

Solving “Hard” Satisfiability Problems Using GridSAT *

Wahid Chrabakh

University of California Santa Barbara
Department of Computer Science
Santa Barbara, CA
chrabakh@cs.ucsb.edu

Rich Wolski

University of California Santa Barbara
Department of Computer Science
Santa Barbara, CA
rich@cs.ucsb.edu

Abstract

We present the latest instantiation of GridSAT [5], a distributed and complete satisfiability solver that is explicitly designed to aggregate grid resources for application performance. GridSAT was previously shown to outperform the state-of-the-art sequential solvers. In this work, we explore the unprecedented solving power GridSAT enables through algorithmic and implementation innovations. We describe the implementation techniques that allow GridSAT to leverage a variety of high-end batch-scheduled resources, clusters, interactive workstations, and personal computing resources through autonomous scheduling, checkpoint scheduling, and work migration. These innovations have allowed GridSAT to solve a set of “hard” and previously unsolved industrial and community satisfiability problems. In addition to this new solution power, GridSAT also outperforms the otherwise highest performance general solvers on the annual SAT competition performance benchmarks.

Keywords: *Parallel, Distributed, Scheduling, Satisfiability, Computational Grid.*

1 Introduction

Grid computing [11] is an emergent field in computer science that focuses, in part, on the aggregation of geographically distributed and federated computational resources. These resource aggregations can be harnessed by grid applications to solve problems in science and engineering [21, 1] which require large computing power. Solving such challenging problems and enabling new scientific results is an integral part of the grid computing vision.

One such challenging problem is propositional satisfiability. This problem involves finding a set of binary assignments to variables that satisfies a set of constraints (i.e. makes a binary expression evaluate to “true”). The problem of solving satisfiability instances is important from both theoretical and practical perspectives and is, in general, NP-complete. In practice, many engineering disciplines require the

*This work was supported by grants from the National Science Foundation, numbered CAREER-0093166, EIA-9975020, ACI-0103759, and CCR-0331654.

solution to domain specific instances of satisfiability. Such disciplines include scheduling, model checking, security, Artificial Intelligence, software verification, and the the area of Electronic Design Automation (EDA) which includes circuit design [29], Field-Programmable Gate Arrays (FPGA) detailed routing [23], combinational equivalence checking [18] and automatic test and pattern generation [20].

Because satisfiability solvers [22, 12, 15, 3] have become more efficient, they are now widely used in many industrial and research settings. There has been an extensive research effort geared towards the development of gradually more efficient satisfiability solvers [22, 12, 15, 3]. These solvers use different techniques to navigate the entire search space of possible truth assignments for the variables of a given expression. The best (fastest and most comprehensive) of these solvers use *learning* optimizations that permit the search space to be “pruned” during execution. Learning [28] introduces new deduced propositions which improve the solver’s efficiency by obviating subtrees in the space of possible variable assignments.

Because learning requires a large, centralized database of intermediate propositions to be searched and updated frequently, the best known solvers are sequential. These sequential solvers are characterized by heavy use of compute power (CPU) as well as the memory of the host machine as the database must be kept memory-resident (or the speed becomes unacceptably low).

Research in parallel solvers [5, 17, 30, 8] , shows that using a large pool of computational resources leads to better performance for most problems. The aggregate CPU power and memory of the hosts allows the solver to navigate the search space faster. Thus a computational grid populated by a large pool of resources offers potential improvements in solver speed. With the exception of those results reported in [5], however, the fastest solutions to the largest number of problems is generated by sequential solvers [26, 25].

By carefully leveraging the resources in grid settings, our goal is to build a parallel and distributed satisfiability (SAT) solver that correctly solves previously infeasible industrial problem instances, the answers for which cannot be determined in any other way. Secondly, we would like to be able to solve faster the problems that sequential solvers find feasible.

Our previous work with GridSAT [5, 4] demonstrates the latter. By dynamically acquiring and releasing resources under the control of an automatic scheduler, GridSAT improves the time-to-solution for various feasible SAT instances. Indeed, GridSAT outperforms the best-known solver on all problems that this leading solver can complete [26, 25]. We have also been able to use GridSAT to solve several previously unsolved problems using non-dedicated, wide-area grid resources. It is these new domain-science results, and the techniques we have employed to achieve them, that are the subject of this paper.

