

3.36pt

Compositional verification of a lazy cache coherence protocol with a weak memory model

Chris Banks, Marco Elver, Vijay Nagarajan, Paul Jackson



PEPA Club, 4th November 2016

- How we verified a lazy cache coherence protocol with a weak memory model:

- How we verified a lazy cache coherence protocol with a weak memory model:
 - Intro to TSO-CC and lazy protocols;

- How we verified a lazy cache coherence protocol with a weak memory model:
 - Intro to TSO-CC and lazy protocols;
 - why previous verification techniques don't work;

- How we verified a lazy cache coherence protocol with a weak memory model:
 - Intro to TSO-CC and lazy protocols;
 - why previous verification techniques don't work;
 - stages of verification:
 - modelling TSO-CC;
 - checking TSO-CC satisfies the memory model (TSO)
 - parameterisation (compositional check for n cores).

- How we verified a lazy cache coherence protocol with a weak memory model:
 - Intro to TSO-CC and lazy protocols;
 - why previous verification techniques don't work;
 - stages of verification:
 - modelling TSO-CC;
 - checking TSO-CC satisfies the memory model (TSO)
 - parameterisation (compositional check for n cores).
- Future directions?

TSO-CC and lazy protocols

The basics (all we need to know for this talk):

TSO-CC and lazy protocols

The basics (all we need to know for this talk):

- TSO-CC is a lazy cache coherence protocol,

TSO-CC and lazy protocols

The basics (all we need to know for this talk):

- TSO-CC is a lazy cache coherence protocol,
 - designed to take advantage of weak memory model
 - scale to large numbers of cores.

TSO-CC and lazy protocols

The basics (all we need to know for this talk):

- TSO-CC is a lazy cache coherence protocol,
 - designed to take advantage of weak memory model
 - scale to large numbers of cores.
- Weak memory models (TSO) only require consistency at synchronisation boundaries.

TSO-CC and lazy protocols

The basics (all we need to know for this talk):

- TSO-CC is a lazy cache coherence protocol,
 - designed to take advantage of weak memory model
 - scale to large numbers of cores.
- Weak memory models (TSO) only require consistency at synchronisation boundaries.
- TSO-CC is lazy:
 - does not eagerly invalidate on writes,

TSO-CC and lazy protocols

The basics (all we need to know for this talk):

- TSO-CC is a lazy cache coherence protocol,
 - designed to take advantage of weak memory model
 - scale to large numbers of cores.
- Weak memory models (TSO) only require consistency at synchronisation boundaries.
- TSO-CC is lazy:
 - does not eagerly invalidate on writes,
 - cache lines self-invalidate after so many reads (eventually consistent),

TSO-CC and lazy protocols

The basics (all we need to know for this talk):

- TSO-CC is a lazy cache coherence protocol,
 - designed to take advantage of weak memory model
 - scale to large numbers of cores.
- Weak memory models (TSO) only require consistency at synchronisation boundaries.
- TSO-CC is lazy:
 - does not eagerly invalidate on writes,
 - cache lines self-invalidate after so many reads (eventually consistent),
 - does not need to track sharers, so is more scalable (storage reqs.).

Classical verification fails

- Classical verification techniques rely on local properties:
 - e.g. Single Writer Multiple Reader (SWMR),
 - Data Value Invariant, ...

Classical verification fails

- Classical verification techniques rely on local properties:
 - e.g. Single Writer Multiple Reader (SWMR),
 - Data Value Invariant, ...
- These don't hold in lazy protocols.

