

Protocol verification using flows: An industrial experience

John O’Leary
Intel

john.w.oleary@intel.com

Murali Talupur
Intel

murali.talupur@intel.com

Mark R. Tuttle
Intel

tuttle@acm.org

Abstract—We prove the parameterized correctness of one of the largest cache coherence protocols being used in modern multi core processors today. Our approach is a generalization of a method we described last year that uses data type reduction and compositional reasoning to iteratively abstract and refine the protocol and uses invariants derived from protocol “flows” to make the abstraction-refinement loop converge. Our prior work demonstrated the value of sequencing information that appeared within the linear flows describing a protocol in design documents. This paper extends the notion of flows to capture intricate scenarios seen in real industrial protocols and demonstrates that there is also valuable information in the interaction among flows. We further show that judicious use of flows is required to make the method converge and identify which flows are most suitable.

I. INTRODUCTION

We validated an extremely complex cache coherence protocol that will soon appear in multi-core processors from Intel. We used a generalization of the method we reported last year [26] based on the CMP method [20], [4], [14] augmented with message flows. This protocol, which we call LCP, is an intricate high-performance protocol that is designed to be scalable to large number of cores. Such intricate distributed protocols are especially susceptible to functional bugs that standard techniques like testing and simulation are unlikely to find and consequently formal verification is indispensable in their validation. We think LCP may be one of the largest, most complicated cache coherence protocol ever validated with formal methods. As one measure, the Flash cache coherence protocol, to which only a handful of formal methods have been successfully applied, has about 10 Boolean state variables per process and 16 different message types in all. In contrast, with over 70 Boolean state variables and around 50 message types, the state space for LCP is several orders of magnitude larger than Flash (see Section II).

While many techniques [22], [15], [11], [3], [18] have been proposed for parametric protocol verification, none of them scale well to large protocols, and those that do scale [9], [21] require an inordinate amount of manual effort to succeed. The CMP method [20], [4], [14], [26] is the only method for parametric verification we are aware of that scales to large protocols and is easy to use. It is an interactive proof method based on compositional reasoning that uses a model checker as a proof assistant. Though it combines the best of theorem proving and model checking, the main difficulty in applying

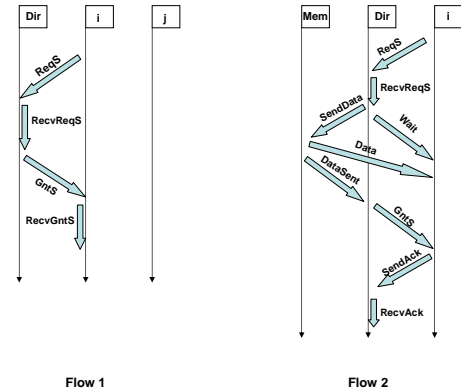


Fig. 1. Message flows as a linear sequence or acyclic graph of events.

this method is coming up the *non-interference lemmas* or invariants to guide the proof. As in theorem proving, this is a time consuming process requiring a thorough knowledge of the protocol. Moreover, adding one wrong invariant can lead the proof astray and render subsequent work useless.

In our earlier paper [26], we showed that the burden of generating the noninterference lemmas required by CMP can be significantly reduced by using the message *flows* typically found in industrial design documents. Flows are linear sequences of system events such as sending and receiving messages in case of distributed message passing systems, as illustrated on the left of Figure 1. We demonstrated the efficacy our method by applying it to several academic protocols, namely, German’s protocol and Flash protocol.

In this paper, we describe a generalization of the method presented in [26] and its application to the LCP cache coherence protocol. The primary contributions of this paper are:

- 1) Generalizing flows from linear traces to directed acyclic graphs, like the flow on the right of Figure 1, and a simple language for describing flows.
- 2) Demonstrating that we can derive powerful noninterference lemmas from constraints on events occurring in different flows, and not just constraints on events occurring within a single flow. Simply stating that two flows cannot be in progress at the same time, for example, can dramatically speed the convergence of the CMP method.
- 3) Demonstrating that not all flows are equally useful, and

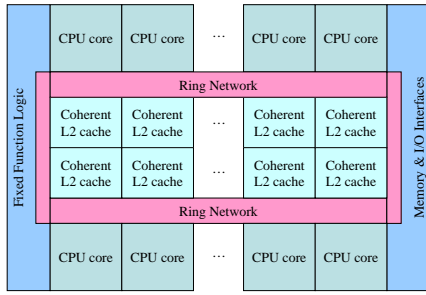


Fig. 2. Schematic of the Larrabee many-core architecture

that a more judicious use of the information in flows can also speed the convergence of the CMP method.

- 4) Parametrically verifying the correctness of the Intel cache coherence protocol LCP for any number of processors.

