirw, ins, outs \end{array} \right)$	
ParentRecvRsp $\frac{\left(\begin{array}{l} d, ch, cs, \\ dir, w, dirw, \\ ins, outs \end{array} \right)}{M_c} \xrightarrow{parent(c, p) \quad ch(c, p, \text{Rp}) = \{(a, y, x, v)\} :: rs \quad dir(p, c, a) = y} \left(\begin{array}{l} d[(p, a) := \text{if}(y = M) \text{ then } v \text{ else } d(p, a)], ch[(c, p, \text{Rp}) := \\ rs], cs, dir[(p, c, a) := x], w, dirw[(p, c, a) := \\ \text{if}(dirw(p, c, a) \geq x) \text{ then } \epsilon \text{ else } dirw(p, c, a)], ins, outs \end{array} \right)$	
Voluntary downgrade for replacement	
VolResp $\frac{\left(\begin{array}{l} d, ch, cs, \\ dir, w, dirw, \\ ins, outs \end{array} \right)}{M_c} \xrightarrow{parent(c, p) \quad (\forall i. parent(i, c) \Rightarrow dir(c, i, a) \leq x) \quad cs(c, a) > x} \left(\begin{array}{l} d, ch[(c, p, \text{Rp}) := (a, cs(c, a), x, \\ \text{if}(cs(c, a) = M) \text{ then } d(c, a) \text{ else } _) :: ch(c, p, \text{Rp})], \\ cs[(c, a) := x], dir, w, dirw, ins, outs \end{array} \right)$	
Dropping request because of voluntary downgrade	
DropReq $\frac{\left(\begin{array}{l} d, ch, cs, dir, w, dirw, ins, outs \end{array} \right)}{M_c} \xrightarrow{parent(c, p) \quad ch(p, c, \text{RR}) = rs :: (\text{Rq}, (a, y, x)) \quad cs(c, a) \leq x} \left(\begin{array}{l} d, ch[(p, c, \text{RR}) := rs], cs, dir, w, dirw, ins, outs \end{array} \right)$	

Fig. 7: LTS for cache-coherent shared-memory system

We now turn our attention toward a more efficient implementation of memory. With the cache hierarchy of Figure 1, we have concurrent interaction of many processors with many caches, and the relationship with the original M_m system is far from direct. However, this intricate concurrent execution is crucial to hiding the latency of main-memory

access. Figure 7 formalizes as an LTS M_c the algorithm we implemented, based on a published implementation [17], for providing the memory abstraction on top of a cache hierarchy. We have what is called an *invalidating directory-based hierarchical cache-coherence protocol*.

We describe a state of the system using fields d , ch , cs , dir , w , $dirw$, ins , $outs$. The ins and $outs$ sets are the interfaces to the processors and are exactly the same as in M_m (Figure 2). We use $parent(c, p)$ to denote that p is the parent of c .

A coherence state is M , S or I , broadly representing permissions to modify, read, or do nothing with an address, respectively, the decreasing permissions denoted by $M > S > I$. More precisely, if a node is in coherence state M or S for some address, then there might be some node in the former's sub-tree which has write or read permissions, respectively, for that address. Coherence state of cache c for address a is denoted by $cs(c, a)$. $d(c, a)$ represents the data in cache c for address a .

$w(c, a)$ stores what permission an address a in cache c is waiting for, if any. That is, cache c has decided to *upgrade* its coherence state for address a to a more permissive value, but is waiting for acknowledgment from its parent before upgrading.

$dir(p, c, a)$ represents the parent p 's notion of the coherence state of the child c for address a . We later prove that this notion is always conservative, *i.e.*, if the parent assumes that a child does not have a particular permission, then it is guaranteed in this system that the child will not have that permission. $dirw(p, c, a)$ denotes whether the parent p is waiting for any downgrade response from its child c for address a , and if so, the coherence state that the child must downgrade to as well.

There are three types of communication channels in the system: (i) $ch(p, c, RR)$ which carries both downgrade request and upgrade response messages from parent p to its child c , (ii) $ch(c, p, Rq)$ which carries upgrade request messages from child c to its parent p and (iii) $ch(c, p, Rp)$ which carries downgrade response messages from child c to its parent p . While the $ch(c, p, Rp)$ and $ch(p, c, RR)$ channels deliver messages between the same pair of nodes in the same order in which the messages were injected (*i.e.*, they obey the FIFO property, indicated by the use of $::$ in Figure 7), $ch(c, p, Rq)$ needn't obey such a property (indicated by the use of \uplus for unordered bags in Figure 7). This asymmetry is because only one downgrade request can be outstanding for one parent-child pair for an address.

