

MiraXT – A Multithreaded SAT Solver

Tobias Schubert Matthew Lewis Natalia Kalinnik Bernd Becker
Institute of Computer Science, Albert-Ludwigs-University of Freiburg, Germany
{schubert, lewis, kalinnik, becker}@informatik.uni-freiburg.de

Abstract

This paper briefly describes MiraXT, a multithreaded SAT solver that was designed to take advantage of current and future shared memory multiprocessor systems. The experimental results in [6] show that already in single threaded mode, MiraXT compares well to other state-of-the-art solvers on a wide range of industrial problems. In threaded mode, it provides cutting edge performance, as speedup is obtained on both satisfiable and unsatisfiable instances. The paper highlights some of the design and implementation details that allow multiple threads to run and cooperate efficiently.

1. MiraXT

1.1. Overview

MiraXT is a significantly enhanced re-implementation of MIRA [4, 5] that is able to run with multiple threads. It contains MIRA's *Early Conflict Detection BCP* as well as *Implication Queue Sorting*. As the decision strategy, a modified VSIDS algorithm is used, in which all scores over 512 are concatenated so that a bucket sort can be used to sort the list in linear time with respect to the number of variables. This allows us to sort the list more frequently keeping it up-to-date, and makes the decision heuristic less greedy. Lastly, MiraXT was implemented in C++ using POSIX threads.

1.2. Shared Clause Database

MiraXT uses one master clause database that stores pointers to the original problem clauses, plus pointers to all the conflict clauses generated by each thread. Each clause is only present once in memory, and is shared between all threads. In order to insure coherency within the database, a lock must be acquired before a thread inserts a pointer to its newly generated conflict clause. As soon as the pointer is inserted and the database clause counter is incremented the lock is released. All clauses, once generated, are read-only, so that sharing can be done without locks. These steps are important as we want to reduce the amount of locks needed by the solver, and remove any lock contention and

wait times that might result from the remaining locks. Also, each thread has one lock associated with it that is used when the thread requests a new clause from the master clause database. This lock is used to increment its current database position pointer. This pointer keeps track of which clauses the thread has already looked at, and those that can still be added. Figure 1 shows a top level diagram with threads inserting pointers to clauses into the master clause database (C_x represents a pointer to a conflict clause).

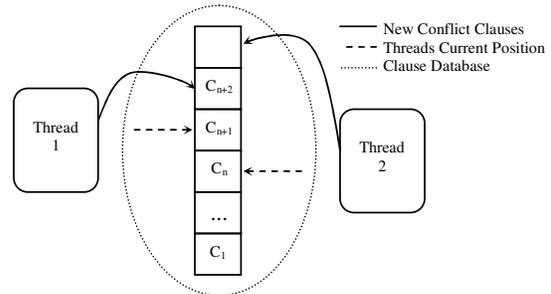


Figure 1. Shared clause database structure

Clause deletion is also an important issue. In MiraXT, each thread deletes clauses using an algorithm similar to Berkmin [2] in which older inactive clauses are easier to delete. To facilitate clause deletion efficiently on a multiprocessor system, each thread has one Boolean variable associated with it for every clause. Also, each clause consists of an array of literals with the first few spots in the array being reserved. These reserved spots specify the clause length and its unique master database reference number. When a thread deletes its references to a clause, it must set its Boolean variable for that clause using the clauses reference number. Because the Boolean variable for the clause is specific for that thread, no global lock is required when deleting clauses.

Once a thread has deleted all the clauses it wants to delete, it will ask the master database to see if a master delete should be run, as the threads only delete their references to clauses, and not the actual clauses themselves. In MiraXT, a simple test based on how many threads there are, and how many deletion processes have been run, is used to decide if a master delete is required. If the master

database needs cleaning, the thread grabs a lock and proceeds to delete clauses that are no longer used by any thread, relinquishing the lock when it is finished. This lock is used to insure that no two threads run a master clause deletion procedure at the same time.

Using the fine grained lock system described above, practically all lock contention issues were removed, and in testing we saw no signs of even light lock contention.