In verifying LCP we used a total of 15 flows, all easily obtained from the design documents, to derive around 36 lemmas. To make the CMP method converge another 5 lemmas had to be supplied by hand. A similar effort earlier [25] where we verified a cache protocol of comparable size using the CMP method required us to supply nearly 25 lemmas manually. Clearly, flows lead to a dramatic reduction in the number manually supplied lemmas and makes it much easier to use the CMP method.

The rest of the paper is structured as follows. In the next section we describe the salient features of the LCP protocol. In the following section III we discuss the possible alternatives to the CMP method and why they are inadequate. An overview of the CMP plus flows method of [26] is given in Section IV followed by a discussion of the extensions required to deal with LCP in the next section. In section VI, we present a new language to capture richer flows and also show how derive stronger constraints than just simple precedence constraints. A detailed description of our experience in using these extended flows in verifying the LCP protocol is given in Sec VII. Section VIII concludes the paper.

II. LARRABEE AND LCP

Larrabee is the code name for a many-core visual computing architecture under development at Intel Corporation [23]. The Larrabee architecture is based on a set of CPU cores that run the x86 instruction set, extended with support for vector processing operations and some specialized scalar instructions. Figure 2 shows a schematic of the architecture. Each core is associated with its own subset of a coherent L2 cache that affords fast, high-bandwidth data access to each core and simplifies data sharing and synchronization. The number of CPU cores and the number and type of co-processors and I/O blocks are implementation dependent, as are the positions of CPU and non-CPU blocks on the chip. To validate the LCP protocol in full generality we need parametric reasoning.

Figure 3 shows the major functional blocks in a single core. Larrabee’s global second-level (L2) cache is divided into




Fig. 3. Larrabee CPU core and associated system blocks

separate subsets, one for each CPU core. Each CPU has direct access to its own subset of the L2 cache. Data read by a CPU core is stored in its L2 cache subset and can be accessed quickly, in parallel with other CPUs accessing their own local L2 cache subsets. Data written by a CPU core is stored in its own L2 cache subset and is flushed from other subsets, if necessary. Larrabee uses a bi-directional ring network to allow agents such as CPU cores, L2 caches and other logic blocks to communicate with each other within the chip. The LCP (Larrabee coherence protocol) runs on the ring network and maintains coherency of shared data.

Our model of the Larrabee coherence protocol is organized as a parameterized number of identical caching agents which talk to a central directory that controls access to the data items as shown in Figure 4. For the purpose of verifying the coherence protocol our model abstracts away the ring structure and assumes point to point communication links between the agents (including links between the caches and from the off-chip memory to individual caches).

Unlike the Flash protocol where the directory distinguishes between local requests and external requests, the LCP makes no such distinction. This means when verifying two index properties it is enough to retain two cache agents concretely in the abstract model whereas for Flash we had to keep three agents, one local agent and two non-local agents [26]. In addition to these agents, there is also a memory controller that talks to the directory and supplies memory lines that have not yet been imported onto the chip.

The high level model we verified preserved much of the internal structure of each cache agent. Thus, apart from the L2 cache we also had the L1 cache and actions of the agent depended on the states of both the caches. Further, the various in- and out- message buffers and related book keeping data structures were also modeled. Other than assuming point to point links between the various agent, we modeled almost every significant detail of the protocol which increased the complexity of flows/transactions considerably.

The complexity of a protocol can be judged by the number of different types of messages that are exchanged by the agents. German’s protocol for example has only 7 different messages, Flash protocol, considered to be hard to verify, has 16 different types and LCP has around 50 different message types (comparable to the protocol we verified in [25]). In terms




Fig. 4. Basic Organization of the LCPmodel

of Boolean variables, each process in LCP has approximately 70 variables. In contrast the Flash protocol has around 10 state variables per process. Murphi description of the protocol actions was about 3000 lines whereas Flash has around 1000 lines.

III. PROTOCOL VERIFICATION TECHNIQUES

Broadly, there are two classes of techniques to verify distributed protocols: model checking based methods that aim for maximum automation and theorem proving style methods that aim for scalability.

A. Model Checking based Methods

Several techniques like Indexed Predicates [15], [16], Counter Abstraction [22] and related methods [10], Regular Model Checking [1], [2], Split Invariants based method [8], Environment Abstraction [6], [7] have been proposed to deal with verification of distributed protocols. While these methods have a higher degree of automation than the one we present, their scalability remains to be proven.

1) *Indexed Predicates Method*: Indexed predicate abstraction generalizes predicate abstraction to predicates that have free index variables. Given a protocol $P(N)$ and collection \mathcal{I} of indexed predicates this method automatically produces the strongest universally quantified invariant of $P(N)$ over \mathcal{I} . Intricate systems like Bakery protocol and Tomasulo's out of order processor has been verified by using this method [15]. But scalability remains an issue. Even for as small a protocol as the German's protocol this method takes couple of hours to produce an invariant.