In particular, by combining different batch-controlled super-computers with interactive workstations and user desktop machines, we have applied GridSAT to *hard* SAT problems – ones that are not only unsolved but for which previous attempts at solution using other general techniques have failed. This pattern of combining different types of resources is new and different from that used by existing *parallel* SAT implementations [17, 30]. Moreover, we know of no *distributed* (i.e. network and/or grid enabled) SAT implementations, efficient or otherwise, at the time of this writing.

The resources in a computational grid may be of two different types: time-shared or batch controlled. In the case of time-shared resources the application will compete with other user applications running simultaneously on the host machine. However, since these resources are always available the application can continue to make progress. Other resources which are controlled by a batch scheduler, will participate intermittently in the application through some of their nodes. But these systems will provide

significant compute power depending on the size of the application's request.

In order to enable a grid implementation of a SAT solver to use many resources simultaneously, we need to address two types of challenges. First the solver's algorithm needs to be modified so that it can run in parallel while ensuring that the parallel components cooperate to improve over-all efficiency. The second challenge is developing a framework capable of running the parallel solver in a very volatile computational environment.

Solving the above two problems was at the core of our methodology in designing the application components and their interactions. Implementing this methodology can be achieved by selecting suitable technologies. Examples of these technologies include those from parallel computing, which predate grid computing, such as MPI [9]. The more relevant technologies are those which were the outcome of grid-specific research projects such as Globus [10], Web Services [33] and related standards. We discuss in this paper the requirements imposed by the application's dynamic behavior and constraints on the technology so that a successful implementation is realized. We also describe the current design and implementation of the application.

We have developed GridSAT, a distributed satisfiability solver capable of running on a computational grid. GridSAT implements a parallel algorithm for solving satisfiability problems based on Chaff [22]. GridSAT distributes and shares the internal proposition database among processors in a way that takes advantage of dynamic resource performance predictions to achieve new levels of solver efficiency.

In this paper, we detail the current, most capable version of GridSAT. Our most recent improvements in the clause sharing and resource scheduling algorithms have made it possible to solve previously unsolved satisfiability problems from the field of FPGA routing as well as artificially generated benchmarks specifically design to foil automatic SAT solvers.

2 GridSAT: SAT Solver for the Grid

A satisfiability problem is expressed as a boolean formula over a set of variables. Most solvers operate on formulas expressed in Conjunctive Normal Form (CNF) in which an expression conjoins (logically "ANDs") a set of *clauses*, each of which may contain disjoined ("ORed") literals. A literal is either an instance of a variable (V) or its complement ($\sim V$) and variables are boolean. A SAT problem instance is termed *satisfiable* if there exists a set of variable assignments that makes the formula evaluate to *true* where "true" corresponds to a boolean 1 algebraically. If such an assignment does not exist the the problem is declared *unsatisfiable*.

GridSAT is based on Chaff [22], a sequential SAT solver algorithm. Chaff, in turn, builds upon the Davis-Putnam-Loveland-Logemann (DPLL) [7] algorithm which solves a SAT instance by making a set of speculative variable assignments (termed "decisions") stored in a *decision stack*. When these decisions are propagated through the clauses they could lead to a cascade of *implications*. Implications are assignments of boolean values to different variables as deductive consequences of previous speculative decisions. These speculative decisions and the resulting implications may lead to logical conflicts – deduced contradictions in which a variable must take on both boolean values because of different clauses in the original problem. In Chaff, as well as other solvers, the performance of the algorithm is enhanced by using techniques for adding new deduced clauses after a conflict occurs. This technique is called *Learning* [27, 19, 28]. Using learning, the algorithm may generate a vast number of additional clauses during execution. These clauses consume memory, possibly overwhelming the capacity of the host, and also may slow the algorithm as they can add to the search complexity of the clause database.