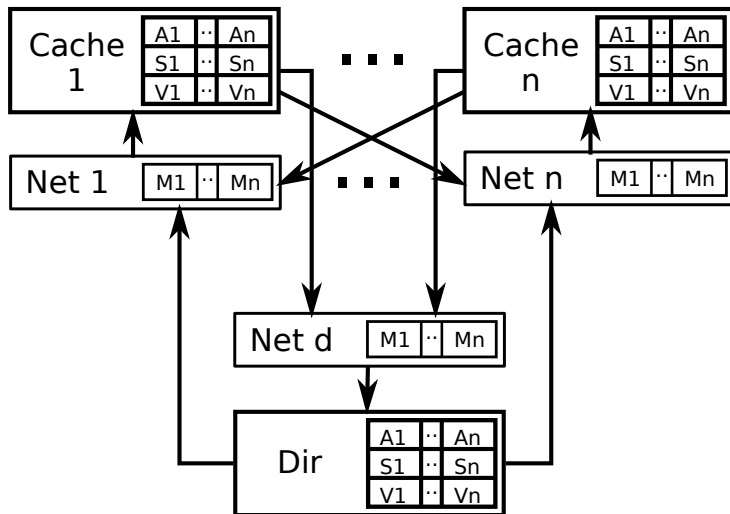
Classical verification fails

- Classical verification techniques rely on local properties:
 - e.g. Single Writer Multiple Reader (SWMR),
 - Data Value Invariant, ...
- These don't hold in lazy protocols.
- What do we verify against then?

Classical verification fails

- Classical verification techniques rely on local properties:
 - e.g. Single Writer Multiple Reader (SWMR),
 - Data Value Invariant, ...
- These don't hold in lazy protocols.
- What do we verify against then?
- Answer: the memory model.

Modelling TSO-CC



Mur ϕ example rules:

$$\begin{array}{l} c[a].state = I \implies \text{SendGetS}(c, Dir, a); \\ c[a].state := WS \qquad \qquad \qquad \text{(Read I)} \end{array}$$
$$\begin{array}{l} c[a].state = E \implies c[a].val := v \\ c[a].state := M; \qquad \qquad \qquad \text{(Write E)} \end{array}$$

Rules to handle processor actions.

Rules to handle receipt of messages, functions for actions upon receipt:

```
DirectoryReceive(msg, a) = if Dir[a].state = I  $\wedge$  msg.type = GetS
    then SendDataS(msg.src, a, ...);
        ReplaceOwner(msg.src, a);
        Dir[a].state := WE1
    else ...
```

We do some clever things like abstracting counters:

$$c[a].state = S \implies \text{SendGetS}(c, Dir, a);$$
$$c[a].state := WS \quad (\text{Read } S \text{ (ACount=MAX)})$$
$$c[a].state = S \implies //\text{do nothing} \quad (\text{Read } S \text{ (ACount < MAX)})$$

(This becomes important later.)

- How do we show TSO-CC \models TSO?

TSO-CC \models TSO

- How do we show TSO-CC \models TSO?
- Answer: show TSO-CC is *weakly simulated by* TSO.

TSO-CC \models TSO

- How do we show TSO-CC \models TSO?
- Answer: show TSO-CC is *weakly simulated by* TSO.
- To show a weak simulation relation we need an operational model of TSO.

- How do we show TSO-CC \models TSO?
- Answer: show TSO-CC is *weakly simulated by* TSO.
- To show a weak simulation relation we need an operational model of TSO.
- Existing operational models of TSO have potentially infinite store buffers.

- How do we show TSO-CC \models TSO?
- Answer: show TSO-CC is *weakly simulated by* TSO.
- To show a weak simulation relation we need an operational model of TSO.
- Existing operational models of TSO have potentially infinite store buffers.
- TSO-LB has finite load buffers.

- How do we show TSO-CC \models TSO?
- Answer: show TSO-CC is *weakly simulated by* TSO.
- To show a weak simulation relation we need an operational model of TSO.
- Existing operational models of TSO have potentially infinite store buffers.
- TSO-LB has finite load buffers.
- TSO-LB \implies TSO (proof).

TSO-CC \models TSO-LB

- $\text{local} : (S \times P \times A) \rightarrow V$ is a function where $\text{local}(s, p, a)$ is the value at address a in the local buffer of p in state s .
- $\text{global} : (S \times A) \rightarrow V$ is a function where $\text{global}(s, a)$ is the value at address a in the global buffer in state s .

$$\frac{\text{local}(s, p, a) = v}{s \xrightarrow{\text{Read}(p, a, v)} s} \text{ READ}$$

$$\frac{\top}{s \xrightarrow{\text{Write}(p, a, v)} s' \wedge \text{local}(s', p, a) \mapsto v \wedge \text{global}(s', a) \mapsto v} \text{ WRITE}$$

$$\frac{\top}{s \xrightarrow{\text{Prop}(p)} s' \wedge \forall a. \text{local}(s', p, a) \mapsto \text{global}(s, a)} \text{ PROPAGATE}$$

