



January 2005

Verifying Safety of a Token Coherence Implementation by Parametric Compositional Refinement

Sebastian Burckhardt
University of Pennsylvania

Rajeev Alur
University of Pennsylvania, alur@cis.upenn.edu

Milo M.K. Martin
University of Pennsylvania, milom@cis.upenn.edu

Follow this and additional works at: http://repository.upenn.edu/cis_papers

Recommended Citation

Sebastian Burckhardt, Rajeev Alur, and Milo M.K. Martin, "Verifying Safety of a Token Coherence Implementation by Parametric Compositional Refinement", *Lecture Notes in Computer Science: Verification, Model Checking, and Abstract Interpretation* 3385, 130-145. January 2005. http://dx.doi.org/10.1007/978-3-540-30579-8_9

From the 6th International Conference, VMCAI 2005, Paris, France, January 17-19, 2005.

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/cis_papers/174
For more information, please contact libraryrepository@pobox.upenn.edu.

Verifying Safety of a Token Coherence Implementation by Parametric Compositional Refinement

Abstract

We combine compositional reasoning and reachability analysis to formally verify the safety of a recent cache coherence protocol. The protocol is a detailed implementation of *token coherence*, an approach that decouples correctness and performance. First, we present a formal and abstract specification that captures the safety substrate of token coherence, and highlights the symmetry in states of the cache controllers and contents of the messages they exchange. Then, we prove that this abstract specification is coherent, and check whether the implementation proposed by the protocol designers is a refinement of the abstract specification. Our refinement proof is parametric in the number of cache controllers, and is compositional as it reduces the refinement checks to individual controllers using a specialized form of assume-guarantee reasoning. The individual refinement obligations are discharged using refinement maps and reachability analysis. While the formal proof justifies the intuitive claim by the designers about the ease of verifiability of token coherence, we report on several bugs in the implementation, and accompanying modifications, that were missed by extensive prior simulations.

Comments

From the 6th International Conference, VMCAI 2005, Paris, France, January 17-19, 2005.

Verifying Safety of a Token Coherence Implementation by Parametric Compositional Refinement*

(Extended Version with Proofs)[†]

Sebastian Burckhardt Rajeev Alur Milo M. K. Martin

University of Pennsylvania

Abstract

We combine compositional reasoning and reachability analysis to formally verify the safety of a recent cache coherence protocol. The protocol is a detailed implementation of *token coherence*, an approach that decouples correctness and performance. First, we present a formal and abstract specification that captures the safety substrate of token coherence, and highlights the symmetry in states of the cache controllers and contents of the messages they exchange. Then, we prove that this abstract specification is coherent, and check whether the implementation proposed by the protocol designers is a refinement of the abstract specification. Our refinement proof is parametric in the number of cache controllers, and is compositional as it reduces the refinement checks to individual controllers using a specialized form of assume-guarantee reasoning. The individual refinement obligations are discharged using refinement maps and reachability analysis. While the formal proof justifies the intuitive claim by the designers about the ease of verifiability of token coherence, we report on several bugs in the implementation, and accompanying modifications, that were missed by extensive prior simulations.

*This research was partially supported by the NSF award CCR0306382, and a donation from Intel Corporation.

[†]The published short version [7] is copyrighted by Springer Verlag. This extended version includes more material, in particular more proofs. Because the readability is impacted, *we recommend that this version is used for reference only*. Both versions can be found online [8]. The numbering of proofs, definitions and examples is consistent with the regular versions — references that do only appear in the extended version are prefixed with E.

1 Introduction

Shared memory multiprocessors have become the most important architecture used for commercial and scientific workloads. Such systems use hardware cache coherence protocols to create the illusion of a single, shared memory without caches. These protocols are important factors of the overall system performance, and numerous optimizations contribute to their complexity. Since hard-to-cover race conditions elude simulations of the protocols, formal methods are often employed to verify their correctness.

Token Coherence is a new approach to cache coherence protocols that decouples correctness requirements from performance choices, claiming to improve both performance and verifiability [23]. Separate correctness mechanisms ensure safety and liveness. *Safety* is achieved by token counting: per memory location, the number of tokens in the system is a global invariant. By requiring at least one token for read access and all tokens for write access, the protocol directly enforces a single-writer, multiple-reader policy. On the other hand, *Liveness* is achieved by persistent requests. This reliable, but slower protocol is used when the regular requests do not succeed within a timeout period. Persistent requests are required because the regular requests, while likely to complete quickly, do not guarantee eventual success.

In this work, we combine compositional verification and model checking to verify the safety of a detailed implementation of a token coherence protocol for an arbitrary number of caches. Our method takes advantage of the opportunities offered by the token coherence design. It proceeds in four steps.

1. We present a formal specification of the safety substrate of token coherence. This abstract protocol is based on rewrite rules and multisets, and expresses the symmetry between components and messages. It applies to arbitrary network topologies, cache numbers, and even cache hierarchies.
2. We prove manually that the abstract protocol is safe (i.e. coherent). The verification problem is thus reduced to checking that the implementation correctly refines the abstract protocol.
3. We prove that the refinement can be verified for each component individually, by replacing its context with an abstraction. We prove that this decomposition into local refinement obligations is sound, using a variant of assume-guarantee reasoning based on *contextual refinement*, and performing an induction on the number of caches.
4. We discharge the local refinement obligations with the conventional model checker Mur φ [13, 12]. To obtain the models, we manually translate, abstract and annotate the implementation code. This procedure reduces the refinement checking to a reachability problem, which Mur φ solves by enumerative state space search.

Even though the protocol implementation had been extensively simulated prior to this work, we discovered a few bugs, and were able to fix them quickly with the help of counterexamples produced by the model checker. The compositional refinement method proved to be effective in avoiding the state space explosion problem [17] which is commonly encountered in system-level models [29].

1.1 Related Work

Prior work on formal verification of cache coherence varies in (1) the protocol complexity and level of detail (2) the coverage achieved (safety, liveness, parametric systems) (3) the underlying tools (enumerative or symbolic model checkers, decision procedures, theorem provers), (4) reduction techniques (symmetry, abstraction, compositional verification), and (5) degree of automation. We refer to Pong and Dubois [29] for a general survey, and to various illustrative efforts [24, 28, 15, 3].

Our proof methodology modifies and combines a variety of ideas in the formal verification literature. These include assume-guarantee reasoning for compositional verification (c.f. [1, 9, 2, 26]), structural induction for proving properties for arbitrary number of processes (c.f. [20, 10, 16, 14, 11, 4]), data abstraction (c.f. [34, 18]), use of term rewrite systems for hardware verification [5], and proving refinement using reachability analysis (c.f. [19]).

2 Process Model

In this section, we define the process model and introduce our assume-guarantee proof rules. We chose to define the process model from scratch, so to keep it concise and self-contained, and to obtain the desired combination of features. Except for the specialized definition of *contextual refinement*, all concepts (traces, composition, refinement) are standard and appear in many variations and combinations in the process algebra literature [30].

A process is defined as the set of its traces, which are finite words over an alphabet Σ of events. Σ is considered fixed and common to all processes. We further partition $\Sigma = \Sigma_e \cup \Sigma_c$ into disjoint subclasses: Σ_e contains events that are visible to external observers of the system only, while Σ_c describes synchronous communication events. Matching events in Σ_c (e.g. sending and receiving of a message) are denoted σ and $\bar{\sigma}$.

Definition 2.1 *A process P over Σ is a non-empty prefix-closed language; i.e. $P \subset \Sigma^*$, $P \neq \emptyset$ and for all $u, v \in \Sigma^* : uv \in P \Rightarrow u \in P$. A process P refines a process Q , written $P \preceq Q$, iff $P \subset Q$. A process P is closed if $P \subset \Sigma_e^*$.*

The refinement relation \preceq is a complete partial order on the processes. The

bottom (silent) process $\{\epsilon\}$ has but one trace: the empty string. The top (universal) process Σ^* includes all possible traces.

When composing processes, we merge their traces by interleaving their events and hiding mutual communication.

Definition 2.2 Let $u, v, w \in \Sigma^*$ be traces. We define the relation $u \mid v \vdash w$ (speak: u, v can combine to form w) by the following inference rules:

$$\frac{}{\epsilon \mid \epsilon \vdash \epsilon} \text{ (EPSILON)} \qquad \frac{u \mid v \vdash w \quad \sigma \in \Sigma_c}{u\sigma \mid v\bar{\sigma} \vdash w} \text{ (COMMUNICATION)}$$

$$\frac{u \mid v \vdash w \quad \sigma \in \Sigma}{u\sigma \mid v \vdash w\sigma} \text{ (L-EVENT)} \qquad \frac{u \mid v \vdash w \quad \sigma \in \Sigma}{u \mid v\sigma \vdash w\sigma} \text{ (R-EVENT)}$$

Example 2.3 Let $\Sigma_e = \{a, b, c, d\}$ and $\Sigma_c = \{e, \bar{e}\}$. Then we have

$$ab \mid cd \vdash acbd \quad ab \mid cd \vdash abcd \quad ae \mid \bar{e}b \vdash ab \quad ae \mid \bar{e}b \vdash ae\bar{e}b$$

but not $ae \mid \bar{e}b \vdash ba$.

Definition 2.4 Let P, Q be processes. Then $P \mid Q \doteq \{w \in \Sigma^* \mid \exists u \in P : \exists v \in Q : u \mid v \vdash w\}$.

Composition is commutative, which follows from Lemma E.1 below. It is also associative, which follows from Lemma E.2 below.

Lemma E.1 Let $u, v, w \in \Sigma^*$. If $u \mid v \vdash w$, then $v \mid u \vdash w$.

PROOF. We prove this by induction on the derivation of $u \mid v \vdash w$, with a case distinction on the type of the last rule used.

case (EPSILON):

In this case, $u = v$, so the claim is immediate.

case (L-EVENT):

In this case, we have a derivation for $u'\sigma \mid v \vdash w'\sigma$.