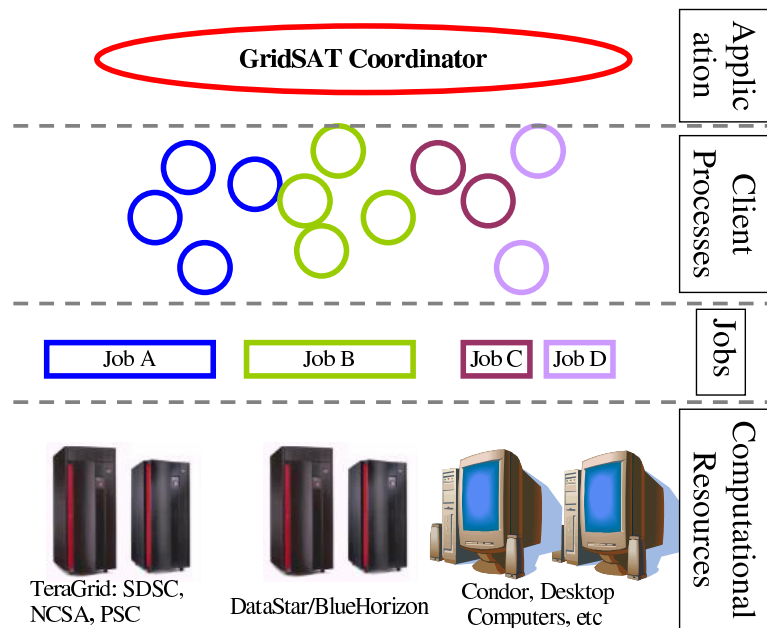


Figure 1. GridSAT resource views

GridSAT’s distributed solver addresses three significant challenges to improving solver performance. First, GridSAT parallelizes the search algorithm that is navigating the space of possible truth assignments. Second, certain learned clauses from the various solvers are selected to be distributed and shared across resources. Finally, the GridSAT application components are dynamically scheduled at runtime to take advantage of those available resources which can enhance the solver’s performance.

To apply a parallel search technique to SAT, we split the original problem into subproblems (having decision stacks with different truth assignments), each of which is independently investigated for satisfiability. Subproblems, themselves, may be split in the same way, forming a recursive tree, each node of which is assigned to a logically distinct processor. Clause sharing is facilitated by identifying and sharing only important clauses.

3 GridSAT Architecture and Resource Scheduling

GridSAT is implemented as a special form of the coordinator/client model where individual clients communicate directly and share clauses (i.e. communication is between peers rather than routed through the master). The GridSAT application uses two views of the computational resources as shown in figure 1. The first view employs jobs to classify processes which belong to the same resource. The second view is flat where all processes are part of a single pool. Both of these views are useful for managing resources under GridSAT

The coordinator (or master), shown in figure 2, reflects the resource views shown in figure 1. It consists of the resource manager, the client manager, the scheduler and the checkpoint server. We now describe the role of these components.

The resource manager is tasked with loading resource information from one or more grid information systems such as Globus MDS [6] and the NWS [37]. The scheduler, however, is responsible for coordinating the interactions between all the components. In addition, it handles interactions with ex-

ternal resources and monitors them to detect failures. For example, the scheduler queries the resource manager for resource types. If the resource is time-shared, then only one GridSAT process is launched. For batch systems, the scheduler instead submits one job request. Additional jobs could be manually submitted and GridSAT will use their resources when they become available. We term this form of scheduling *active queuing*; jobs waiting in queue logically execute on the interactive resources until the batch-controlled resources become available. At that time, the scheduler migrates work into the newly available resources. Thus, the application makes progress using the slower, shared resources while it waits in queue. It is the client manager that maintains a list of all GridSAT processes (active and queued) and monitors their progress.

The GridSAT scheduler is the focal point and is responsible for coordinating the rest of the components and launching new processes, also termed clients. The scheduler uses a progressive scheme for starting additional clients on remote resources and adding them to the active resources' pool. Resources which are no longer performing a task on behalf of GridSAT are released immediately when possible. The reason for this approach is the variability and unpredictability of resource usage for a particular SAT problem. Some problems are solved easily using a single host after a short time period. Other problems, however, might be harder and require a large number of hosts and a longer time period. By starting with a small resource pool and expanding the set of used resources, GridSAT achieves three goals. First, a small number of resources will be used to solve the easy problems which results in a smaller communication overhead and therefore shorter time to solve the problem. Second, GridSAT can adapt resource usage to how difficult the problem is perceived. If at a particular stage the problem is perceived difficult, the size of the resource pool used will grow. At another stage, the same problem might be perceived to be easy, a smaller resource set will be used, and excess resources will be released. Lastly, by remaining as small as possible at any given point in the execution, GridSAT promotes allocation stability and sharing. The scheduler does not waste resources needlessly thus the maximum number of GridSAT instances can co-exist since each is attempting to use as few resources as possible for its own problem instance.