Weak simulation checking

Add invocation of TSO-LB from TSO-CC (in the model checker).

$$c[a].\text{state} = E \implies c[a].\text{val} := v$$
$$c[a].\text{state} := M;$$

TSOStore(c, a, v)

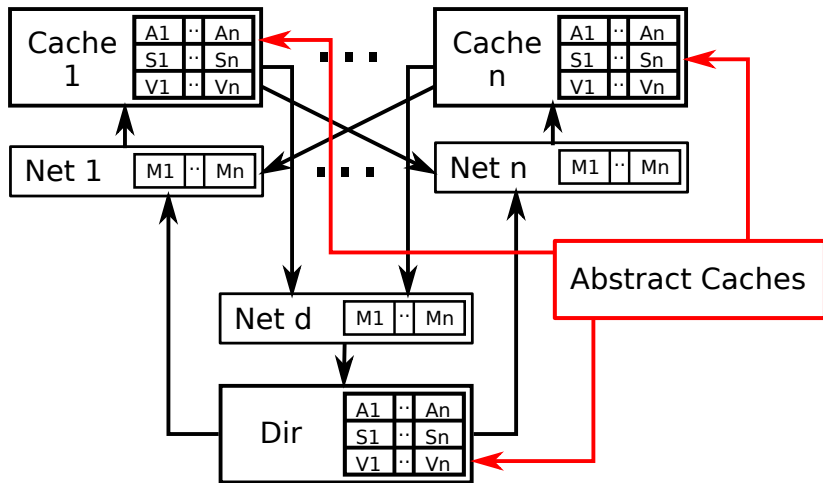
 (Write E)
$$c[a].\text{state} = S \implies //\text{do nothing};$$

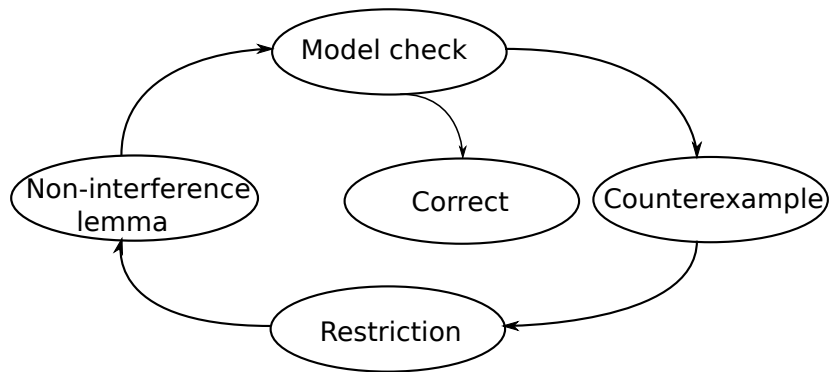
Assert(TSOVerify($c, a, c[a].\text{val}$))

 (Read S)

Thus we show weak simulation by model checking.

Parameterisation





Add rules for messages received from the abstract cache, e.g.:

$c[a].\text{state} = WX$

$\implies \text{SendAck}(c, Dir, a);$

$c[a].\text{state} := M;$

$\text{InvalidateAllOtherLines}(c, a)$

(Cache Recv DataX Abs)

+ all combinations of node state and message type.

Abstract cache initially produces all possible messages
(*over-approximation*).

Parameterisation

Model check; get counterexample; add restriction:

$$c[a].state = WX \quad \boxed{\wedge \text{ IsOwner}(c, a)}$$
$$\implies \text{SendAck}(c, Dir, a);$$
$$c[a].state := M;$$
$$\text{InvalidateAllOtherLines}(c, a)$$

(Cache Recv DataX Abs)

Add non-interference lemma:

$$\forall n \forall a \forall i. \text{net}[n][a][i].msgType = \text{DataX} \implies \text{IsOwner}(n, a)$$

Repeat. (31 restriction/lemma pairs!)

Problems solved:

- How to verify against TSO – *weak simulation*.
- How to show weak simulation in the model checker.
- Finite state operational model of TSO (TSO-LB).
- How to extend this to compositional model checking.

Validation:

- bug found: only in the full n -process model!
- Trivial bug in model, but shows it catches more bugs.