

# Verifying Multiple Compare and Swap


Joe Hurd

`joe.hurd@cl.cam.ac.uk`

University of Cambridge

# Introduction

This work shows how a parallel architecture can be formalized in HOL, with many processors acting on a shared memory.



Relevant to possible future work of the ARM project (individual processors could be retargeted to ARM instruction set).

# The Formalization Task

Create a logical model of a parallel architecture, supporting:

- independent execution of many processors communicating only through a shared memory; ✓
- arbitrary interleaving of instructions at the granularity of the primitive atomic operations; ✓
- plugging-in different memory models, with different rules for access reordering and barrier instructions; ✓
- automatic tools and techniques for specification and verification of concurrent programs.  $\rightsquigarrow$  **current work**

# Parallel Architecture: Overview

- Identify a processor with its register file:

$$\textit{processor} = \textit{register} \rightarrow \textit{value}$$

- The whole parallel architecture can then be modelled by a list of processors and a shared memory:

$$\textit{architecture} = \textit{memory} \times \textit{processor list}$$

- The global state advances by any non-halted processor executing a primitive atomic instruction.

# Parallel Architecture: Instruction Set

An *instruction* corresponds to a primitive atomic operation:

*instruction* =

NOP

| UP of *value*  $\times$  *register*

| RD of *register*  $\times$  *register*

| WR of *register*  $\times$  *register*

| CAS of *register*  $\times$  *register*  $\times$  *register*  $\times$  *register*

| ALLOC of  $\mathbb{N}$   $\times$  *register*

| MB

Some memory models (such as the Alpha) use memory barrier (MB) instructions to impose an ordering on accesses.

# The CAS(n) Algorithm

- A case study to investigate verification in our model.
- CAS(n) stands for Multiple Compare And Swap:

if the  $n$  memory addresses  $a_1, \dots, a_n$   
contain the expected values  $x_1, \dots, x_n$   
then replace them with the values  $y_1, \dots, y_n$   
(and this must all happen atomically!)

- A fast CAS(n) algorithm was recently developed by Harris, Fraser and Pratt in Cambridge.
- So far has led to:
  - instruction macros to simplify programming;
  - and a simulation engine to test arbitrary interleaving.

# Conclusion

- We have laid out a formalization of a parallel architecture in HOL.
- We're working on case studies to develop techniques for specification and verification for a simple instantiation.
- There's scope for making the memory model more realistic (c.f. Gordon's model of the Alpha architecture),
- and also the processor instruction set (c.f. Fox's model of the ARM).
- The intention is that proofs completed for the simple instantiation lift up to more realistic models (we'll see).