We have a derivation subtree for $u' \mid v \vdash w'$.

By induction, $v \mid u' \vdash w'$.

Apply (R-EVENT) to get $v \mid u \vdash w$.

case (R-EVENT):

In this case, we have a derivation for $u \mid v'\sigma \vdash w'\sigma$.

We have a derivation subtree for $u \mid v' \vdash w'$.

By induction, $v' \mid u \vdash w'$.

Apply (L-EVENT) to get $v \mid u \vdash w$.

case (COMMUNICATION):

In this case, we have a derivation for $u'\sigma \mid v'\bar{\sigma} \vdash w$.

We have a derivation subtree for $u' \mid v' \vdash w$. \square

By induction, $v' \mid u' \vdash w$.

Apply (COMMUNICATION) to get $v \mid u \vdash w$.

Lemma E.2 *Let $u_1, u_2, u_3, v, w \in \Sigma^*$. If the following two conditions hold:*

$$u_1 \mid u_2 \vdash v \tag{1}$$

$$v \mid u_3 \vdash w \tag{2}$$

Then there exists a word $z \in \Sigma^$ such that both of the following hold:*

$$u_1 \mid z \vdash w \tag{3}$$

$$u_2 \mid u_3 \vdash z \tag{4}$$

PROOF. The proof proceeds by induction on the sum of the depths of the derivation trees for (1) and (2). We do a case distinction on what inference rules are used at the bottom of those derivations. For each combination of inference rules, there is at least one matching proof case.

case (EPSILON),(EPSILON):

We are given derivations for (1),(2) of the form

$$\epsilon \mid \epsilon \vdash \epsilon \text{ and } \epsilon \mid \epsilon \vdash \epsilon$$

Set $z = \epsilon$.

This gives (3) by (EPSILON) and (4) by (EPSILON).

case (EPSILON),(L-EVENT):

case (EPSILON),(COMMUNICATION):

case (R-EVENT),(EPSILON):

case (L-EVENT),(EPSILON):

these combinations are not possible, because the variable v can not match both ϵ and $v'\sigma$.

case (COMMUNICATION),(any):

We are given derivations for (1),(2) of the form

$$u'_1\sigma \mid u'_2\bar{\sigma} \vdash v \text{ and } v \mid u_3 \vdash w$$

Furthermore, we have a derivation subtree for

$$u'_1 \mid u'_2 \vdash v$$

By the induction hypothesis, there exists a $z' \in \Sigma^*$ such that

$$(a) u'_1 \mid z' \vdash w \quad (b) u'_2 \mid u_3 \vdash z'$$

Now set $z = z'\bar{\sigma}$. Then

we get (3) from (a) by (COMMUNICATION),

and we get (4) from (b) by (L-EVENT).

case (any),(R-EVENT):

We are given derivations for (1),(2) of the form

$$u_1 \mid u_2 \vdash v \text{ and } v \mid u'_3\sigma \vdash w'\sigma$$

Furthermore, we have a derivation subtree for

$$v \mid u'_3 \vdash w'$$

By the induction hypothesis, there exists a $z' \in \Sigma^*$ such that

$$(a) u_1 \mid z' \vdash w' \quad (b) u_2 \mid u'_3 \vdash z'$$

Now set $z = z'\sigma$. Then

we get (3) from (a) by (R-EVENT),

and we get (4) from (b) by (R-EVENT).

case (L-EVENT),(L-EVENT):

We are given derivations for (1),(2) of the form

$$u'_1\sigma \mid u_2 \vdash v'\sigma \text{ and } v'\sigma \mid u_3 \vdash w'\sigma$$

Furthermore, we have derivation subtrees for

$$u'_1 \mid u_2 \vdash v' \text{ and } v' \mid u_3 \vdash w'$$

By the induction hypothesis, there exists a $z \in \Sigma^*$ such that

$$(a) u'_1 \mid z \vdash w' \quad (b) u_2 \mid u_3 \vdash z$$

Now we get (3) from (a) by (L-EVENT),
and (4) is (b).

case (R-EVENT),(L-EVENT):

We are given derivations for (1),(2) of the form

$$u_1 \mid u'_2\sigma \vdash v'\sigma \text{ and } v'\sigma \mid u_3 \vdash w'\sigma$$

Furthermore, we have derivation subtrees for

$$u_1 \mid u'_2 \vdash v' \text{ and } v' \mid u_3 \vdash w'$$

By the induction hypothesis, there exists a $z' \in \Sigma^*$ such that

$$(a) u_1 \mid z' \vdash w' \quad (b) u'_2 \mid u_3 \vdash z'$$

Now set $z = z'\sigma$. Then

we get (3) from (a) by (R-EVENT),

and we get (4) from (b) by (L-EVENT).

case (L-EVENT),(COMMUNICATION):

We are given derivations for (1),(2) of the form

$$u'_1\sigma \mid u_2 \vdash v'\sigma \text{ and } v'\sigma \mid u'_3\bar{\sigma} \vdash w$$

Furthermore, we have derivation subtrees for

$$u'_1 \mid u_2 \vdash v' \text{ and } v' \mid u'_3 \vdash w$$

By the induction hypothesis, there exists a $z' \in \Sigma^*$ such that

$$(a) u'_1 \mid z' \vdash w \quad (b) u_2 \mid u'_3 \vdash z'$$

Now set $z = z'\bar{\sigma}$. Then

we get (3) from (a) by (COMMUNICATION),

and we get (4) from (b) by (R-EVENT).

case (R-EVENT),(COMMUNICATION):

We are given derivations for (1),(2) of the form

$$u_1 \mid u'_2\sigma \vdash v'\sigma \text{ and } v'\sigma \mid u'_3\bar{\sigma} \vdash w$$

Furthermore, we have derivation subtrees for

$$u_1 \mid u'_2 \vdash v' \text{ and } v' \mid u'_3 \vdash w$$

By the induction hypothesis, there exists a $z \in \Sigma^*$ such that

$$(a) u_1 \mid z \vdash w \quad (b) u'_2 \mid u'_3 \vdash z$$

Now (a) already is (3),

and we get (4) from (b) by (COMMUNICATION).

□

Like in CCS [27], composition does not restrict its components: for processes P, Q we always have $P \preceq P \mid Q$. This follows from Lemma E.3. Furthermore, for all processes P , $P \mid \{\epsilon\} = P$. This follows from the above, and from Lemma E.4.

Lemma E.3 *For any $u \in \Sigma^*$, $u \mid \epsilon \vdash u$ and $\epsilon \mid u \vdash u$.*

PROOF. By induction on the length of u . If $u = \epsilon$, then the result is direct from rule (EPSILON). Otherwise, $u = u'\sigma$ for some $\sigma \in \Sigma$. By induction, $u' \mid \epsilon \vdash u'$ and $\epsilon \mid u' \vdash u'$. Application of the rules (L-EVENT) and (R-EVENT), respectively, then implies the result. \square

Lemma E.4 *Let $u, v \in \Sigma^*$. If $u \mid \epsilon \vdash v$ or $\epsilon \mid u \vdash v$, then $u = v$.*

PROOF. We prove that $u \mid \epsilon \vdash v$ implies $u = v$ (the other follows with Lemma E.1, using induction on the derivation of $u \mid \epsilon \vdash v$, with a case distinction on the last rule used).

case (EPSILON):

$u = v$ is immediate from the matching.

case (L-EVENT):

In this case, $u = u'\sigma$ and $v = v'\sigma$ for some $\sigma \in \Sigma$.

We have a derivation subtree for $u' \mid \epsilon \vdash v'$.

By induction, $u' = v'$, and therefore $u = v$.

case (R-EVENT):

case (COMMUNICATION):

These are impossible — they don't match ϵ .

\square

Refinement is preserved by composition: if $P' \preceq P$, then $P' \mid Q \preceq P \mid Q$. We can use this fact to prove that a system implementation refines its specification

$$P' \mid Q' \preceq P \mid Q \tag{5}$$

from the simpler, local refinement conditions

$$P' \preceq P \quad \text{and} \quad Q' \preceq Q \tag{6}$$

However, this method is not very powerful, because the refinements (6) do often not hold because of implicit assumptions on the context. Assume-guarantee reasoning remedies this shortcoming. We provide the context as an explicit subscript to the refinement relation, enabling us to conclude (5) from

$$P' \preceq_Q P \quad \text{and} \quad Q' \preceq_P Q \tag{7}$$

Most process models used for compositional refinement of hardware [2, 25] can express the contextual refinement $P' \preceq_Q P$ directly as $P' \parallel Q \preceq P$ (using synchronous parallel composition). The same does not work in our context (as exemplified by the observation 5 below), so we use a direct definition instead.

Definition 2.5 (Contextual refinement.) *Let P, P', C be processes. Then P' is said to refine P in context C , written $P' \preceq_C P$, iff for all traces $u \in P'$ the following condition holds: if there is a trace $v \in C$ such that $u \uparrow \Sigma_c = v \uparrow \Sigma_c$ (i.e. the communication events in u, v match up), then $u \in P$.*

Intuitively, we require that all behaviors of P' that are actually possible within an environment that adheres to C are allowed by P .

The following observations provide insight about contextual refinement.

1. For any process C , \preceq_C is a pre-order on processes.
2. If $P' \preceq_C P$, and $C' \preceq C$, then $P' \preceq_{C'} P$.
3. Refinement in a universal context corresponds to regular refinement:
 $P' \preceq_{\Sigma^*} P \Leftrightarrow P' \preceq P$.
4. Refinement in a silent context corresponds to refinement of closed processes:
 $P' \preceq_{\{\epsilon\}} P \Leftrightarrow (P' \cap \Sigma_e^*) \preceq (P \cap \Sigma_e^*)$
5. The refinement $P' \mid C \preceq_{\{\epsilon\}} P \mid C$ does not imply $P' \preceq_C P$, because the traces of $P' \mid C$ do not indicate what mutual communication takes place. However, the converse always holds.