2) *Cutoffs based approach*: Another approach to verifying a parameterized system $P(N)$ is to find a cutoff k such that verifying $P(k)$ is enough to guarantee the correctness of $P(N)$ for any value of N . There has been some work on this topic [11], [12] and related topics [13], [5] but the cutoffs are large making them impractical to use. For example, in [12] a cutoff of 7 was found for a directory based protocol. But real protocols are so large that even verifying a system with 3 agents is often not possible.

3) *Counter and Environment Abstractions*: Counter Abstraction [22] and its generalization Environment Abstraction [6] are based on the idea of partitioning a collection of identical agents into equivalent classes based on the predicates they satisfy and for each partition tracking only one representative. The abstract model produced by these methods tend to be very detailed and consequently too large as we look at bigger protocols. Environment abstraction for example is just able to handle a simplified version of the Flash protocol [24].

B. Theorem Proving style methods

Apart from classical theorem proving there are methods like aggregated transactions [21] that are user-guided techniques. These suffer from the well known problem of having to provide guidance in minute detail to get the proof through. It is unlikely that theorem proving style methods can be used practically to verify large protocols given that just to verify Flash protocol the aggregated transactions method took couple of days of effort.

C. CMP method

The CMP method straddles both the above categories of model checking and theorem proving based methods: it uses model checker as a proof assistant to carry out user guided proof. The crucial advantage of the CMP method is that the user supplied lemmas don't have to add up to an inductive invariant [26]. This means the amount of guidance provided is a lot less than in theorem proving methods. Using CMP method we have earlier verified Flash protocol [26] and another large cache coherence protocol within Intel [25]. The latter protocol is comparable in size to the LCP protocol. Clearly, CMP method is the only viable method currently for handling large protocols and our effort was to make it more usable by reducing the lemma burden as much as possible.

IV. DESCRIPTION OF THE CMP METHOD

For the rest of the paper we will use the same system model as in [26]. In particular, we consider a symmetric protocol P with N processors $[1..N]$ whose transition relation is given as a collection of rules. Each rule is a *guarded command* written as

$$rl : \forall i, j. \rho \rightarrow a \quad \text{or} \quad rl(i, j) : \rho(i, j) \rightarrow a(i, j)$$

where rl is the *rule name*, ρ is an expression called the *guard*, a is a list of assignments called the *action* and i, j are the process index variables.

The CMP method is a compositional reasoning based method and it consists of two basic steps — *abstraction* and *strengthening* — that are applied iteratively to a protocol as shown in Figure 5.

Given a property $I = \forall i, j \in [1..N]. I(i, j)$ that we want to prove is an invariant of P , the CMP method creates an abstract model \hat{P} that retains two agents, say 1, 2 without loss of generality, and replaces the rest of the processes with a highly non-deterministic process *Other*. Intuitively, in a




Fig. 5. The CMP method

symmetric system if a property $\phi(1, 2)$ holds for processes 1, 2 then $\phi(i, j)$ holds for any other pair of processes i, j as well. Thus, it is enough to consider an abstract model with detailed information on 1, 2 and check if $I(1, 2)$ holds. The abstraction process in CMP method is inexpensive as it is almost syntactic and bulk of the state space of \hat{P} comes from 1, and 2. On the flip side \hat{P} tends to be very coarse as the behavior of *Other* is completely unconstrained. To get rid of the resulting spurious counter examples, the CMP method requires the user to provide *non-interference lemmas* or invariants that are used to refine the abstract model. This process is continued iteratively till we find a real counter example or prove that $I(1, 2)$ holds. Pseudo-code for the method is given below.

```

CMP(P, I) =
  P# = P; I# = I
  while abstract(P#)  $\not\models$  abstract(I#(1, 2)) do
    examine counterexample cex
    exit if cex is a real counterexample
    find L =  $\forall i, j. L(i, j)$  ruling out cex
    P# = strengthen(P#, L)
    I#(i, j) = I#(i, j)  $\wedge$  L(i, j)
  end
  
```

If the loop terminates normally, the method and protocol symmetry allow us to conclude that $I^\#$ and consequently I are invariants of P . If the loop terminates via the exit, then either I or one of the proposed lemmas L is not an invariant of the protocol, and the user must back up and try again.

In McMillan's work [20] the abstraction operation used was data type reduction [19] which essentially throws away all the state variables of processes $[3..N]$. Our analysis of the CMP method in [26] allows us to use richer abstractions than data type reduction. In particular this allows meaningful abstraction of auxiliary variables used to track flows.

A. Making CMP better using Flows

Not surprisingly the main difficulty in applying the CMP method is coming up with the non-interference lemmas. As demonstrated in [26], message flows or simply flows yield powerful invariants that can be used as non-interference lemmas in the CMP method and thus reduce the number of lemmas we need to supply by hand.