The GridSAT scheduler uses the first available client immediately to start solving the problem. Each client records the time it took to receive the problem data. Clients also monitor their memory usage. The decision for splitting a problem is made locally by the client and not by a centralized scheduler. A client notifies the master that it wants to split its assigned subproblem with another client when its memory usage exceeds a specified limit (currently 80% of available memory) or after running for a specific period of time. This time period is determined as twice the duration of the communication period the client used to obtain the problem data. Using this method, the scheduler allows for computation time to offset the communication overhead by using the previous communication period as a prediction of

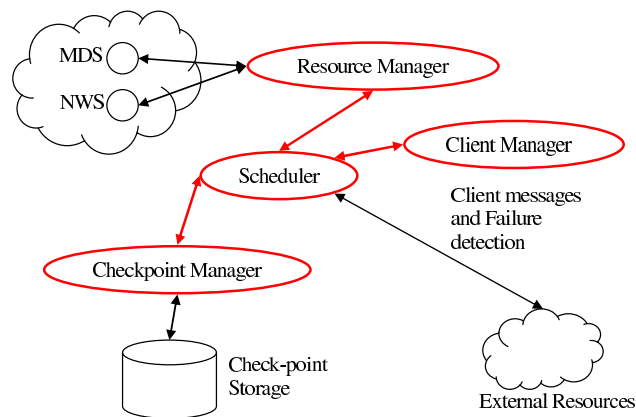


Figure 2. GridSAT components and their internal and external interactions. The external components and systems which GridSAT uses, such as the Globus MDS and the NWS, are shown in clouds.

future overhead. The clients, therefore, do not spend most of their time splitting instead of doing useful computation. The splitting process is performed by the cooperation of the master, the splitting client and an *idle* client. The *idle* client is a process which is not currently assigned a subproblem to investigate.

The GridSAT solver terminates when all subproblems are solved or one of the clients finds a satisfying assignment. In the latter case the client which finds the satisfying assignment sends its stack to the master. Finally, the master saves the final solution, terminates all running clients and cancels any pending resource requests. Most solvers in the literature are evaluated based on the time the first satisfiable instance is found. However there are cases [16] where knowing all satisfiable instances is helpful. GridSAT can also enumerate all the instances where a problem is satisfiable.

3.1 Active Queuing: Efficient Use of Batch Jobs

In GridSAT, initial batch job requests are large with a high number of nodes and long duration. This leads to a long waiting period in the scheduler's batch queue. Thus, if a job is not solved after this long waiting period, then it most probably is a hard problem. Therefore batch jobs are only used when the problem is hard. When a batch job starts execution, GridSAT migrates work (as a checkpoint file) to achieve more efficient use of batch nodes. Remote GridSAT nodes, which are numerous, will migrate immediately to occupy batch nodes. After, migration takes place and since networks are fast within super-computing nodes, splitting happens at higher rates especially after the above mentioned reductions in communication overhead. Moreover the GridSAT scheduler senses the additional bandwidth between clients executing on a supercomputer or cluster. It then increases the size and number of clauses shared by subproblems inside the tightly coupled resource as a further improvement. Note that the number of active nodes (i.e. those with subproblems) will increase exponentially. This happens because the number of new subproblems is increased in proportion to the number of existing active solvers. Problem migration leads to a more efficient use of batch jobs.

4 Grid Implementation

4.1 Application Characteristics

The GridSAT application is different from most high-performance computing applications. In general, these applications are composed of alternating steps involving computation and communication. The computation and communication intervals do not overlap. Also the communication steps are used as synchronization barriers which enable the various components of the application to exchange information. Moreover, these applications use a predetermined set of compute resources throughout their execution.

Our application differs in much of the above aspects. The GridSAT application has variable resource requirements depending on the problem instance. The number of resources and duration of use of those resources cannot be predicted in general for satisfiability instances. In fact, the set of active resources which are assigned parts of the search space during runtime is dynamic. Resources are added each time the problem is split. Also resources are released immediately after a subproblem is solved. There can be many simultaneous acquiring of new resources, through problem splitting, and release of other unneeded resources at any given instance. Moreover, the application components shared intermediate results as soon as they are produced. These results are asynchronously used by all the receiving clients.

Therefore, all the GridSAT segments are event driven and events are produced and consumed asynchronously. The solver components, for instance, can simultaneously perform communication and computation. All application modules are designed and implemented to allow for efficient management and responsiveness to these events.