The following Lemma shows how contextual refinement and composition are related to each other.

Lemma E.5 *Let $u, v \in \Sigma^*$. Then $u \uparrow \Sigma_c = \overline{v \uparrow \Sigma_c}$ iff there exists a $w \in \Sigma_e^*$ such that $u \mid v \vdash w$.*

PROOF. First, we show the implication ' \Rightarrow ', using induction on the sum of the lengths of u and v . We make a case distinction based on the kind of the last events in u and v .

- If $u = \epsilon$, then $u \uparrow \Sigma_c = \epsilon$, therefore $\overline{v \uparrow \Sigma_c} = \epsilon$, and therefore $v \in \Sigma_e^*$. Now choose $w = v$, and because we know $\epsilon \mid v \vdash v$ by Lemma E.3, we get $u \mid v \vdash w$.
- If $v = \epsilon$, then $\overline{v \uparrow \Sigma_c} = \epsilon$, so $u \uparrow \Sigma_c = \epsilon$, and therefore $u \in \Sigma_e^*$. Now choose $w = u$, and because we know $u \mid \epsilon \vdash u$ by Lemma E.3, we get $u \mid v \vdash w$.
- If $u = u'\sigma$ with $\sigma \in \Sigma_e$, then $u \uparrow \Sigma_c = u' \uparrow \Sigma_c$ and by induction, there must exist a $w' \in \Sigma_e^*$ such that $u' \mid v \vdash w'$. Now choose $w = w'\sigma$, and by applying (L-EVENT) we get $u \mid v \vdash w$.
- If $v = v'\sigma$ with $\sigma \in \Sigma_e$, then $v \uparrow \Sigma_c = v' \uparrow \Sigma_c$ and by induction, there must exist a $w' \in \Sigma_e^*$ such that $u \mid v' \vdash w'$. Now choose $w = w'\sigma$, and by applying (R-EVENT) we get $u \mid v \vdash w$.
- If $u = u'\sigma$ and $v = v'\tau$ with $\sigma, \tau \in \Sigma_c$, then (because $u \uparrow \Sigma_c = \overline{v \uparrow \Sigma_c}$) it must be the case that $\tau = \bar{\sigma}$ and $u' \uparrow \Sigma_c = \overline{v' \uparrow \Sigma_c}$. By induction, there must exist a $w \in \Sigma_e^*$ such that $u' \mid v' \vdash w$. Now we simply apply (COMMUNICATION) to get $u \mid v \vdash w$.

Now, we show the implication ‘ \Leftarrow ’, using induction on the derivation of $u \mid v \vdash w$, with a case distinction on the last rule used.

case (EPSILON):

In this case, $u = v = \epsilon$ and therefore trivially
 $u \uparrow \Sigma_c = \overline{v \uparrow \Sigma_c}$.

case (L-EVENT):

In this case, $u = u'\sigma$ and $w = w'\sigma$ for some $\sigma \in \Sigma$.
 Because $w \in \Sigma_e^*$, we must have $\sigma \in \Sigma_e$ and $w' \in \Sigma_e^*$.
 We have a derivation subtree for $u' \mid v \vdash w'$.
 By induction, $u' \uparrow \Sigma_c = \overline{v \uparrow \Sigma_c}$,
 from which it follows that $u \uparrow \Sigma_c = \overline{v \uparrow \Sigma_c}$.

case (R-EVENT):

In this case, $v = v'\sigma$ and $w = w'\sigma$ for some $\sigma \in \Sigma$.
 Because $w \in \Sigma_e^*$, we must have $\sigma \in \Sigma_e$ and $w' \in \Sigma_e^*$.
 We have a derivation subtree for $u \mid v' \vdash w'$.
 By induction, $u \uparrow \Sigma_c = \overline{v' \uparrow \Sigma_c}$,
 from which it follows that $u \uparrow \Sigma_c = \overline{v \uparrow \Sigma_c}$.

case (COMMUNICATION):

In this case, $u = v'\sigma$ and $v = v'\bar{\sigma}$ for some $\sigma \in \Sigma_c$.
 We have a derivation subtree for $u' \mid v' \vdash w$.
 By induction, $u' \uparrow \Sigma_c = \overline{v' \uparrow \Sigma_c}$,
 from which it follows that $u \uparrow \Sigma_c = \overline{v \uparrow \Sigma_c}$.

□

To avoid circularity in the assume-guarantee reasoning, we conservatively require that the specification processes can always engage in a subset of communication events $\Sigma_r \subset \Sigma_c$ that is sufficiently large, i.e. $\Sigma_r \cup \overline{\Sigma_r} = \Sigma_c$; in our case, we will take care of this requirement by having specification processes accept any message at any time¹. We use the following definition to formalize this property of processes.

Definition 2.6 *Let P be a process over Σ , and $\Sigma_r \subset \Sigma$ be an event subset. P is called Σ_r -enabled iff $\forall u \in P : \forall \sigma \in \Sigma_r : u\sigma \in P$.*

We now give the two proof rules for compositional refinement. The first rule is simpler, but restricted to two components. The second rule is a generalization suited for induction.

Theorem 2.7 *Let P, P', Q, Q', C be processes over $\Sigma = \Sigma_e \cup \Sigma_c$. Let $\Sigma_r \subset \Sigma_c$ such that $\Sigma_r \cup \overline{\Sigma_r} = \Sigma_c$. Then the following proof rules are sound:*

¹If this is not true by default, we could extend the specification to generate a special error event if it receives an unexpected message.

$$\frac{P' \preceq_Q P \quad P, Q \text{ are } \Sigma_r\text{-enabled} \quad Q' \preceq_P Q}{P' \mid Q' \preceq_{\{\epsilon\}} P \mid Q}$$

$$\frac{P' \preceq_{Q|C} P \quad P, Q \text{ are } \Sigma_r\text{-enabled} \quad Q' \preceq_{P|C} Q}{P' \mid Q' \preceq_C P \mid Q}$$

PROOF. We need only prove the second rule since the first can be obtained from the second by setting $C = \{\epsilon\}$.

Suppose we are given processes P, Q, P', Q', C that satisfy the hypotheses of the proof rule. Suppose that $u \in P' \mid Q'$ and $v \in C$ and $u \uparrow \Sigma_c = \overline{v \uparrow \Sigma_c}$. There must be a derivation for $u \in P' \mid Q'$, say $p \mid q \vdash u$ for some $p \in P'$ and $q \in Q'$. By Lemma E.5, there exists a $w \in \Sigma_c^*$ such that $u \mid v \vdash w$. Now Lemma E.6 tells us that in this situation, $p \in P$ and $q \in Q$. Because $p \mid q \vdash u$, this implies that $u \in P \mid Q$. Therefore, $P' \mid Q' \preceq_C P \mid Q$. \square

Lemma E.6 *Let P, Q, P', Q', C be processes, let $\Sigma_r \subset \Sigma_c$ such that $\Sigma_r \cup \overline{\Sigma_r} = \Sigma_c$, let both P, Q be Σ_r -enabled, let $P' \preceq_{Q|C} P$, let $Q' \preceq_{P|C} Q$ and let $p \in P'$, $q \in Q'$, $v \in C$.*

Now, if $p \mid q \vdash u$ for some $u \in \Sigma^$, and $u \mid v \vdash w$ for some $w \in \Sigma_c^*$, then $p \in P$ and $q \in Q$.*

PROOF. We prove this by induction on the derivation of $p \mid q \vdash u$, with a case distinction on the kind of rule used last.

case (EPSILON):

In this case, $p = q = \epsilon$ and therefore $p \in P$ and $q \in Q$.

case (L-EVENT):

In this case, we have $p = p'\sigma$, $u = u'\sigma$ and derivations

$$p'\sigma \mid q \vdash u'\sigma \text{ and } u'\sigma \mid v \vdash w$$

We have a derivation subtree for $p' \mid q \vdash u'$.

By Lemma E.7, there are $c' \in C$ and $w' \in \Sigma_c^*$ such that $u' \mid c' \vdash w'$.

By induction, $p' \in P$ and $q \in Q$.

By Lemma E.2, there is a $z' \in \Sigma^*$ such that

$$(a) p \mid z \vdash w \quad (b) q \mid c \vdash z$$

By Lemma E.5, (a) implies $p \uparrow \Sigma_c = \overline{z \uparrow \Sigma_c}$

Furthermore, (b) and $q \in Q$ imply that $z \in Q \mid C$.

Therefore, by $P' \preceq_{Q|C} P$, we have $p \in P$.

case (R-EVENT):

In this case, we have $q = q'\sigma$, $u = u'\sigma$ and derivations

$$p \mid q'\sigma \vdash u'\sigma \text{ and } u'\sigma \mid v \vdash w$$

We have a derivation subtree for $p \mid q' \vdash u'$.

By Lemma E.7, there are $c' \in C$ and $w' \in \Sigma_c^*$ such that $u' \mid c' \vdash w'$.

By induction, $p \in P$ and $q' \in Q$.

By Lemma E.1 and Lemma E.2, there is a $z' \in \Sigma^*$ such that

$$(a) \ q \mid z \vdash w \quad (b) \ p \mid c \vdash z$$

By Lemma E.5, (a) implies $q \uparrow \Sigma_c = \overline{z \uparrow \Sigma_c}$

Furthermore, (b) and $p \in P$ imply that $z \in P \mid C$.

Therefore, by $Q' \preceq_{P \mid C} Q$, we have $q \in Q$.

case (COMMUNICATION):

In this case, we have $p = p'\sigma$, $q = q'\bar{\sigma}$, and derivations

$$p'\sigma \mid q'\sigma \vdash u \text{ and } u \mid c \vdash w$$

We have a derivation subtree for $p' \mid q' \vdash u$.

By induction, $p' \in Q$ and $q' \in Q$.

Now, at least one of $\sigma \in \Sigma_r$ or $\bar{\sigma} \in \Sigma_r$.