Here is an intuition on how the transitions work in the common case. A cache can decide, spontaneously, to upgrade its coherence state, in which case it sends an upgrade request to its parent. The parent then makes a local decision on whether to send a response to the requesting child or not, based on its directory approximation and its own coherence state cs . If cs is lower than the requested upgrade, then it cannot handle the request, and instead must decide to upgrade cs . Once the parent's cs is not lower than the requested upgrade, it makes sure that the rest of its children are "compatible" with the requested upgrade (given by the $dirCompat$ definition below). If not, the parent must send requests to the incompatible children to downgrade. Finally, when the cs 's upgrade and children's downgrade responses are all received, the original request can be responded to. A request in ins can be processed by an L1 cache only if it is in the appropriate state, otherwise it has to request an upgrade for that address.

Definition 6 $dirCompat(p, c, x, a) = \begin{cases} x = M \Rightarrow \forall c' \neq c. dir(p, c', a) = I \\ x = S \Rightarrow \forall c' \neq c. dir(p, c', a) \leq S \end{cases}$

A complication arises because a cache can voluntarily decide to downgrade its state. This transition is used to model invalidation of cache lines to make room for a different location. As a result, the parent's *dir* and the corresponding *cs* of the child may go out of sync, leading to the parent requesting a child to downgrade when it already has. To handle this situation, the child has to drop the downgrade request when it has already downgraded to the required state (Rule DropReq in Figure 7), to avoid deadlocks by not dequeuing the request.

8.1 Proving M_c is Store Atomic

We must prove the following theorem, *i.e.*, the cache-based system is sound with respect to the simple memory.

Theorem 6. $M_c \sqsubseteq M_m$

We present the key theorem needed for this proof below. Throughout this section, we say *time* to denote the number of transitions that occurred before reaching the specified state.

Theorem 7. *A is store atomic, i.e., $A \sqsubseteq M_m$ and $M_m \sqsubseteq A$ iff for any load request $\text{ToM}(t, \text{Ld}, a)$ received, the response $\text{ToP}(t, \text{Ld}, v)$ sent at time T is such that*

1. $v = v_0$ (the initial value of any memory address) and no store request $\text{ToM}(\text{St}, a, v')$ has been processed at any time T' such that $T' < T$ or
2. There is a store request $\text{ToM}(\text{St}, a, v)$ that was processed at time T_q such that $T_q < T$ and no other store request $\text{ToM}(\text{St}, a, v')$ was processed at any time T' such that $T_q < T' < T$.

The proof that M_c obeys the properties in Theorem 7 is involved enough that we will only state key lemmas that we used.

Lemma 1. *At any time T , if address a in cache c obeys $cs(c, a) \geq S$ and $\forall i. \text{dir}(c, i, a) \leq S$, then a will have the **latest value**, *i.e.*,*

1. $d(c, a) = v_0$ and no store request $\text{ToM}(\text{St}, a, v)$ has been processed at any time T' such that $T' < T$ or
2. There is a store request $\text{ToM}(\text{St}, a, v)$ that was processed at time T_q such that $T_q < T \wedge d(c, a) = v$ and no other store request $\text{ToM}(\text{St}, a, v')$ was processed at any time T' such that $T_q < T' < T$.

It is pretty straightforward to prove the properties of Theorem 7 given Lemma 1. To prove Lemma 1, it has to be decomposed further into the following, each of which holds at any time.

Lemma 2. *If data for an address a are in transit (*i.e.*, $\forall T. T_s \leq T \leq T_r$ where T_s is the time of sending the data and T_r the time of receiving the data), no cache can process a store request for a , and the data must be sent from a clean cache.*

Lemma 3. *At any time, $\forall p, \forall c, \forall a. \text{parent}(c, p) \Rightarrow$
 $cs(c, a) \leq \text{dir}(p, c, a) \wedge \text{dirCompat}(p, c, \text{dir}(p, c, a), a) \wedge \text{dir}(p, c, a) \leq cs(p, a)$*

The same proof structure can be used to prove other invalidation-based protocols with inclusive caches (where any address present in a cache will also be present in its parent) like MESI, MOSI, and MOESI; we omit the discussion of extending this proof to these for space reasons. The MSI proof is about 12,000 lines of Coq code, of which 80% can be reused as-is for the other protocols.