The flows used in [26] are linear ordering of events usually involving two agents, see `reffi:flow`. Each entry in the flow is either a simple event, corresponding to a single rule firing

of the protocol or a *sub-flow* where a sub-flow is itself a flow composed of simple events. The notion of sub-flow serves to chop up a complicated flow into smaller units such that each unit shows interaction between two agents.

The constraints derived from flows are the implicit precedence constraints between events occurring in the flows. For example, according to the first flow in Figure 1, *ReqS* has to happen before *RecvReqS* can happen. This can be converted into a precise lemma by having a set $Aux(i)$ of auxiliary variables for each process i that track all the flows i is involved in and for each flow the last rule that was fired by i . In case of *RecvReqS* action the precedence constraint is simply that if, for process i , the *RecvReqS* action is enabled then there must be an auxiliary variable recording the fact that i was involved in *ReqS* action earlier. These simple precedence constraints turn out to be surprisingly powerful as invariants.

The advantage of the flows is that they are intuitive to understand and readily available in the design documents. Flows in fact allow us to state the core ideas that go into the design of a protocol in terms of higher level concepts while avoiding the specific implementation details. This means flows are quite robust and resistant to changes in the protocol design which makes them very attractive as user supplied annotations

V. EXTENSIONS TO FLOWS

Apart from the local caching agents and the directory, real cache protocols have other types of agents, like the off-die memory controller, which add to the complexity of the interaction between the agents. Flows designed to capture two agent interactions are no longer sufficient to capture complicated scenarios scene.

Consider Flow 2 of Figure 1 which calls for an extension to our notion of flows. Here a process i requests access to an item that has not yet been fetched onto the chip. The on-chip directory forwards the request to the off-chip memory controller along with id of the requesting agent. The directory also sends an acknowledge message to the caching agent. The memory controller for its part sends the required item to the caching agent and also sends a message completion to the directory. On receiving the message completion the directory sends a grant message to the agent. On receiving both the grant message from directory and the data message from memory controller agent i transitions to shared state and sends a completion message to the directory. The transaction ends when the directory receives the completion message.

This scenario is similar to, and typical of, the complex interactions present in LCP and it differs from the flow shown in the left of Figure 1 in crucial ways. The interaction between the three agents is tightly coupled and it is not possible to identify meaningful sub-flows that have only two agents involved. Earlier, in [26] even though we had more than two agents involved, it was too easy to see that the flow was made up of logical sub-units consisting of only two agents.

Further, an event might have multiple preceding actions. For instance, event *SendAck* can happen only after the agent

has received both the data and grant messages. Receiving only one of these is not sufficient to enable the *SendAck* message. Linear flows cannot capture multiple dependency.

Similarly, an event in the flow might have more than one succeeding event. For instance, the *RecvReqS* event leads to two further events *SendData* and *Wait*. With linear flows then the number of events depending on a given event is at most one.

Finally, unlike in the Flow 1 of Figure 1 which total orders on events, Flow 2 is a partial order. For instance, there is no ordering between the events *Data* and *GntS*, they can be received by *i* in any order. One way to deal with this is to flatten the partial order into a collection of total orders. But this runs into two issues: firstly, the number of resulting flows might be large and unnecessarily obscure the simplicity present in the dag representation. Secondly, having more flows also means we will have to introduce more auxiliary variables or extend their ranges and thus, we will also end up making the augmented model bigger.

Thus, it is clear that the new flows or rather flow language has to be expressive enough to capture directed acyclic graphs (dags). Moreover, note that in Flow 2 of Figure 1 there are three primary agents in the flow but that each event involves at most two agents. So apart from specifying events we also have to specify which agents are involved in the events. A new flow language that takes into account these extensions is described in the next section.

VI. GENERALIZED FLOWS

We now extend our prior work to incorporate the generalized flows and flow-based invariants described in previous sections. We omit treatment of subflows in this section since they are not required for LCP.

A. Language for Flows

In our language for flows, a flow is denoted by

$$(flow, conflicts) : \{prec_1, \dots, prec_n\}$$

where *flow* is the name of the flow, *conflicts* is a set of flows, and each *prec_i* is a *precedence* of the form

$$(rule, id, agents) : \{(rule_1, id_1, agents_1), \dots, (rule_m, id_m, agents_m)\}$$

The meaning of a precedence is that each rule *rule_i* instantiated with the list of agents *agents_i* must occur in an execution before the rule *rule* instantiated with the list of agents *agents* can occur. We say each $(rule_i, id_i, agents_i)$ *precedes* $(rule, id, agents)$ in the flow. For example, in the flow on the right side of Figure 1, the first precedence is

$$(ReqS, id_1, \langle Dir, i \rangle) : \{ \}$$

and the third precedence is

$$(SendData, id_2, \langle Mem, Dir \rangle) : \{(RecvReqS, id_3, \langle Dir \rangle)\}$$




Fig. 6. Flows can share a common event (left) and a common prefix (right).