Therefore, $p'\sigma \in P$ or $q'\bar{\sigma} \in Q$, and therefore $p \in P$ or $q \in Q$.

Say $q \in Q$.

Proceed as in the second part of the (L-EVENT)-clause above.

Say $p \in P$.

Proceed as in the second part of the (R-EVENT)-clause above.

□

Lemma E.7 *Let $u, v, w \in \Sigma^*$ such that $u \mid v \vdash w$. If u' is a prefix of u , then there exist prefixes v' of v and w' of w such that $u' \mid v' \vdash w'$.*

PROOF. We prove this first for the special case where $u = u'\sigma$. The general case can then be obtained by induction on the length difference between u and its prefix u' . For the special case $u = u'\sigma$, we use induction on the derivation of $u \mid v \vdash w$, with a case distinction on the type of rule used last.

case (EPSILON):

This case is impossible.

case (L-EVENT):

In this case, we have a derivation of $u'\sigma \mid v \vdash w'\sigma$.

We have a derivation subtree for $u' \mid v \vdash w'$.

This directly confirms the claim since w' is a prefix of w .

case (R-EVENT):

In this case, we have a derivation of $u'\sigma \mid v'\tau \vdash w'\tau$.

We have a derivation subtree for $u'\sigma \mid v' \vdash w'$.

By induction, there exist prefixes u'' of u' and v'' of v' such that $u' \mid v'' \vdash w''$.

This confirms the claim since v'' is a prefix of v and w'' is a prefix of w .

case (COMMUNICATION):

In this case, we have a derivation of $u'\sigma \mid v'\bar{\sigma} \vdash w$.

We have a derivation subtree for $u' \mid v' \vdash w$.

This directly confirms the claim since v' is a prefix of v .

□

For example, consider again the local refinement obligations (7). Suppose that

the specification processes P, Q can receive messages at any time. We can then apply the first proof rule to conclude that $P' \mid Q'$ refines $P \mid Q$, if there is no external communication, i.e., there are no other components in the system.

3 Token Coherence

In this section, we introduce a formal specification of the safety substrate of token coherence. This abstract protocol is a generalization of the MOESI token counting rules in Martin’s dissertation [21]. We then justify its use as a specification, by proving that it is coherent, and with its implementation that refines it.

3.1 Background: Cache Coherence

Cache coherence describes the contract between the memory system and the processor in a shared-memory multiprocessor. It is typically established at the granularity of a cache block. A memory system is cache coherent if for each block, writes are serialized, and reads get the value of the last write.

Definition 3.1 *Let V be the set of values of a fixed cache block, and $v_0 \in V$ the initial value. Let $\Sigma_{rw} = \{rd(v), wr(v) \mid v \in V\}$ be the alphabet of events, describing accesses to the block by some processor. Then the coherent traces of the system are given by the following regular language over Σ_{rw} :*

$$Coh = rd(v_0)^* \left(\bigcup_{v \in V} wr(v) rd(v)^* \right)^*$$

Token coherence, like many contemporary coherence protocols such as the popular MOESI protocol family [33], provides this strong form of coherence by enforcing a “single writer, multiple reader” policy².

Whether a cache block can be read and/or written is determined by its permission state, one of $\{L, L_r, L_{rw}\}$ (see also figure 1 for more detail). The cache coherence protocol then enforces the following important mutual exclusion property: *If one cache is in L_{rw} state, all other caches are in L state.*

For example, the popular MOESI protocol family [33] fits into our permission scheme as follows: the I state is L , both S and O are L_r , and both M and E are L_{rw} .

Token coherence guarantees mutual exclusion by requiring since an agent need by assigning a fixed number m of tokens to each cache block. These tokens are

²We are considering only the interface between the memory system and the processor here. Independently, the contract between the processor and the programmer may use weaker forms of coherence that involve temporal reordering of events, as specified by the memory model.




Figure 1: LTS model of caches.

stored in the caches and the memory, and may be carried around by protocol messages, but their total number is constant. The permission state of a cache corresponds to how many tokens it has: L for 0 tokens, L_r for $1 \dots (m - 1)$ tokens and L_{rw} for m tokens.

The mutual exclusion has additional applications beyond just cache coherence: some processor designs directly exploit the per-address mutual exclusion provided by the cache coherence protocol, to implement synchronization primitives, or for optimizations such as speculative loads or speculative retirement [31].

3.2 The abstract protocol

In our abstract protocol, system components and messages are of the same type and treated completely symmetrically: both are represented by token bags. Token bags are finite multisets (or bags) over some set T of tokens, and may be required to satisfy some additional constraints (well-formedness). The tokens in the bag constitute the state of the component, or the contents of the message.

The state of the entire system is represented as yet another bag that encloses the token bags of the individual components and messages. The sending of a message is modeled as a division, where a bag separates into two bags, dividing its tokens. The receipt of a message, symmetrically, is modeled as a fusion of token bags. Change is expressed by local reactions: tokens within a bag can be

consumed, produced or modified according to rewrite rules.

We now formalize the general behavior of such token transition systems, which prepares us for defining the actual abstract protocol.

Definition 3.2 (Multisets) *Let T be a set. Two words $u, v \in T^*$ are equivalent if one is a permutation of the other. The induced equivalence classes $\{[u] \mid u \in T^*\}$ are called finite multisets over T , or T -bags. Multiset union is defined as concatenation $[u] \smile [v] \doteq [uv]$. The set of all T -bags is denoted $\mathcal{M}(T)$. For $x \in \mathcal{M}(T)$, let $|x|$ denote the set of elements of T that occur in x .*

For example, for any $t_1, t_2 \in T$, all of the following denote the same T -bag: $[t_1^2 t_2] = [t_1 t_1 t_2] = \{t_1 t_1 t_2, t_1 t_2 t_1, t_2 t_1 t_1\}$. The exponent is a convenient notation for repeated symbols, and often used with regular languages.

Definition 3.3 (Token Transition System) *A TTS is a tuple (T, B, I, Σ_e, W) where T is a set of tokens, $B \subset \mathcal{M}(T)$ defines the set of well-formed T -bags, $I \in \mathcal{M}(B)$ is the initial configuration, Σ_e is a set of local events, and $W \subset \Sigma_e \times \mathcal{M}(T) \times 2^T \times \mathcal{M}(T)$ is a set of rewrite rules.*

A rewrite rule $(a, x, H, y) \in W$ is denoted $a: x \xrightarrow[H]{} y$. It describes a reaction labeled a that can occur whenever all the tokens in x are together in a bag, and the bag does not contain any of the inhibiting tokens listed in H . When the reaction fires, the tokens x are replaced by the tokens y . If H is empty, we omit it from the notation.

A TTS defines a process over the alphabet $\Sigma = \Sigma_e \cup \Sigma_c$, with $\Sigma_c = \{\text{snd}(b), \text{rcv}(b) \mid b \in B\}$, with the traces $\{u \in \Sigma^* \mid \exists C \in \mathcal{M}(B) : I \xrightarrow{u} C\}$, where we define the transition relation $C \xrightarrow{u} C'$ with the inference rules³ below.

³The variables in the rule templates range over the following domains: $u, v, w \in \Sigma^*$, $x, y, z \in B$, and $C, C', C'' \in \mathcal{M}(B)$. Furthermore, as a syntactic shortcut, we allow C, C', C'' to match several positions in a multiset of token bags: for example, $[C z]$ can match $[x y z]$ by setting $C = [x y]$.

$$\begin{array}{c}
\frac{}{C \xrightarrow{\epsilon} C} \text{ (STUTTER)} \qquad \frac{C \xrightarrow{u} C' \quad C' \xrightarrow{v} C''}{C \xrightarrow{uv} C''} \text{ (TRANS)} \\
\\
\frac{x \smile y \in B}{[C \ x \ y] \xrightarrow{\epsilon} [C \ x \smile y]} \text{ (FUSION)} \qquad \frac{}{[C \ x \smile y] \xrightarrow{\epsilon} [C \ x \ y]} \text{ (DIVISION)} \\
\\
\frac{a: x \xrightarrow{H} y \quad |z| \cap H = \emptyset \quad y \smile z \in B}{[C \ x \smile z] \xrightarrow{a} [C \ y \smile z]} \text{ (REACTION)} \\
\\
\frac{}{[C \ x] \xrightarrow{\text{snd}(x)} [C]} \text{ (SEND)} \qquad \frac{}{[C] \xrightarrow{\text{rcv}(x)} [C \ x]} \text{ (RECEIVE)}
\end{array}$$

Token transition systems have a feel of concurrency much like a biological system where reactive substances are contained in cells that can undergo fusion and division. Chemical abstract machines [6] capture the same idea (with molecules, membranes, and solutions instead of tokens, bags, and configurations), but are also different in many ways (for example, they do not have fusion or division).

Example E.8 *Suppose our tokens are electrons \ominus , positrons \oplus , and gamma rays γ , and we have a reaction where electrons and positrons annihilate each other, emitting a gamma ray:*

$$\text{fire:} \quad [\ominus \oplus] \Longrightarrow [\gamma]$$

Then the following is a trajectory of the system:

$$\begin{array}{l}
[\gamma \oplus \oplus][\ominus] \\
\rightarrow [\gamma \oplus][\oplus][\ominus] \\
\rightarrow [\gamma \oplus][\oplus \ominus] \\
\stackrel{\text{fire}}{\rightarrow} [\gamma \oplus][\gamma]
\end{array}$$

What happens is that the left bag splits off a bag containing a single \oplus . This bag is then absorbed by the right bag. The latter now contains the ingredients needed for the “fire” reaction, so it can react.

Definition 3.4 (The abstract protocol.) *The safety substrate T_m (where m is the number of tokens, a fixed parameter) is a TTS (T, B, I, Σ_e, W) where*

- *T contains the following tokens:*
 - R is a regular token as used by token coherence.
 - $O(s)$ is a owner token in one of two states $s \in \{C, D\}$ (clean or dirty).
 - $D(v)$ is an instance of the data, with value $v \in V$.
 - $M(v)$ is a memory cell containing the value $v \in V$.