1.3. Watched Literal Reference List

In most solvers, to keep track of the watched literals, the original clause is modified in some way (e.g., by using the first two literals in the clause). This is not possible in MiraXT, because clauses are read-only. So, each thread creates a second data structure called the *Watched Literal Reference List* (WLRL). For each clause, this structure contains two watched literals and a third literal, that is checked first by the BCP procedure when searching for a new watched literal. The WLRL basically allows each thread in MiraXT to have a condensed reference or copy of every clause. In early experiments on a selection of BMC problems we observed, that 84% of the clauses with 3 literals or more can be directly evaluated with only the WLRL. This means the original clauses are not needed 84% of the time! Also, on many problems, clauses with 3 literals or less are fairly common, meaning the entire clause can be stored here. In any case, this allows MiraXT to better utilize each CPU's cache and local memory.

1.4. Preprocessing

Currently, two forms of preprocessing are integrated in MiraXT: first, all the simplification routines from MiniSat [1], which are executed only once during the initialization phase. And second, a Boolean *Unit Propagation Lookahead* mechanism similar to [3], which is executed not only during the initialization but also during the SAT solving phase when a thread receives a new subproblem, or when the solver has assigned a large number of variables to decision level 0. In these situations, the unit propagation lookahead procedure will only look at variables which could be directly affected by the new variable assignments such as free variables that are in unsolved clauses which contain decision level 0 variables.

1.5. Multithreaded Solver Control

MiraXT contains no controlling master process unlike most other parallel solvers. Instead, a master control object (MCO) allows the threads to communicate with each other. All communication is done in a passive way, such that the MCO will not interfere with the threads. It will only store messages and suspend threads which ask for it to do so. The solver threads poll the MCO occasionally to check for messages, such as idle threads waiting for a new subproblem.

1.6. Multithreaded Conflict-Driven Learning

The conflict analysis procedure in MiraXT is based on zChaff's *IUIP* scheme [7]. However, a separate clause addition procedure was added. In MiraXT, the conflict analysis procedure will add a clause pointer to the master clause database. Then the clause addition procedure will be run, asking the master clause database for all new clauses; this includes clauses generated by other threads and its newly generated conflict clause. It will then process these clauses, deciding which clauses should be added. Currently, all conflict clauses, undefined clauses, or really short clauses (10 literals or less), are added. The clause addition procedure will assign watched literals, search for implications, and perform conflict driven backtracking as needed. Both the conflict analysis procedure and the clause addition procedure can signal that the current subproblem is unsatisfiable. Note, that sometimes the thread might decide not to add the conflict clause it just generated because clauses generated by other threads were better, allowing the solver to backtrack further.

2. SAT Competition 2007

Three different versions of MiraXT are going to enter this year's SAT competition. *MiraXTv1* and *MiraXTv2* are configured as sequential SAT solvers, but slightly differ with respect to the parameters of the VSIDS decision strategy, while *MiraXTv3* always starts two threads (even if both of them have to be executed on the same CPU).

Finally, the source code of MiraXT can be obtained at <http://ira.informatik.uni-freiburg.de/~schubert/>.

References

- [1] N. Eén and A. Biere. Effective Preprocessing in SAT through Variable and Clause Elimination. In *8th International Conference on Theory and Applications of Satisfiability Testing*, 2005.
- [2] E. Goldberg and Y. Novikov. BerkMin: A Fast and Robust SAT-Solver. In *Design, Automation, and Test in Europe*, 2002.
- [3] D. Le Berre. Exploiting the Real Power of Unit Propagation Lookahead. In *LICS 2001 Workshop on Theory and Applications of Satisfiability Testing*, 2001.
- [4] M. Lewis, T. Schubert, and B. Becker. Early Conflict Detection Based BCP for SAT Solving. In *7th International Conference on Theory and Applications of Satisfiability Testing*, 2004.
- [5] M. Lewis, T. Schubert, and B. Becker. Speedup Techniques Utilized in Modern SAT Solvers – An Analysis in the MIRA Environment. In *8th International Conference on Theory and Applications of Satisfiability Testing*, 2005.
- [6] M. Lewis, T. Schubert, and B. Becker. Multithreaded SAT Solving. In *12th Asia and South Pacific Design Automation Conference*, 2007.
- [7] M. Moskewicz, C. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an Efficient SAT Solver. In *IEEE/ACM Design Automation Conference*, 2001.