The meaning of a conflict set is that a flow *f* and a flow *f'* in the conflict set of *f* cannot be alive at the same time. A flow is said to be *alive* if some rule in the flow has occurred and one of its successors has not occurred. The conflict set of *f* might contain *f* itself. In this case the meaning is that there can be only one instantiation of *f* alive in the system. The constraints resulting from these conflict sets turn out to be very powerful in constraining the *Other* process.

Finally, we associate an id with each event since a particular instantiation of a rule may occur in multiple flows. The left side of Figure 6 illustrates that we must distinguish the occurrences of *ev* in the flows *f₁* and *f₂*, so that the occurrence of *ev* in *f₁* will enable the event *e₃* in *f₁* and not *e₇* in *f₂*. The right side of Figure 6 illustrates that we cannot use the flow name as the event id, since different flows can share a common prefix. As long as the system is in the prefix, we don't know whether the system is in flow *f₁* or *f₂*. In fact, because we don't know which flow the system is in, an event in the prefix must have the same event id in both flows.

B. Tracking Flows

We use auxiliary variables to track active flows as in [26]. We will assume that a pair (rl, id) appears only once in a flow.

Define $last(rl, id, p)$ in a flow *f* to be the set of pairs (rl', id') for which there exist agent lists *ag* and *ag'* both containing *p* such that (rl', id', ag') precedes (rl, id, ag) in *f*. We require that $last(rl, id, p)$ be the same for all flows containing (rl, id) . Intuitively, this means the prefix for (rl, id) must be the same in all flows in which it appears.

Define $next(rl, id, p)$ in a flow *f* to be the set of pairs (rl', id') for which there exist agent lists *ag* and *ag'* both containing *p* such that (rl, id, ag) precedes (rl', id', ag') in *f*.

Define the *out degree* of (rl, id) for *p* in *f* to be the size of $next(rl, id, p)$ in *f*. We require that the out degree of (rl, id) for *p* be the same in all flows containing (rl, id) .

Define $flows(rl, id)$ to be the set of flows *f* containing an event (rl, id, ag) for some agent list *ag*. This is the set of all flows containing (rl, id) and by definition one of them is active when an agent executes (rl, id) .

For each process *p* we have a set $Aux(p)$ of tuples of the form $(rl, id, out_deg, fl_set)$ where *rl* is the rule name, *id* is the associated id, *out_deg* is an out degree, *fl_set* is a set of flows. Intuitively, *fl_set* represents the set of flows that might be active when *p* executes (rl, id) given what we have

```

update(p, rl(i, j), id) =
  for each (rl', id') ∈ last(rl, id, p)
    aux = choose (rl', id', od, fl_set) ∈ Aux(p)
      such that flows(rl, id) ∩ fl_set ≠ {}
    new_aux = if od > 1
      then {(rl', id', od - 1, fl_set ∩ flows(rl, id))}
      else {};
    Aux(p) := Aux(p) \ {aux} ∪ new_aux
  if (rl, id) precedes no other rule then
    return Aux(p) else
    return Aux(p) ∪ (rl, id, outdegree(rl, id), flows(rl, id))

```

Fig. 7. Tracking flows with auxiliary variables. This procedure describes how p updates $Aux(p)$ when an event (rl, id, ag) involving p is executed. The choose operator throws an error if there is nothing to choose.

seen thus far; and out_deg is the number of rules (rl', id') preceded by (rl, id) that p has not yet executed. The set fl_set is initialized to $flows(rl, id)$. The out degree out_deg is initialized to the out degree of (rl, id) for p . (It is sufficient for us to maintain a count like out_deg , even though we could track flows more precisely by maintaining the actual set of rules in $next(rl, id)$ yet to fire.)

Whenever an instantiation $rl(p_1, p_2)$ of a rule $(rl(i, j))$ fires, we update the auxiliary variables of p_1 and p_2 as follows¹. To associate this rule firing with an (rl, id) pair mentioned in the flows, we identify the set I of all the ids id such that (rl, id) appears in the flows. For which ever $id \in I$ the update procedure shown in Figure 7 goes through without raising an error that is the id we associate with the rule firing and we let the effects of the update procedure stay (after undoing the effects of the previous tries). For LCP, corresponding to each rule rl there was only one id that appeared in the flows, so updating auxiliary variables was simpler than the general procedure given here.

C. Lemmas from Flows

We derive two classes of lemmas from flows. The first class is extension of the “precedence” lemmas derived in [26] to the richer flows. The second class uses the conflict sets.