- B is defined by imposing two conditions on a token bag $x \in \mathcal{M}(T)$:
 - if x contains data $D(v)$, then it must contain at least one regular token R or an owner token $O(s)$.
 - if x contains a dirty owner token $O(D)$, then it must contain data $D(v)$.
- $I \doteq [[R^{m-1} O(C) M(v_0)]]$.
- $\Sigma_e \doteq \{rd(v), wr(v), memread, memwrite, copy, drop \mid v \in V\}$.
- W consists of the rewrite rules shown in table 1.

Table 2 shows an example trajectory of the abstract protocol. Next, we explain the reaction rules and their interaction in some more detail.

$rd(v)$ reads a value from a data instance (it can be applied at any time, and does not modify the state). $wr(w)$ modifies a data token, and can only be applied if all m tokens (one owner token and $m - 1$ regular tokens) are present, and no other data copies are in the same bag (which guarantees that the data token being modified is the only one in the system).

To guarantee proper writebacks of modified data, a special owner token is used. The owner token records the clean/dirty state, i.e. whether the memory value is stale. When modifying data, the owner token is set to dirty. When the memory writes back the data ($memwrite$), the owner token is cleaned. $memread$ loads data from the memory only if there is a clean owner token, and thereby avoids reading stale data.

The rules $copy$ and $drop$ imply that data instances $D(v)$ can be freely copied or destroyed, subject only to the restriction enforced by B that all bags are well-formed — for example, whoever has the dirty owner token must keep at least one data instance.

We can now prove that the abstract protocol is coherent.

Theorem 3.5 *The closed system $T_m \cap \Sigma_e^*$ is coherent:*

$$(T_m \cap \Sigma_e^*) \uparrow \Sigma_{rw} \subset Coh$$

PROOF. To prove this, verify that (1) all of the following invariants hold in the initial state I and (2) prove (by induction on derivations) that if the invariants hold for a state C , they hold for any state C' such that $C \xrightarrow{u} C'$ for some $u \in \Sigma_e^*$.

1. The number of regular tokens R in the system is $m - 1$.
2. There is always exactly one owner token $O(s)$.
3. There is always exactly one memory cell $M(v)$.
4. All data instances $D(v)$ have the same values.

$rd(v):$	$[D(v)]$	\implies	$[D(v)]$
$wr(w):$	$[R^{m-1} O(s) D(v)]$	$\xRightarrow{\{D(v)\}}$	$[R^{m-1} O(D) D(w)]$
$memread:$	$[M(v) O(C)]$	\implies	$[M(v) O(C) D(v)]$
$memwrite:$	$[M(v) O(D) D(w)]$	\implies	$[M(w) O(C) D(w)]$
$copy:$	$[D(v)]$	\implies	$[D(v) D(v)]$
$drop:$	$[D(v)]$	\implies	$[]$

Table 1: The reaction rules of the abstract protocol.

Description	System trajectory
initial state	$[[M(v_0) O(C) R^{m-1}]_D []_{C_1} []_{C_2}]$
C_1 requests M	(requests are abstracted away)
D responds	
— read memory data	$\xrightarrow{memread} [[M(v_0) D(v_0) O(C) R^{m-1}]_D []_{C_1} []_{C_2}]$
— send data w/ tokens	$\xrightarrow{\epsilon} [[M(v_0)]_D [D(v_0) O(C) R^{m-1}] []_{C_1} []_{C_2}]$
C_1 receives response	$\xrightarrow{\epsilon} [[M(v_0)]_D [D(v_0) O(C) R^{m-1}]_{C_1} []_{C_2}]$
C_1 writes value v_1	$\xrightarrow{wr(v_1)} [[M(v_0)]_D [D(v_1) O(D) R^{m-1}]_{C_1} []_{C_2}]$
C_2 requests S	(requests are abstracted away)
C_1 responds	
— copy data	$\xrightarrow{copy} [[M(v_0)]_D [D(v_1) D(v_1) O(D) R^{m-1}]_{C_1} []_{C_2}]$
— send data w/ token	$\xrightarrow{\epsilon} [[M(v_0)]_D [D(v_1) O(D) R^{m-2}]_{C_1} [D(v_1) R] []_{C_2}]$

Table 2: A short example trajectory of the abstract protocol, representing a system with a memory D and two caches C_1 and C_2 . For clarification, token bags carry subscripts indicating the component that they represent. Those subscripts are *not* part of the abstract protocol.

5. If the owner token is clean, any data instances present have the same value as the memory cell.
6. If there is a data token, it contains the value of the last write. Otherwise, the memory does.

We proceed by induction on the derivation of $C \xrightarrow{u} C'$, with a case distinction on the last derivation used. Since the invariants are predicates on states, it is immediately clear that they are preserved in the case (STUTTER), and the case (TRANS) is a straightforward induction. The cases (FUSION) and (DIVISION) preserve the invariants because the latter are insensitive to how the tokens are partitioned into bags. The cases (SEND) and (RECEIVE) can not apply because we assumed $u \in \Sigma_e^*$. Finally, the rule (REACTION) is the only interesting case, and we will examine it in detail.

First, invariants 1, 2 and 3 must be preserved because the reactions preserve the number of R , O and M tokens. We examine the remaining invariants 4, 5, and 6 separately for each reaction.

Invariant 4 is obviously preserved by $rd(v)$, $memwrite$, $copy$ and $drop$. The rule $memwrite$ also preserves it, because invariant 5 guarantees that in this situation, the memory has the same value as all existing data tokens. Finally, the rule $wr(w)$ preserves the invariant because the data token that participates in the reaction must be the only data token in the whole system. This is because any other tokens would have to either be in a different bag (which, by B , they can't be without an accompanying R or O token, and we know that there are no R or O tokens in any other bag because by invariants 1 and 2, there is a fixed number of them, and they are all in this bag), or in the same bag (which they can't be because they would inhibit the reaction).

Invariant 5 is obviously preserved by $rd(v)$, $copy$ and $drop$. It is preserved by $wr(w)$, because it is vacuously satisfied after the reaction. Similarly, it becomes true after $memwrite$. And it is also preserved by $memread$, which creates a new data token with the same value as the memory.

Invariant 6 is true after $wr(w)$, because there is just one data token (as argued for Invariant 4). It is clearly preserved by $rd(v)$, $copy$ and $memwrite$. It is preserved by $drop$ because either the token being dropped is not the last one (in which case the remaining data tokens keep the value), or it is the last one (in which case the memory must be up-to-date, because of invariant 5 and the fact that the owner token must be clean, which is because a dirty owner token requires the presence of a data token, because of B). It is preserved by $memread$ because either the newly created token is not the only one (in which case invariant 4 helps us), or it is the only one (in which case we know that the value read from memory is the value of the last write).

Together, these invariants guarantee that all data instances $D(v)$ are always up-to-date; therefore, reads get the correct value which implies coherence. \square

All state is modeled by tokens, and there is no distinction between components

and messages. This symmetry points out interesting design directions. For example, we consider the memory cell $M(v)$ to be stationary. However, the formal token rules do not impose this restriction and could be used as an implementation guideline for a system with home migration.

3.3 How a TTS is an LTS

For completeness, and for use by later proofs, we show how our token transition systems can be understood as LTSs, and that our processes compose in the same way that the LTSs do. We include this section only in the extended version since it is quite technical in a boring way.

To define the LTS corresponding to a TTS, we use exactly the same inference rules as for defining the traces of a TTS, except that we do not include the (STUTTER) and (TRANS) rule (since they are redundant once we take the trace set of the LTS). To avoid possible confusion when referencing these rules, we added primes to the rule names in this section.

Definition E.9 *The token transition system $S = (T, B, I, \Sigma_e, W)$ defines a labeled transition system $LTS(S) \doteq (Q, q0, \Sigma \cup \{\epsilon\}, \delta)$ where*

$$\begin{aligned} Q &\doteq \mathcal{M}(B) \\ q0 &\doteq I \\ \Sigma &= \Sigma_e \cup \Sigma_c \\ \Sigma_c &\doteq \{\text{snd}(b), \text{rcv}(b) \mid b \in B\} \end{aligned}$$

and δ is defined by the inference rules

$$\begin{array}{c} \frac{x \smile y \in B}{[C \ x \ y] \xrightarrow{\epsilon}_\delta [C \ x \smile y]} \text{ (FUSION')} \qquad \frac{}{[C \ x \smile y] \xrightarrow{\epsilon}_\delta [C \ x \ y]} \text{ (DIVISION')} \\ \\ \frac{a: x \xrightarrow[H]{} y \quad |z| \cap H = \emptyset \quad y \smile z \in B}{[C \ x \smile z] \xrightarrow{a}_\delta [C \ y \smile z]} \text{ (REACTION')} \\ \\ \frac{}{[C \ x] \xrightarrow{\text{snd}(x)}_\delta [C]} \text{ (SEND')} \qquad \frac{}{[C] \xrightarrow{\text{rcv}(x)}_\delta [C \ x]} \text{ (RECEIVE')} \end{array}$$

Definition E.10 *For a labeled transition system $(Q, q0, \Sigma \cup \{\epsilon\}, \delta)$, states $q_1, q_2 \in Q$ and a word $v \in \Sigma^*$ we define: $q \xrightarrow{v} q'$ iff there exists a $k \geq 0$ and a sequence of transitions $q_0 \xrightarrow{v_1}_\delta q_1 \xrightarrow{v_2}_\delta \dots \xrightarrow{v_k}_\delta q_k$ such that $q_0 = q$, $q_k = q'$ and $v_1 v_2 \dots v_k = v$ (where $v_1 v_2 \dots v_k$ is ϵ for $k = 0$).*

Definition E.11 For an LTS $M = (Q, q_0, \Sigma \cup \{\epsilon\}, \delta)$, define the trace set

$$\mathcal{L}(M) \doteq \{w \in \Sigma^* \mid \exists q \in Q : q_0 \xRightarrow{w} q\}$$

Lemma E.12 Let P be the process defined by the token transitions systems S . Then $\mathcal{L}(\text{LTS}(S)) = P$.