Dynamic resource usage are needed, in general, to efficiently solve any satisfiability problem [5]. Solving “hard” satisfiability problems represents further challenges. For “hard” problems, a small number of resources would be exhausted in a relatively short time. The CPU and memory resources would be saturated and additional resources are required in order to make progress in solving the problem under investigation. Therefore, we wanted to use all computational resources at our disposal, in order to render the solution of the hardest problems more plausible. The set of available resources varied from desktop machines, to small-size clusters, to supercomputers. This collection of resources was heterogeneous in terms of hardware, Operating Systems and resource management software. This heterogeneity represents a further challenge to the deployment of the application.

These application characteristics described above represent a true Computational Grid application. Moreover, these characteristics are not unique to GridSAT. Other branch-and-bound or coordinator-worker applications can benefit from a similar use of computational resources.

A major challenge before implementing the various application components was to develop an implementation strategy. The final implementation aims at using all the available grid resources efficiently while dynamically adjusting to the application behavior and resource needs.

4.2 Implementation Strategy

Given these resource usage patterns, which are typical for a true Grid application, we had to choose an implementation strategy which would satisfy these requirements. There are several technology choices to select for the implementation of the application. Such options include, among others, MPI [9], Globus [10], vanilla Web Services [33] and later improvements such as WSRF [24].

According to our experience with GridSAT we have learned that a successful implementation technology should allow for three pivotal capabilities: dynamic resource pool, error detection and universal deployment.

The first capability is to allow the use of dynamic a resource pool. This feature, for example, was not available in MPI-I which did not allow for dynamic Communicators. MPI-2 has introduced extensions to allow for dynamic creation and destruction of communicators. Globus and Web services also allow for a dynamic set of resources.

The second capability is error detection and reporting. Since GridSAT runs for extended periods of time using a set of geographically distributed resources, then network and resource failures are more frequent. Therefore in order to implement this application we need a technology which allows for the detection of these errors. From the perspective of the application, the distinction between resource and network failures is not important. It suffices for the application to obtain a feedback if a certain operation is not successful after a certain time period.

Error detection and recovery is very important because in our experience all resources experience a failure at some point. Even those resources which are professionally maintained can become unresponsive from the application’s perspective. Those resources that do not experience hardware and software failures usually have scheduled routine preventive maintenance periods or a combination of software and hardware upgrades. From the point of view of the application these are “scheduled” or “anticipated”

failures. Without rigorous error handling the application would not be able to run for extended periods as shown later in the results section.

Different technologies provide some form of error handling. MPI-I allows for error handling in a limited scope which is expanded further in MPI-2. Globus GRAM allows for error handling and call-back functions for job management. In Web Services, WS-Notification [14], WS-BaseFaults [34] and related standards could be used to provide this functionality. The desirable error handling for our application is to provide a time period for some actions after which some form of error handling should be performed. Sometimes if an action fails, then all is needed is to retry it. In other cases, it is assumed that the resource (or the connecting network) has failed. This form of error handling is not available for the grid technologies mentioned above and can be implemented at the application level.

The last desirable capability for a suitable grid technology is universal deployment. This is not entirely a characteristic of the technology but of the computational environment as well. A widely deployed technology is advantageous because it reduces the development overhead since one version can be deployed on all available resources. In our experience there was no grid technology that was universally adopted and deployed which would enable us to combine all computational resources at our disposal.

Furthermore, in order to deploy our application over a large set of resources, we had to interface with many types of resource managers. For example, resources could be managed by one of many Batch schedulers, Condor [31] or simply shared. Our goal was to use all these resources simultaneously regardless of what systems they originate from. This is accomplished by determining a general job description which can be instantiated differently using specific launchers for each resource manager. For instance, shared resources can be accessed directly using SSH. Batch systems, however, are accessed by submitting a batch script with syntax tailored to the scheduler used. Whenever, Globus is deployed we use it to launch and monitor job submissions.

4.3 GridSAT Implementations

We believe that many of these technologies could be used to develop GridSAT. In fact, we have developed a previous versions of GridSAT called GrADSAT [4] (note the “A” in the spelling) using *GrADSoft*. GrADSoft is a set of programming abstractions where the baseline grid infrastructure is provided by Globus and the NWS. GrADSoft is part of the **Grid Application Development Software** (GrADS) project [2, 13] which is a comprehensive research effort studying grid programming tools and application development. To facilitate experimental application research and testing, the project maintains a nationally distributed grid of resources for use as a production testbed. Since the GrADS tools were universally deployed on this testbed we were able to deploy our application with little overhead on the entire testbed.