1) *Precedence Constraints*: For a rule $rl(i, j)$, we find all pairs (rl, id) associated with it. For index i (and similarly for j) and for all pairs (rl, id) we compute the set $last(rl, id, i)$. This set essentially is the set of all (rl', id') pairs that must fire before (rl, id) can fire. For each (rl', id') we have a constraint

$$\begin{aligned} \exists od, fl_set. \quad & (rl', id', od, fl_set) \in Aux(i) \\ & \wedge od > 0 \\ & \wedge fl_set \cap flows((rl, id)) \neq \{\} \end{aligned}$$

Define the precondition $pre_i(rl, id)$ for i firing (rl, id) to be the conjunction of the above constraints for all $(rl', id') \in last(rl, id, i)$. Define $pre_j(rl, id)$ in the same way. Define the precondition $pre(rl, id)$ for firing (rl, id) to be the conjunction

$$pre_i(rl, id) \wedge pre_j(rl, id).$$

¹The case where rl has more/less than 2 agents is handled similarly.

Define the precondition $pre(rl)$ for firing rl to be the disjunction

$$\bigvee_{id} pre((rl, id)).$$

If the guard for the rule $rl(i, j)$ is $\rho(i, j)$ the non-interference lemma is

$$\rho(i, j) \Rightarrow pre(rl).$$

2) *Conflict Constraints*: Given two flows f_1 and f_2 that are conflicting, we derive constraints as follows. Let the top level events, that is those that have no preceding events, appearing in f_2 be $(rl_1^2, id_1^2, ag_1^2), \dots, (rl_n^2, id_n^2, ag_n^2)$. Let $\mathcal{R} = \{(rl_1^1, id_1^1), \dots, (rl_m^1, id_m^1)\}$ be the set of rule, id pairs appearing in flow f_1 . Conflict between f_1 and f_2 means while f_1 is alive f_2 cannot start and vice-versa. Considering the first case if any of the events (rl_k^1, id_k^1) from f_1 has occurred (and flow f_1 has not ended) then guards of rl_1^2, \dots, rl_n^2 cannot be enabled for any process. Mathematically,

$$\begin{aligned} (\exists p. (rl_k^1, id_k^1, _, _) \in Aux(p) \text{ for some } (rl_k^1, id_k^1) \in \mathcal{R}) \\ \Rightarrow \forall rl_m^2 \forall i, j. \neg \rho_m^2(i, j) \end{aligned}$$

where $\rho_m^2(i, j)$ is the guard of the rule rl_m^2 .

This lemma prevents the start of f_2 if f_1 is still alive. We can derive similar invariant for the case other way around.

We can derive another class of invariants that are quite useful in constraining the abstract model though at first glance these invariants don't seem useful. Suppose f belongs to its own conflict set. Consider a rule, id pair (rl_k, id) appearing in f and a process i that has just fired that rule. It is clear that i cannot fire the same rule again until the current flow has ended because f is in conflict with itself. Formally,

$$\exists (rl_k, id, _, _) \in Aux(i) \Rightarrow \neg j. \rho_k(i, j)$$

This is similar in spirit to the conflict constraints presented above except that instead of preventing rules at the beginning of other flows from firing this constraint prevents same rule from firing twice in a flow. This turns out to be surprisingly useful in removing spurious behaviors from of the *Other* process in the abstract model.

VII. VERIFYING THE LCP PROTOCOL

In this section we describe our experience applying the CMP method in conjunction with flow based lemmas to verify the LCP protocol.

A. Obtaining the Flows

The flows we considered during verification were all readily available in the design documents written by the architects. In fact, the scenarios listed in the design documents had more information than we needed. Intuitively, only those parts of the flows that involve the directory (in other words the place where all the coordination happens) yield useful invariants.

Apart from identifying the useful parts of the flows, we had to annotate the events with agents involved and also identify the incompatibility set for each flow. Both these steps were straight-forward.

B. Protocol Model

The Murphi model that we verified was quite hierarchical with each rule having an extremely large body covering a variety of cases (as it was semi-automatically generated from tables describing the protocol). For instance, there is only one rule specifying the behavior of a cache agent in case it receives an invalidate message. The guard of the rule only checks the type of the incoming message and the body of the rule considers the various sub-cases depending on the state of the L1/L2 caches and the data structures. The behavior of the cache in each of the sub-cases might be quite different with differing messages being sent out.

In other words, the events associated with each of the sub-cases are quite different though they all belong to the same Murphi rule syntactically. This is not conducive to our flow based methodology which depends on being able to track events precisely. To allow precise tracking of flows, we broke up large rules into smaller ones so that each rule corresponded to a specific event mentioned in the flows. This was accomplished by simply lifting some of the branch conditions in the body of a rule to the guard². After this step each event mentioned in the flows mapped to a Murphi rule.

C. Proof details

The properties that we proved were the standard safety property requiring that if there is a cache with exclusive access to a data item then no other cache has access to that item and properties specifying consistency between the directory's list of caches with access to a data item and the real access each local cache has.