PROOF. For the ‘ \subset ’ direction, we are given a word w and a transition sequence as in definition E.10. For each transition, we have a primed rule as listed under definition E.9. We can then transform this transition sequence into a derivation for $I \xRightarrow{w} C$, by taking, for each transition, the corresponding non-primed rule (as listed under definition 3.3), and glueing them together with (TRANS). This then shows that $w \in P$.

For the ‘ \supset ’ direction, we show that for all derivations of $C \xrightarrow{u} C'$, there exists a state sequence as in definition E.9, with $w = u$ and $q_0 = C$ and $q_k = C'$ (we do so by induction on the derivation of $I \xrightarrow{u} C$). From that, we conclude that for all u such that $I \xrightarrow{u} C$ for some C , we must have $u \in \mathcal{L}(\text{LTS}(S))$. \square

Definition E.13 For two LTSs $M_i = (Q_i, q_{0_i}, \Sigma \cup \{\epsilon\}, \delta_i)$ (with $i \in \{1, 2\}$) we define the composition $M_1 \mid M_2$ to be the LTS $(Q_1 \times Q_2, (q_{0_1}, q_{0_2}), \Sigma \cup \{\epsilon\}, \delta)$ where δ is defined by the inference rules

$$\begin{array}{c} \frac{q_1 \xrightarrow{u}_{\delta_1} q'_1 \quad u \in \Sigma \cup \{\epsilon\}}{(q_1, q_2) \xrightarrow{u}_{\delta} (q'_1, q_2)} \quad (\text{L-EVENT}') \\ \\ \frac{q_2 \xrightarrow{u}_{\delta_2} q'_2 \quad u \in \Sigma \cup \{\epsilon\}}{(q_1, q_2) \xrightarrow{u}_{\delta} (q_1, q'_2)} \quad (\text{R-EVENT}') \\ \\ \frac{q_1 \xrightarrow{\sigma}_{\delta_1} q'_1 \quad q_2 \xrightarrow{\bar{\sigma}}_{\delta_2} q'_2 \quad \sigma \in \Sigma_c}{(q_1, q_2) \xrightarrow{\epsilon}_{\delta} (q'_1, q'_2)} \quad (\text{COMMUNICATION}') \end{array}$$

Lemma E.14 For two LTSs $M_i = (Q_i, q_{0_i}, \Sigma \cup \{\epsilon\}, \delta_i)$ (with $i \in \{1, 2\}$), the following holds:

$$\mathcal{L}(M_1) \mid \mathcal{L}(M_2) = \mathcal{L}(M_1 \mid M_2)$$

Definition E.15 For two LTSs $M_i = (Q_i, q_{0_i}, \Sigma \cup \{\epsilon\}, \delta_i)$ (with $i \in \{1, 2\}$) we say that M_2 can simulate M_1 if there exists a weak simulation relation. A weak simulation relation is a relation $S \subset Q_1 \times Q_2$ such that

- (S1) $(q_{0_1}, q_{0_2}) \in S$
- (S2) for all $(q_1, q_2) \in S$, and for all transitions $q_1 \xrightarrow{u}_{\delta_1} q'_1$, there exists a state $q'_2 \in Q_2$ such that $(q'_1, q'_2) \in S$ and $q_2 \xRightarrow{u}_{\delta_2} q'_2$.

Lemma E.16 If M_2 can simulate M_1 , then $\mathcal{L}(M_1) \subset \mathcal{L}(M_2)$.

	GETX	GETS	Lockdown	Unlockdown	Data Owner	Data Shared	Ack	Ack Owner	Exclusive Compl.	Owned Compl.	Shared Compl.
O	d b j /NO	d j /NO	dd l/L			f q k	f q k		v i	w i	x i
NO	a b j	c j	a l/L		m f p k/O	f q k	f q k	n f p k/O	v i	w i	x i
L	j	j	l	l/NO	r k	r k	r k	s k	v i	w i	x i

Figure 2: The SLICC table for the memory controller. Rows show controller states, columns show events, and cells show transitions. For example, consider the upper left box. It states that if a Request-Exclusive message arrives while the controller is in state O, the actions d, b and j are executed in sequence, and the controller transitions to the NO state. Shaded cells indicate that an event is not expected to occur in the given state.

4 Implementation

In this section, we describe how we verified the safety of a detailed implementation of token coherence for an arbitrary number of caches. We describe how we used compositional verification to deal with the parametric character, and how we employed abstraction to handle the fine level of detail. We conclude with a list of discovered bugs.

4.1 The Protocol Implementation

The protocol implementation was developed by Martin et al. for architecture research on token coherence [21], and was extensively simulated prior to this work. It consists of finite state machines (FSM) for the cache and memory controllers, augmented with message passing capabilities. The FSMs are specified using the domain-specific language SLICC (Specification Language for Implementing Cache Coherence) developed by Martin et al.

The FSMs include all necessary transient states that arise due to the asynchronous nature of the protocol. The memory and cache controller amount to 600 and 1800 lines of SLICC code, respectively, a scale on which purely manual analysis methods are impractical, in particular because these low-level specifications are usually changed over time.

The SLICC compiler generates (1) executables for the simulation environment and (2) summary tables containing the control states, events and transitions in a human-readable table format⁴. Figure 2 shows the summary table for the memory controller, with its 3 states and 11 events. Note that some parts of

⁴More about the table format can be found in Sorin et al. [32]

the state, such as the number of tokens, or the actual data values, are stored in variables that are not visible in the summary table.

Due to lack of space, we can not reproduce the summary table for the cache controller (17 states and 20 events), and we can not explain further the meaning of the states and events. The complete SLICC code and interactive HTML-tables are online [22], along with implementations of three other cache coherence protocols.

4.2 Parametric Compositional Refinement Proof

Consider the system S'_n consisting of n caches C' , a directory controller D' (which is attached to the memory, and sometimes called memory controller), and an interconnection network N' . We consistently use primes for implementation processes to distinguish them from specification processes:

$$S'_n \doteq \underbrace{C' | C' | \dots | C'}_n | N' | D' \quad (8)$$

In the beginning, the memory holds all tokens. We define local specification processes as token transition systems:

$$\begin{aligned} D &\doteq T_m = (T, B, I, \Sigma_e, W) \\ C &\doteq (T, B, [[]], \Sigma_e, W) \\ N &\doteq (T, B, [[]], \Sigma_e, W) \end{aligned}$$

Because a token transition system already models all conceivable distributions of state over the system, no new behavior emerges if we compose. This “self-similarity” is an important feature in our induction proof, and expressed in the following general proposition:

Proposition E.17 *Let P_i be three processes ($i \in \{1, 2, 3\}$) defined by the token transitions systems $S_i = (T, B, I_i, \Sigma_e, W)$ (which are identical except for their initial states). Then, if $I_1 \sim I_2 = I_3$, we have $P_1 | P_2 = P_3$.*

PROOF. We adopt the LTS view of transition systems for this proof, as developed in section 3.3. We denote the respective labeled transition systems by $M_i \doteq \text{LTS}(S_i)$.

For the direction ‘ \supset ’, note that the transition system $M_1 | M_2$ can simulate M_3 because (1) we can first send all bags from M_2 to M_1 , which does not produce any observed behavior, and makes the states of M_1 match the state of M_3 , and (2) from there on, we can simulate all transitions of M_3 with identical transitions of M_1 , using (L-EVENT’) to embed M_1 ’s behavior in $M_1 | M_2$.

For the direction ‘ \subset ’, we let M_3 simulate $M_1 | M_2$ by the following relation:

$$S \doteq \{((q_1, q_2), q_3) \in (Q_1 \times Q_2) \times Q_3 \mid q_1 \sim q_2 = q_3\}$$

To show that S is a weak simulation, we need to show that it satisfies the properties (S1) and (S2) in definition E.15. (S1) is direct from the assumptions. For (S2), we make a case distinction on the kind of transition of $M_1 \mid M_2$.

First, assume we are dealing with (L-EVENT') or (R-EVENT'). We can then match all possible transitions (one of (FUSION'), (DIVISION'), (REACTION'), (SEND') or (RECEIVE')) because $q_1 \smile q_2 = q_3$ means that the bags that are involved in producing the transition in q_1 or q_2 can produce the same transition in q_3 .

Second, assume we need to match a (COMMUNICATION') transition. It turns out that we can simulate this by doing nothing at all: consider $q_1 \xrightarrow{\text{snd}(b)} q'_1$ and $q_2 \xrightarrow{\text{rcv}(b)} q'_2$. Then we must have $q_1 \smile q_2 = q_1' \smile q_2' = q_3$, so we are o.k. because $q_3 \xrightarrow{\epsilon} q_3$.

□

We now state the central result which (together with Theorem 3.5) allows us to verify the implementation components D' , C' and N' individually, each within an abstracted context rather than a fully instantiated system.

Theorem 4.1 *If the implementation processes satisfy the local refinement obligations*

$$D' \preceq_C D \tag{9}$$

$$C' \preceq_D C \tag{10}$$

$$N' \preceq_D C \tag{11}$$

then for all $n \in \mathbb{N}$, we have $S'_n \preceq_{\{\epsilon\}} T_m$, i.e., the system refines the formal token coherence protocol.

PROOF. Using proposition E.17, we obtain the following equalities, which allow us to split and merge specification processes as needed during the proof:

$$C \mid D = D \tag{12}$$

$$C \mid C = C \tag{13}$$

The following expresses the fact that a collection of n caches behaves just like a single cache, and we prove it by induction:

$$C'^{(n)} \preceq_D C \tag{14}$$

For $n = 1$, it coincides with (10). As for $n \rightsquigarrow (n + 1)$: the equation (12) can be expanded in the induction hypothesis (14) to yield

$$C'^{(n)} \preceq_{C \mid D} C \tag{15}$$

and the same expansion in (10) produces

$$C' \preceq_{C \mid D} C \tag{16}$$

Now, (15) and (16) are plugged into theorem 2.7 to get the following

$$C'^{(n)} \mid C' \preceq_D C \mid C$$

from which we can complete the induction by using (13).

