Parametric Verification of Industrial Cache Protocols (working title)

Distributed protocols like cache coherence protocols form the bedrock on which modern multiprocessor systems are built. Such distributed protocols are typically designed *parametrically*, that is, independent of the precise number of processors involved. Given that distributed programs are hard to reason about for humans and that no amount of testing/simulation can cover all scenarios, it becomes necessary that we find methods to formally and parametrically verify the correctness of such systems.

In this talk we will relate our practical experience with parameterized verification of an on-die cache coherence protocol for a many-core microprocessor design in progress at Intel. The protocol contains complexity that is not present in standard examples such as the FLASH or German protocols. To give an idea: German protocol (the standard academic version [4]) has 7 different messages; the FLASH protocol, which is considered very hard and only 2 or 3 methods have ever been successfully applied to it, has 16 different messages. In contrast, our protocol with 54 different types of messages is vastly more complex.

The verification technique we used is based on a method first described by McMillan [3] and subsequently elaborated by Chou, Mannava and Park [1] and Krstić [2]. This method, which we call the CMP method, is based on circular compositional reasoning and uses model checkers as proof assistants. Briefly, a parameterized system containing a directory and N cacheing agents is abstracted to a system containing a directory, two cacheing agents, and a third, highly nondeterministic, process representing “all the other agents”. The user then supplies a series of lemmas that refine the behavior of the third agent; these lemmas are used mutually to prove one another and also the final property of interest. Coming up with these lemmas is a time consuming process requiring a deep understanding of the protocol. It took us about a month and 25-odd lemmas to prove the safety property of the protocol. As far as we are aware this is the first time a protocol of this size and complexity has been verified parametrically.

The next step of the project was to make the CMP method easy to use by automating as much of it as possible. The CMP method is composed of three main stages: (i) creating the inital abstraction, (ii) coming up with lemmas after examining counter example traces, and (iii) refining the abstract model. While writing the lemmas requires ingenuity the other two parts can be automated. We have built a tool that creates an initial abstraction and refines the abstract model with user supplied lemmas automatically, and our talk will also describe the principles behind this tool.

The most important limitation of the CMP method as presented by Chou et al and earlier by McMillan is that it does not deal with systems in which processes are ordered by their indices. Commonly occuring algorithms break symmetry between processes by ordering them according to indices either linearly (as in the bakery algorithm), or placing them on a ring, grid or other network topologies. The third element of our talk will explain how we extended the theory to handle such asymmetric systems as well and incorporated these extensions in our tool. Extending the theory to handle asymmetric systems led us to discover what we believe are new circular compositional reasoning principles. Besides extending the CMP method to asymmetric systems, these new principles also allow to formulate and reason about the CMP method much more intuitively and succinctly.

References

[1] C.-T. Chou, P. K. Mannava, and S. Park. A simple method for parameterized verification of cache coherence protocols. In A. J. Hu and A. K. Martin, editors, *Proc. Conf. on Formal Methods in Computer-Aided Design (FMCAD04)*, volume 3312 of *Lecture Notes in Computer Science*, pages 382–398. Springer, 2004.

[2] S. Krstić. Parameterized system verification with guard strengthening and parameter abstraction. In *Fourth Int. Workshop on Automatic Verification of Finite State Systems*, 2005. To appear in *Electronic Notes in Theoretical Computer Science*.

[3] K. L. McMillan. Parameterized verification of the FLASH cache coherence protocol by compositional model checking. In *Correct Hardware Design and Verification Methods (CHARME'01)*, volume 2144 of *LNCS*, pages 179–195, 2001.

[4] A. Pnueli, J. Xu, and L. D. Zuck. Liveness with (0, 1, infty)-counter abstraction. In *CAV*, pages 107–122, 2002.