To carry out the abstraction and refinement with lemmas we used a tool called *Abster* written in Ocaml. *Abster* takes in a parameterized protocol written in Murphi and creates an intermediate AST for the protocol. All the abstraction and refinement operations take place on the AST. We specify the number of agents, usually 2 or 3, to be retained concretely in the abstract model depending on the protocol to be verified. The flows are specified in a separate file. *Abster* automatically constructs flow lemmas and adds them to the guards of the appropriate rules (a rule whose guard appears in the flow invariant gets modified by that invariant).

To carry out the proof we used 15 flows from design documents. These led to 36 flow lemmas, with 25 of them coming from precedence constraints and the rest from conflict constraints. At first we used only the precedence constraints and proceeded to refine the abstract model with hand supplied lemmas. But we soon realized that several counter examples were caused by two conflicting flows starting at the same time. More precisely, suppose one of the concrete agents, say 1, is

involved in a *Request Shared* transaction. Since the *Other* process is not fully constrained, it might start sending conflicting messages corresponding to *Request Exclusive* transaction to the directory. One way to prevent this is to write lemmas by hand and refine the *Other* process. The simpler option is to use the conflict constraints that prevent a *Request Exclusive* transaction from starting while a *Request Shared* transaction is alive. Together with precedence constraints this ensures that no rule in a conflicting flow can fire.

Another cause of spurious counter examples was the *Other* process repeatedly firing the same rule from a flow. Suppose (rl', id') precedes (rl, id) and some other event. That is, the out degree of (rl', id') is 2. Since we are only tracking the number of successors of (rl', id') and not the precise names, (rl, id) can get fired twice by *Other* process. This does not happen for the concrete agents because they have state variables controlling their execution which is not the case with *Other* process. Adding conflict constraints fixes this problem as well.

To complete the proof it was necessary to add 5 lemmas by hand. Since we had in- and out- message buffers, characterizing when a cache had access to a given item was hard. A grant shared message might be sitting in the in-message buffer, for instance, though the L2 state may not reflect it. We used the CMP method (without the flow invariants) earlier to verify a protocol of similar size though less intricate because of simpler internal structure of each cache [25]. There we had to add 25 lemmas by hand. Compared to that effort the reduction obtained using flows is dramatic and clearly makes the CMP method much more usable. This experience once again confirms that flows do yield powerful invariants that get to the heart of protocol correctness. The running time for the final abstract model was around 5.5 hours.

D. Flows and State Explosion

A surprising discovery during this proof process was that, even after we have chosen important flows involving the directory, using all the flows is not the right thing to do as it leads to state explosion in the abstract model. Apart from various flows for requesting shared and exclusive accesses we had a collection of flows for write backs and invalidates. Unlike the former flows the latter flows were not incompatible with themselves, that is, there could well be many of these flows alive at the same time. Thus, in the abstract model, the *Other* process can fire rules from these flows multiple times and saturate the auxiliary variables leading to an explosion in the number of states. It took us some time to understand this phenomenon, especially since augmenting the concrete model with auxiliary variables increased the state space by at most a factor of 2. That is, the auxiliary variables don't increase the number of states in the concrete model by much – they only widen the state. But this is not so for the abstract model, especially if we have flows that don't have too many conflict constraints. This experience indicates that flows that appear in their own conflict sets might be best ones to use.

²This had to be done only for some of the rules. Similarly, not all branch conditions had to be lifted.

E. Flows as Validation Collateral

Apart from helping us prove the safety properties of interest by yielding invariants, the flows themselves are valuable validation collateral. The CMP method not only uses lemmas but also validates them in the process. Thus, we are not only using flows but also proving them correct. To see the crucial flows of the protocol exercised and to have an assurance that important actions of the protocol, like the directory sending grant exclusive/shared access messages, happen precisely as specified by flows, constitutes a much stronger validation of the protocol than just verifying a final global property. In fact, the architects who saw our proofs were impressed more with the fact that we validated the flows than they were with the global safety properties verified!

VIII. CONCLUSION

Finding invariants is one of the central problems of formal verification and extensive research is being carried out to find invariants automatically or via user provided annotations. While the flow based technique falls into the latter category, it has the advantage that flows arise naturally and are readily available. Ideal annotations (or user supplied information) should be 1) easy to find, 2) provide relevant information precisely and in an easy to understand manner and 3) stay stable as the design details change. Flows have all three properties which makes them so attractive to use.

Apart from message passing distributed system handled here, flows or partial order on system events can also be applied to other types of distributed system. Lamport [17] has used partial orders analogous to flows in reasoning about mutual exclusion algorithms. In [17] the partial order is defined over *operations* which consist of series of atomic events. In addition to the precedence relation over events used in this paper, Lamport also defines a *can influence* relation over operations. While we used flows to derive invariants, Lamport (manually) reasons directly in terms of the partial order to prove the mutual exclusion property (which is also defined in terms of the precedence relation). A natural extension to our work is to generalize the notion of flows along the lines suggested by [17] and use it verify other types of distributed systems.