Now, we add the network. First, we expand (12) in (11) to get

$$N' \preceq_{C \mid D} C \tag{17}$$

which can be plugged into theorem 2.7, along with (15), to get

$$C'^{(n)} \mid N' \preceq_D C \mid C \tag{18}$$

which can be reduced to the following using (13):

$$C'^{(n)} \mid N' \preceq_D C \tag{19}$$

Finally, we plug (9) and (19) into the theorem and get

$$C'^{(n)} \mid N' \mid D' \preceq_{\{\epsilon\}} C \mid D \tag{20}$$

whose right-hand side can be collapsed by (12) to complete the proof. \square

4.3 Discharging the Obligations

To discharge the remaining obligations, we used manual translation, abstraction, and annotation, and the explicit model checker Mur φ [13, 12]. The following steps give an overview of the method.

1. *Obtain models D' , C' for the memory and cache controller implementations.* This step involves translating the SLICC code to Mur φ , instrumenting it with the read/write events relevant for coherence, and abstracting both the state space and the message format. Fig. 3 shows snippets of translated code. The SLICC instructions that fell prey to the abstraction are in *slanted* face. For example, only a single cache block is modeled, therefore the code dealing with addresses is abstracted away. Also, message source and destination fields are irrelevant due to the deep symmetry of formal token coherence. Furthermore, two data values are sufficient⁵.
2. *Obtain good encodings for the specification/environment processes D , C .* We can take advantage (1) of the global system invariants established earlier and (2) of the fact that fusion and division are not observable. For example, the flattening map $[b_1 \ b_2 \ \dots \ b_k] \mapsto b_1 \smile b_2 \ \dots \ \smile b_k$ provides a canonical representative state. This means that *a single T -bag, rather than a multiset of T -bags, is sufficient to model the context.* The models we obtain this way are compact and contribute much to the state-space economy of our approach.

⁵Restricting the set of values is justified by the *data-independence* [34], which implies that we can freely substitute values in the traces.

```

rule "get Request-Excl in 0 state" transition(0, RequestExcl, NO) {
  (I_DirectoryState = state_0)      d_sendDataWithAllTokens;
  ==>                                b_forwardToSharers;
begin                                j_popIncomingRequestQueue;
  d_sendDataWithAllTokens();        }
  I_DirectoryState := state_NO;
endrule;                             action(d_sendDataWithAllTokens, "d") {
                                     peek(requestNetwork_in, RequestMsg) {
procedure d_sendDataWithAllTokens(); enqueue(responseNetwork_out, ResponseMsg) {
var                                     out_msg.Address := address;
  out_msg: I_message;                 out_msg.Type := CoherenceResponseType:DATA_OWNER;
begin                                  out_msg.Sender := id;
  out_msg.RType := DATA_OWNER;       out_msg.SenderMachine := MachineType:Directory;
  if !(I_Tokens > 0) then              out_msg.Destination.add(in_msg.Requestor);
    error "d: assertion failed. ";    out_msg.DestMachine := MachineType:L1Cache;
  endif;                               assert(directory[address].Tokens > 0);
  out_msg.Tokens := I_Tokens;          out_msg.Tokens := directory[in_msg.Address].Tokens;
  out_msg.DataBlk := I_DataBlk;       out_msg.DataBlk := directory[in_msg.Address].DataBlk;
  out_msg.Dirty := false;              out_msg.Dirty := false;
  I_Tokens := 0;                       out_msg.MessageSize := MessageSizeType:Response_Data;
  EVENT_MEMLOAD();                     }
  EVENT_SEND(out_msg);                 }
  EVENT_DROP();                         directory[address].Tokens := 0;
end;                                    }

```

Figure 3: The murphi code (left) is obtained from the SLICC code (right).

3. *Annotate the transitions of the implementation with matching specification transitions, and provide refinement maps.* For each transition of the implementation process, the annotations specify a sequence of transitions of the specification process. Fig. 3 shows such annotations in uppercase. The refinement maps are functions that map a controller state to its corresponding token bag.
4. *Run the model checker Mur ϕ separately for the two relevant obligations⁶ $D' \preceq_C D$ and $C' \preceq_D C$.*

Proposition 4.3 listed below describes how the contextual refinement is discharged. The state enumeration performed by the model checker effectively constructs and verifies the relation R , which describes the reachable states of the implementation process I within the abstract context C . The annotations provided by the user eliminate the need for existential quantification. The model checker also validates the assertions present in the implementation code.

Definition 4.2 *For a labeled transition system $(Q, q_0, \Sigma \cup \{\epsilon\}, \delta)$, states $q_1, q_2 \in Q$ and a word $v \in \Sigma^*$ we define: $q \xrightarrow{v} q'$ iff there exists a $k \geq 0$ and a sequence of transitions $q_0 \xrightarrow{v_1} q_1 \xrightarrow{v_2} \dots \xrightarrow{v_k} q_k$ such that $q_0 = q$, $q_k = q'$ and $v_1 v_2 \dots v_k = v$ (where $v_1 v_2 \dots v_k$ is ϵ for $k = 0$).*

⁶Theorem 4.1 lists three obligations, but we skip $N' \preceq_D C$ because it reduces to checking the reliability of the network, which is trivial at the given abstraction level

Proposition 4.3 *Let I , S and C be processes defined by the trace sets of the labeled transition systems $L_i \doteq (Q_i, q0_i, \Sigma \cup \{\epsilon\}, \delta_i)$ with $i \in \{I, S, C\}$. Let $\phi : Q_I \rightarrow Q_S$ be a function (the refinement map). If $R \subset Q_I \times Q_C$ is a relation such that*

- (R1) $(q0_I, q0_C) \in R$, and $\phi(q0_I) = q0_S$
- (R2) If $(q_I, q_C) \in R$ and $q_C \xrightarrow{u} q'_C$ for some $u \in \Sigma_e \cup \{\epsilon\}$, then $(q_I, q'_C) \in R$.
- (R3) If $(q_I, q_C) \in R$ and $q_I \xrightarrow{u} q'_I$ for some $u \in \Sigma_e \cup \{\epsilon\}$, then $(q'_I, q_C) \in R$ and $\phi(q_I) \xrightarrow{u} \phi(q'_I)$.
- (R4) If $(q_I, q_C) \in R$ and $q_I \xrightarrow{\sigma} q'_I$ and $q_C \xrightarrow{\bar{\sigma}} q'_C$ for some $\sigma \in \Sigma_c$, then $(q'_I, q'_C) \in R$ and $\phi(q_I) \xrightarrow{\sigma} \phi(q'_I)$.

then $I \prec_C S$.

PROOF. We first prove the following statement. Let the circumstances be as described above.

If $u, v \in \Sigma^*$ and $w \in \Sigma_e^*$ such that

$$u \mid v \vdash w \quad (21)$$

$$q0_I \xrightarrow{u} q_I \quad (22)$$

$$q0_C \xrightarrow{v} q_C \quad (23)$$

then both of the following hold:

$$(q_I, q_C) \in R \quad (24)$$

$$q0_S \xrightarrow{u} \phi(q_I) \quad (25)$$

The proposition then follows: for given $u \in I$, $v \in C$ such that $u \uparrow \Sigma_c = \overline{v \uparrow \Sigma_c}$, we know by Lemma E.5 that there exists a $w \in \Sigma_e$ such that $u \mid v \vdash w$. The above statement then implies $u \in S$. Because this is true for any such u, v , it implies $I \prec_C S$.

For the proof of the statement, we do a simultaneous induction on the structures of (21), (22) and (23). To make it more clear that this is a well founded induction, we can define an integer metric that strictly decreases: the sum of the depth of the derivation and the lengths of the sequences satisfies this property. The induction proceeds by case distinction on the types of the last derivation of (21) and the last transitions in the sequences (22) and (23).

case (any, any, ϵ):

In this case, (23) has the form $q0_C \xrightarrow{v} q'_C \xrightarrow{\epsilon} q_C$

By induction, $(q_I, q'_C) \in R$ and $q0_S \xrightarrow{u} \phi(q_I)$.

The first implies (24) by (R2).

The second is (25).

case (any, ϵ , any):

In this case, (22) has the form $q0_I \xrightarrow{u} q'_I \xrightarrow{\epsilon} q_I$

By induction, $(q'_I, q_C) \in R$ and $q0_S \xrightarrow{u} \phi(q'_I)$.

By (R3), we get (24) and $\phi(q'_I) \xrightarrow{\epsilon} \phi(q_I)$, and from there (25).

case ((EPSILON), any, any):

Note that the cases for nontrivial sequences (22) and (23)

are already covered (since they can only contain ϵ transitions).

Therefore, it suffices to consider

$$q_I = q0_I \text{ and } q_C = q0_C$$

For which (R1) implies the claims (24) and (25).

case ((R-EVENT), any, σ):

In this case, we have a derivation for $u \mid v'\sigma \vdash w'\sigma$

and we know $q0_C \xrightarrow{v'} q'_C \xrightarrow{\sigma} q_C$ with $v = v'\sigma$.

Clearly, $w' \in \Sigma_e^*$.

We have a derivation subtree for $u \mid v' \vdash w'$.

By induction, $(q_I, q'_C) \in R$ and $q0_S \xrightarrow{u} \phi(q_I)$.

The first implies (24) by (R2).

The second is (25).

case ((L-EVENT), σ , any):

In this case, we have a derivation for $u'\sigma \mid v \vdash w'\sigma$

and we know $q0_I \xrightarrow{u'} q'_I \xrightarrow{\sigma} q_I$ with $u = u'\sigma$.

