Multi-Core Model Checking for Biological Applications

Jaco van de Pol

Centre for Telematics and Information Technology, Universiteit Twente,
P.O. Box 217, 7500 AE Enschede, The Netherlands
vdpol@cs.utwente.nl

Multi-core model checking algorithms aim at speeding up verification tasks, by using multiple processor cores running in shared memory. Using shared memory avoids communication overhead due to message passing, but it is far from trivial to obtain ideal speedups, since the underlying graph algorithms are memory intensive and irregular.

I will present some key ideas in multi-core model checking, which have enabled us to provide scalable solutions for reachability, LTL model checking, and symbolic model checking. Key ingredients are a scalable shared hashtable, parallel random search algorithms, and an efficient work stealing scheme to implement multi-core decision diagrams.

We implemented our ideas in the LTSmin toolset, providing a high performance model checker through the PINS interface, independent of the specification language. We instantiated this for Promela, mCRL, DVE and UPPAAL timed automata. I will explain how a slight extension of PINS enables LTL model checking for timed automata.