In summary, flows succinctly capture the core ideas that go into the design of a protocol and open up an promising avenue to pursue in our search for powerful system invariants. In conjunction with the CMP method, they lead to a technique that can handle the largest of protocols.

Acknowledgments We thank Ching-Tsun Chou for building and validating the original LCP Murphi model that we validated parametrically in this paper.

REFERENCES

- [1] P. Abdullah, A. Buojjani, B. Jonsson, and M. Nilsson. Handling Global Conditions in Parameterized Verification. In *Proc. CAV*, 1999.
- [2] P. Abdullah and B. Jonsson. On the Existence of Network Invariants for Verifying Parameterized Systems. In *Correct Systems Design-Recent Insights and Advances*, 1999.
- [3] J. Bingham. Automatic invariant generation for parameterized verification. In *Submitted to FMCAD*, 2008.
- [4] C.-T. Chou, P. K. Mannava, and S. Park. A simple method for parameterized verification of cache coherence protocols. In *Proc. FMCAD*, 2004.
- [5] E. Clarke, M. Talupur, T. Touilli, and H. Veith. Verification by Network Decomposition. In *Proc. CONCUR*, 2004.
- [6] E. Clarke, M. Talupur, and H. Veith. Environment Abstraction for Parameterized Verification. In *Proc. VMCAI*, 2006.
- [7] E. Clarke, M. Talupur, and H. Veith. Proving Ptolemy Right: Environment Abstraction Principle for Parameterized Verification. In *Proc. TACAS*, 2008.
- [8] A. Cohen and K. Namjoshi. Local Proofs for Global Safety Properties. In *Proc. CAV*, 2007.
- [9] S. Das, D. L. Dill, and S. Park. Experience with Predicate Abstraction. In *CAV*, 1999.
- [10] G. Delzanno. Automated Verification of Cache Coherence Protocols. In *Proc. CAV*, 2000.
- [11] A. E. Emerson and V. Kahlon. Reducing Model Checking of the Many to the Few. In *Proc. CADE*, pages 236–254, 2000.
- [12] A. E. Emerson and V. Kahlon. Model Checking Large-scale and Parameterized Resource Allocation Systems. In *Proc. TACAS*, pages 251–265, 2002.
- [13] E. A. Emerson and K. S. Namjoshi. Reasoning about Rings. In *Proc. POPL*, 1995.
- [14] S. Krstic. Parameterized system verification with guard strengthening and parameter abstraction. In *Automated Verification of Infinite State Systems*, 2005.
- [15] S. K. Lahiri and R. Bryant. Constructing Quantified Invariants. In *Proc. TACAS*, 2004.
- [16] S. K. Lahiri and R. Bryant. Indexed Predicate Discovery for Unbounded System Verification. In *Proc. CAV*, 2004.
- [17] L. Lamport. A new approach to proving correctness of multiprocess programs. *Transaction on Programming Languages and Systems*, 1(1):84–97, 1979.
- [18] Y. Lv, H. Liu, and H. Pan. Computing invariants for parameter abstraction. In *MEMOCODE*, 2007.
- [19] K. L. McMillan. Verification of infinite state systems by compositional model checking. In *CHARME*, 1999.
- [20] K. L. McMillan. Parameterized verification of the FLASH cache coherence protocol by compositional model checking. In *Proc. Conf. on Correct Hardware Design and Verification Methods (CHARME '01)*, volume 2144 of *LNCS*, pages 179–195. Springer, 2001.
- [21] S. Park and D. L. Dill. Verification of flash cache coherence protocol by aggregation of distributed transactions. In *SPAA '96: Proceedings of the eighth annual ACM symposium on Parallel algorithms and architectures*, pages 288–296. ACM Press, 1996.
- [22] A. Pnueli, J. Xu, and L. Zuck. Liveness with $(0, 1, \infty)$ Counter Abstraction. In *Proc. CAV*, 2002.
- [23] L. Seiler, D. Carmean, E. Sprangle, T. Forsyth, M. Abrash, P. Dubey, S. Junkins, A. Lake, J. Sugarman, R. Cavin, R. Espasa, E. Grochowski, T. Juan, and P. Hanrahan. Larrabee: A many-core x86 architecture for visual computing. *ACM Transactions on Graphics*, 27(3), Aug. 2008.
- [24] M. Talupur. *Abstraction Techniques for Infinite State Verification*. PhD thesis, SCS, CMU, 2006.
- [25] M. Talupur, S. Krstic, J. O'Leary, and M. R. Tuttle. Parameterized Verification of Industrial Strength Cache Coherence Protocols. In *Proc. Workshop on Design of Correct Circuits (DCC)*, 2008.
- [26] M. Talupur and M. R. Tuttle. Going with the Flow: Parameterized Verification using Message Flows. In *Proc. FMCAD*, 2008.