9 The Final Result

With our two main results about optimized processors and memories, we can complete the correctness proof of the composed optimized system.

First, we need to know that, whenever the simple memory can generate some trace of messages, it could also generate the same trace with all speculative messages removed. We need this property to justify the introduction of speculation, during our final series of refinements from the optimized system to SC.

Theorem 8. $M_m \sqsubseteq_{\text{noSpec}^n} M_m$

Proof. By induction on traces, with an identity abstraction function. □

That theorem turns out to be the crucial ingredient to justify placing a speculative processor in-context with simple memory.

Theorem 9. $P_{so}^n + M_m \sqsubseteq P_{ref}^n + M_m$

Proof. Follows from Theorem 3 (our result about +), Corollary 1, and Theorem 8. □

The last theorem kept the memory the same while refining the processor. The next one does the opposite, switching out memory.

Theorem 10. $P_{so}^n + M_c \sqsubseteq P_{so}^n + M_m$

Proof. Follows from Theorems 6 and 3 plus reflexivity of \sqsubseteq (Theorem 1). □

Theorem 11. $P_{so}^n + M_c \sqsubseteq SC$

Proof. We twice apply \sqsubseteq transitivity (Theorem 1) to connect Theorems 10, 9, and 4 □

10 Conclusions and Future Work

In this paper we developed a modular proof structure for distributed shared-memory hardware systems, corresponding naturally to the modularization seen in hardware implementations. Our approach allows verification of the system to be carried out throughout implementation without disrupting the architect's design process.

While we provide a clean interface for an SC system, we are working on encompassing relaxed memory models commonly used in modern processors.

Finally, to make full use of our results for implementations, hardware descriptions must be translated to and from our labeled transition system representation. As we implied in the discussion of our cache-coherence system, Bluespec's rule-based representation is very close to LTSes and in some cases can be transliterated directly. We are working on automating this along the lines of Braibant *et al.* [11].

References

1. Jade Alglave. A formal hierarchy of weak memory models. *Form. Methods Syst. Des.*, 41(2):178–210, October 2012.
2. Jade Alglave and Luc Maranget. Stability in weak memory models. In *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, pages 50–66, 2011.
3. Jade Alglave, Luc Maranget, Susmit Sarkar, and Peter Sewell. Fences in weak memory models. In *Proceedings of the 22Nd International Conference on Computer Aided Verification, CAV'10*, pages 258–272, Berlin, Heidelberg, 2010. Springer-Verlag.
4. Jade Alglave, Luc Maranget, Susmit Sarkar, and Peter Sewell. Litmus: Running tests against hardware. In *Proceedings of the 17th International Conference on Tools and Algorithms for the Construction and Analysis of Systems: Part of the Joint European Conferences on Theory and Practice of Software, TACAS'11/ETAPS'11*, pages 41–44, Berlin, Heidelberg, 2011. Springer-Verlag.
5. Jade Alglave, Luc Maranget, and Michael Tautschnig. Herding cats: Modelling, simulation, testing, and data-mining for weak memory. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, page 7. ACM, 2014.
6. Arvind and Xiaowei Shen. Using term rewriting systems to design and verify processors. *Micro, IEEE*, 19(3):36–46, 1999.
7. Lennart Augustsson, Jacob Schwarz, and Rishiyur S. Nikhil. Bluespec Language definition, 2001. Sandburst Corp.
8. Ritwik Bhattacharya, Steven German, and Ganesh Gopalakrishnan. Symbolic partial order reduction for rule based transition systems. In Dominique Borrione and Wolfgang Paul, editors, *Correct Hardware Design and Verification Methods*, volume 3725 of *Lecture Notes in Computer Science*, pages 332–335. Springer Berlin Heidelberg, 2005.
9. Ritwik Bhattacharya, Steven M. German, and Ganesh Gopalakrishnan. Exploiting symmetry and transactions for partial order reduction of rule based specifications. In *In Antti Valmari, editor, SPIN, volume 3925 of Lecture Notes in Computer Science*, pages 252–270. Springer, 2006.
10. Bluespec, Inc., Waltham, MA. *Bluespec SystemVerilog Version 3.8 Reference Guide*, November 2004.
11. Thomas Braibant and Adam Chlipala. Formal verification of hardware synthesis. In *CAV 2013, 25th International Conference on Computer Aided Verification*, volume 8044 of *Lecture Notes in Computer Science*, pages 213–228. Springer, 2013.
12. Jerry R Burch and David L Dill. Automatic verification of pipelined microprocessor control. In *Computer Aided Verification*, pages 68–80. Springer, 1994.
13. H.W. Cain and M.H. Lipasti. Memory ordering: a value-based approach. In *Computer Architecture, 2004. Proceedings. 31st Annual International Symposium on*, pages 90–101, June 2004.
14. C. Chang, J. Wawrzynek, and R.W. Brodersen. Bee2: a high-end reconfigurable computing system. *Design Test of Computers, IEEE*, 22(2):114–125, March 2005.
15. Xiaofang Chen, Yu Yang, Ganesh Gopalakrishnan, and Ching-Tsun Chou. Efficient methods for formally verifying safety properties of hierarchical cache coherence protocols. *Form. Methods Syst. Des.*, 36(1):37–64, February 2010.
16. Ching-Tsun Chou, Phanindra K. Mannava, and Seungjoon Park. A simple method for parameterized verification of cache coherence protocols. In *Formal Methods in Computer Aided Design*, pages 382–398. Springer, 2004.
17. Nirav Dave, Man Cheuk Ng, and Arvind. Automatic Synthesis of Cache-Coherence Protocol Processors Using Bluespec. In *Proceedings of Formal Methods and Models for Codesign (MEMOCODE)*, Verona, Italy, 2005.