Clearly, $w' \in \Sigma_e^*$.

We have a derivation subtree for $u' \mid v \vdash w'$.

By induction, $(q'_I, q_C) \in R$ and $q0_S \xrightarrow{u'} \phi(q'_I)$.

By (R3), we get (24) and $\phi(q'_I) \xrightarrow{\sigma} \phi(q_I)$, and from there (25).

case ((COMMUNICATION), $\sigma, \bar{\sigma}$):

In this case, we have a derivation for $u'\sigma \mid v'\bar{\sigma} \vdash w$

and we know

$$q0_I \xrightarrow{u'} q'_I \xrightarrow{\sigma} q_I \text{ and } q0_C \xrightarrow{v'} q'_C \xrightarrow{\bar{\sigma}} q_C$$

We have a derivation subtree for $u' \mid v' \vdash w$.

By induction, $(q'_I, q'_C) \in R$ and $q0_S \xrightarrow{u'} \phi(q'_I)$.

By (R4), we get (24) and $\phi(q'_I) \xrightarrow{\sigma} \phi(q_I)$, and from there (25).

□

The full Mur φ code is available online [8].

4.4 Results

The translation required about two days of work. This estimate assumes familiarity with token coherence, and some knowledge of the implementation. We found several bugs of varying severity, all of which were missed by prior random

simulation tests similar to those described by Wood et. al. [35]. Seven changes were needed to eliminate all failures (not counting mistakes in the verification model):

1. The implementation included assertions that do not hold in the general system. Although they were mostly accompanied by a disclaimer like “remove this for general implementation”, the latter was missing in one case.
2. The implementation was incorrect for the case where a node has only one token remaining and answers a Request-Shared. This situation was not encountered by simulation, probably because the number of tokens always exceeded the number of simulated nodes. We fixed the implementation, which involved adding another state to the finite state control.
3. Persistent-Request-Shared messages (which are issued if the regular Request-Shared is not answered within a timeout period) suffered from the same problem, and we applied the same fix.
4. The implementation copied the dirty bit from incoming messages even if they did not contain the owner token. Although this does not compromise coherence, it can lead to suboptimal performance due to superfluous write-backs. This performance bug would have gone undetected had we only checked for coherence, rather than for refinement of the abstract protocol.
5. After fixing bug 4, a previously masked bug surfaced: the dirty bit was no longer being updated if a node with data received a dirty owner token.
6. Two shaded boxes (i.e. transitions that are specified to be unreachable) were actually reachable. This turned out to be yet another instance of the same kind of problem as in bug 2.
7. Finally, another (last) instance of bug 2 was found and fixed.

As expected, the compositional approach heavily reduced the number of searched states. This kept computational requirements low, in particular considering that the results are valid for an arbitrary number of caches. Measurements were carried out on a 300MHz Pentium III ThinkPad.

# tokens	component	# states	# transitions	time
4	memory controller	92	1692	0.3s
8	memory controller	188	5876	0.6s
32	memory controller	764	83396	7.49s
4	cache controller	700	23454	1.4s
8	cache controller	1308	76446	4.6s
32	cache controller	4956	1012638	65.2s

5 Conclusions and Future Work

We make three main contributions. First, we formally verified the safety of a system-level implementation of token coherence, for an arbitrary number of caches. Second, we developed a general and formal specification of the safety substrate of token coherence, and prove its correctness. Third, we demonstrated that token coherence’s “design for verification” approach indeed facilitates the verification as claimed.

Future work may address the following open issues. First, the methodology does not currently address liveness. Second, other protocols or concurrent computations may benefit from the high-level abstraction expressed by token transition systems, and offer opportunities for compositional refinement along the same lines. Third, much room for automation remains: for example, we could attempt to integrate theorem provers with the SLICC compiler.

References

- [1] M. Abadi and L. Lamport. Conjoining specifications. *ACM Transactions on Programming Languages and Systems*, 17(3):507–535, May 1995.
- [2] R. Alur and T. A. Henzinger. Reactive modules. In *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science*, page 207. IEEE Computer Society, 1996.
- [3] R. Alur and B. Wang. Verifying network protocol implementations by symbolic refinement checking. In *Proceedings of the 13th International Conference on Computer-Aided Verification*, 2001.
- [4] T. Arons, A. Pnueli, S. Ruah, J. Xu, and L. D. Zuck. Parameterized verification with automatically computed inductive assertions. In *Computer Aided Verification*, pages 221–234, 2001.
- [5] Arvind and X. W. Shen. Using term rewriting systems to design and verify processors. *IEEE Micro*, 19(3):36–46, /1999.
- [6] G. Berry and G. Boudol. The chemical abstract machine. *Theoretical Computer Science*, 96(1):217–248, 1992.
- [7] S. Burckhardt, R. Alur, and M. M. K. Martin. Verifying safety of a token coherence protocol by compositional parametric refinement. In *Sixth International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI)*, LNCS 3385. Springer Verlag, 2005.
- [8] S. Burckhardt et al. Verifying safety of a token coherence implementation by parametric compositional refinement: Extended version. <http://www.seas.upenn.edu/~sburckha/token/>, 2004.

-
- [9] K. M. Chandy and J. Misra. *Parallel program design: a foundation*. Addison-Wesley Longman Publishing Co., Inc., 1988.
- [10] E. M. Clarke, O. Grumberg, and S. Jha. Verifying parameterized networks. *ACM Trans. Program. Lang. Syst.*, 19(5):726–750, 1997.
- [11] G. Delzanno. Automatic verification of parameterized cache coherence protocols. In *Computer Aided Verification*, pages 53–68, 2000.
- [12] D. L. Dill. The murphi verification system. In *Proceedings of the 8th International Conference on Computer Aided Verification*, pages 390–393. Springer-Verlag, 1996.
- [13] D. L. Dill, A. J. Drexler, A. J. Hu, and C. H. Yang. Protocol verification as a hardware design aid. In *International Conference on Computer Design*, pages 522–525, 1992.
- [14] E. A. Emerson and V. Kahlon. Reducing model checking of the many to the few. In *Conference on Automated Deduction*, pages 236–254, 2000.
- [15] S. M. German. Formal design of cache memory protocols in IBM. *Formal Methods in System Design*, 22(2):133–141, 2003.
- [16] S. M. German and A. P. Sistla. Reasoning about systems with many processes. *J. ACM*, 39(3):675–735, 1992.
- [17] G. J. Holzmann. Algorithms for automated protocol verification. *AT&T Tech. J.*, Jan./Feb. 1990.
- [18] Y. Kesten and A. Pnueli. Control and data abstraction: The cornerstones of practical formal verification. *International Journal on Software Tools for Technology Transfer*, 2(4):328–342, 2000.
- [19] R. P. Kurshan. *Computer-aided verification of coordinating processes: the automata-theoretic approach*. Princeton University Press, 1994.
- [20] R. P. Kurshan and K. McMillan. A structural induction theorem for processes. In *Proceedings of the Eighth Annual ACM Symposium on Principles of Distributed Computing*, pages 239–247. ACM Press, 1989.
- [21] M. M. K. Martin. *Token Coherence*. PhD thesis, University of Wisconsin-Madison, 2003.
- [22] M. M. K. Martin et al. Protocol specifications and tables for four comparable MOESI coherence protocols: Token coherence, directory, snooping, and hammer. http://www.cs.wisc.edu/multifacet/theses/milo_martin_phd/, 2003.
- [23] M. M. K. Martin, M. D. Hill, and D. A. Wood. Token coherence: decoupling performance and correctness. In *Proceedings of the 30th Annual International Symposium on Computer Architecture*, pages 182–193. ACM Press, 2003.

-
- [24] K. McMillan and J. Schwalbe. Formal verification of the Encore Gigamax cache consistency protocol. In *Proceedings of the International Symposium on Shared Memory Multiprocessing*, pages 242–51, Tokyo, Japan, 1991.
- [25] K. L. McMillan. A compositional rule for hardware design refinement. In *Proceedings of the 9th International Conference on Computer-Aided Verification*, pages 24–35, June 1997.
- [26] K. L. McMillan. Verification of an implementation of tomasulo’s algorithm by compositional model checking. In A. J. Hu and M. Y. Vardi, editors, *Computer Aided Verification*, volume 1427 of *Lecture Notes in Computer Science*, pages 110–121. Springer, 1998.
- [27] R. Milner. *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press, 1999.
- [28] S. Park and D. L. Dill. Verification of FLASH cache coherence protocol by aggregation of distributed transactions. In *Proceedings of the Eighth Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 288–296. ACM Press, 1996.
- [29] F. Pong and M. Dubois. Verification techniques for cache coherence protocols. *ACM Computing Surveys*, 29(1):82–126, 1997.
- [30] A. Ponse, S. A. Smolka, and J. A. Bergstra. *Handbook of Process Algebra*. Elsevier Science Inc., 2001.
- [31] P. Ranganathan, V. S. Pai, and S. V. Adve. Using speculative retirement and larger instruction windows to narrow the performance gap between memory consistency models. In *Proceedings of the Ninth Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 199–210. ACM Press, 1997.
- [32] D. J. Sorin, M. Plakal, A. E. Condon, M. D. Hill, M. M. K. Martin, and D. A. Wood. Specifying and verifying a broadcast and a multicast snooping cache coherence protocol. *IEEE Transactions on Parallel and Distributed Systems*, 13(6):556–578, 2002.
- [33] P. Sweazey and A. J. Smith. A class of compatible cache consistency protocols and their support by the IEEE futurebus. In *Proceedings of the 13th Annual International Symposium on Computer Architecture*, pages 414–423. IEEE Computer Society Press, 1986.
- [34] P. Wolper. Expressing interesting properties of programs in propositional temporal logic. In *Proceedings of the 13th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 184–193. ACM Press, 1986.

- [35] D. A. Wood, G. A. Gibson, and R. H. Katz. Verifying a multiprocessor cache controller using random test generation. *IEEE Design & Test*, 7(4):13–25, 1990.