18. Giorgio Delzanno. Automatic verification of parameterized cache coherence protocols. In E.Allen Emerson and AravindaPrasad Sistla, editors, *Computer Aided Verification*, volume 1855 of *Lecture Notes in Computer Science*, pages 53–68. Springer Berlin Heidelberg, 2000.
19. D.L. Dill, A.J. Drexler, A.J. Hu, and C.H. Yang. Protocol verification as a hardware design aid. In *Computer Design: VLSI in Computers and Processors, 1992. ICCD '92. Proceedings, IEEE 1992 International Conference on*, pages 522–525, Oct 1992.
20. E. Allen Emerson and Vineet Kahlon. Exact and efficient verification of parameterized cache coherence protocols. In *Correct Hardware Design and Verification Methods, 12th IFIP WG 10.5 Advanced Research Working Conference, CHARME 2003, L'Aquila, Italy, October 21-24, 2003, Proceedings*, pages 247–262, 2003.
21. James C. Hoe and Arvind. Synthesis of Operation-Centric Hardware Descriptions. In *Proceedings of ICCAD'00*, pages 511–518, San Jose, CA, 2000.
22. James C. Hoe and Arvind. Operation-Centric Hardware Description and Synthesis. *IEEE TRANSACTIONS on Computer-Aided Design of Integrated Circuits and Systems*, 23(9), September 2004.
23. Chung-Wah Norris Ip, David L. Dill, and John C. Mitchell. State reduction methods for automatic formal verification, 1996.
24. Ranjit Jhala and Kenneth L. McMillan. Microarchitecture verification by compositional model checking. In *CAV*, pages 396–410. Springer-Verlag, 2001.
25. Rajeev Joshi, Leslie Lamport, John Matthews, Serdar Tasiran, Mark R. Tuttle, and Yuan Yu. Checking cache-coherence protocols with TLA⁺. *Formal Methods in System Design*, 22(2):125–131, 2003.
26. Roope Kaivola, Rajnish Ghughal, Naren Narasimhan, Amber Telfer, Jesse Whittmore, Sudhinda Pandav, Anna Slobodová, Christopher Taylor, Vladimir Frolov, Erik Reeber, et al. Replacing testing with formal verification in Intel[®] Core[™] i7 processor execution engine validation. In *Computer Aided Verification*, pages 414–429. Springer, 2009.
27. Asif Khan, Muralidaran Vijayaraghavan, Silas Boyd-Wickizer, and Arvind. Fast and cycle-accurate modeling of a multicore processor. In *2012 IEEE International Symposium on Performance Analysis of Systems & Software, New Brunswick, NJ, USA, April 1-3, 2012*, pages 178–187, 2012.
28. J. Kuskin, D. Ofelt, M. Heinrich, J. Heinlein, R. Simoni, K. Gharachorloo, J. Chapin, D. Nakahira, J. Baxter, M. Horowitz, A. Gupta, M. Rosenblum, and J. Hennessy. The Stanford FLASH multiprocessor. In *Computer Architecture, 1994., Proceedings the 21st Annual International Symposium on*, pages 302–313, Apr 1994.
29. Leslie Lamport. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2002.
30. Sela Mador-Haim, Luc Maranget, Susmit Sarkar, Kayvan Memarian, Jade Alglave, Scott Owens, Rajeev Alur, Milo MK Martin, Peter Sewell, and Derek Williams. An axiomatic memory model for POWER multiprocessors. In *Computer Aided Verification*, pages 495–512. Springer, 2012.
31. K. L. McMillan. Parameterized verification of the FLASH cache coherence protocol by compositional model checking. In *In CHARME 01: IFIP Working Conference on Correct Hardware Design and Verification Methods, Lecture Notes in Computer Science 2144*, pages 179–195. Springer, 2001.
32. Kenneth L McMillan. Verification of an implementation of Tomasulo's algorithm by compositional model checking. In *Computer Aided Verification*, pages 110–121. Springer, 1998.
33. K.L. McMillan. Verification of infinite state systems by compositional model checking. In Laurence Pierre and Thomas Kropf, editors, *Correct Hardware Design and Verification Methods*, volume 1703 of *Lecture Notes in Computer Science*, pages 219–237. Springer Berlin Heidelberg, 1999.

34. KL McMillan and James Schwalbe. Formal verification of the Gigamax cache consistency protocol. In *Proceedings of the International Symposium on Shared Memory Multiprocessing*, pages 111–134, 1992.
35. J. Strother Moore. An ACL2 proof of write invalidate cache coherence. In *Proc. CAV'98, volume 1427 of LNCS*, pages 29–38. Springer, 1998.
36. Seungjoon Park and David L. Dill. Verification of FLASH cache coherence protocol by aggregation of distributed transactions. In *Proceedings of the 8th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 288–296. ACM Press, 1996.
37. Susmit Sarkar, Kayvan Memarian, Scott Owens, Mark Batty, Peter Sewell, Luc Maranget, Jade Alglave, and Derek Williams. Synchronising C/C++ and Power. In *PLDI*, volume 47, pages 311–322. ACM, 2012.
38. Susmit Sarkar, Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams. Understanding POWER multiprocessors. In *PLDI*, volume 46, pages 175–186. ACM, 2011.
39. Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. x86-TSO: A rigorous and usable programmer's model for x86 multiprocessors. *Communications of the ACM*, 53(7):89–97, July 2010.
40. Xiaowei Shen, Arvind, and Larry Rudolph. Commit-reconcile & fences (CRF): a new memory model for architects and compiler writers. In *Proceedings of the 26th annual international symposium on Computer architecture*, pages 150–161. IEEE Computer Society, 1999.
41. M. Talupur and Mark R. Tuttle. Going with the flow: Parameterized verification using message flows. In *Formal Methods in Computer-Aided Design, 2008. FMCAD '08*, pages 1–8, Nov 2008.
42. Phillip J. Windley. Formal modeling and verification of microprocessors. *Computers, IEEE Transactions on*, 44(1):54–72, 1995.
43. Meng Zhang, Jesse D. Bingham, John Erickson, and Daniel J. Sorin. Pvcoherence: Designing flat coherence protocols for scalable verification. In *20th IEEE International Symposium on High Performance Computer Architecture, HPCA 2014, Orlando, FL, USA, February 15-19, 2014*, pages 392–403. IEEE Computer Society, 2014.
44. Meng Zhang, Alvin R. Lebeck, and Daniel J. Sorin. Fractal coherence: Scalably verifiable cache coherence. In *Proceedings of the 2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '13*, pages 471–482, Washington, DC, USA, 2010. IEEE Computer Society.