The current version of GridSAT uses EveryWare [36] a very portable communication library. EveryWare has been designed explicitly to manage the heterogeneity and dynamism inherent in grid resource environments. EveryWare can be easily deployed as library on all the resources. In addition, all communication calls use a timeout argument, as desired, for error detection.

The resource management system interfaces with resources which use batch systems as well as desktop machines which are accessed through SSH. All resource related operations have been implemented to allow for a specific timeout. If the resource is not responsive after the timeout period expires, then the resource is considered unreachable.

In the future, we will explore other technologies as they become more widely used. Our goal would

be to make GridSAT implementation independent where we can use an API for interfacing the application with the underlying communication infrastructure. As a result different grid technologies can be substituted without affecting the application.

5 Experimental Apparatus and Results

Since GridSAT is a true grid application, (robust, portable, heterogeneous, pervasive, etc. [11]) we ran a set of experiments to show that GridSAT can run for extended periods of time robustly using a wide variety of resources and also solve previously unsolved hard satisfiability instances. In these experiments we simultaneously use computational resources that belong to collections of individual machines, small size research clusters and super-computing scale clusters. The computational resources we use are composed from four main sources:(1) 40 machines from the VGrADS [35] testbed located at UTK, UCSD and UCSB, (2) Blue Horizon at SDSC, (3) TeraGrid site at SDSC, (4) TeraGrid site at NCSA and (5) DataStar at SDSC.

The TeraGrid [32] project is a multi-site national scale project which is aimed at building the worlds largest distributed infrastructure for open scientific research.

During our experiments, none of the resources we used were dedicated to our use. As such, other applications shared the computational resources with our application. It is, in fact, difficult to determine the degree of sharing that might have occurred across all of the available machines after the fact. In batch controlled system such as Blue Horizon, Data Star and the TeraGrid, the queue wait time incurred is highly variable because of jobs submitted by other users.

Thus, if it were possible to dedicate all of the VGrADS resources to GridSAT, we believe that the results would be better. As they are, they represent what is currently possible using non-dedicated Grids in a real-world compute setting.

These experiments also use a more diverse set of resources for longer periods of time (up to a month in duration) and multiple job requests. We chose a set of challenge problems from both SAT2002 [25] and SAT2003 [26] benchmarks. These benchmarks are used to judge and compare the performance of automatic SAT solvers at the annual SAT conference. All the problems in the benchmarks are shuffled to insure that submitted benchmarks are not biased in favor or against any solver. These benchmarks are used to rate all competing solvers. They include industrial and hand-made or randomly generated problem instances that can be roughly divided into two categories: *solvable* and *challenging*. The solvable category contains problem instances that some SAT solvers have solved correctly. They are used for comparing the speed of competing solvers. Alternatively, the challenging problem suite contains problem instances that have yet to be solved by an automatic method or which have only been solved by one or two automatic methods, but are nonetheless interesting to the SAT community. Some of these problems have known solutions that are known through analytical methods (i.e. the problem has a known solution by construction), but several of these problems are open questions in the field of satisfiability research.

In these experiments, we only chose problems from the challenging set. These problems were deemed hard by all participating solvers in both the 2002 and 2003 SAT competitions. We investigate seven previously unsolved problems where three instances are from the SAT 2003 benchmark category, and four are instances from the SAT 2002 benchmark category, all of which we have not been able to solve using previous versions of GridSAT.

This group of problems represent a variety of fields where problems are reduced to instances of sat-

File name	Description	SAT/UNSAT/*	Time	GridSAT Result
3bitadd-31.cnf	theoretical	UNSAT	8 days	-
k2fix-gr-rcs-w8.cnf	FPGA Routing	*	83261 sec (23 hours)	UNSAT
k2fix-gr-rcs-w9.cnf	FPGA Routing	*	14 days and 8 hours	UNSAT
cnt10.cnf	Theoretical	SAT	13134 sec (4hours)	SAT
comb1.cnf	Model Checking	*	11 days	-
f2clk50.cnf	Model Checking	*	9 days	-
hanoi6.cnf	Theoretical	SAT	23 days	-

(*): problem solution initially unknown

Table 1. GridSAT results using VGrADS testbed, Blue Horizon, Data Star and TeraGrid. All these problems were not previously solved by any other solver.

isfiability and solvers are used to determine the solutions. The problems contain a pair of problems in FPGA routing and model checking. These two disciplines benefit heavily from efficient SAT solvers. The remaining problems are of theoretical nature. In addition, we set the absolute minimum size of shared clauses to two and absolute maximum to 15. This range allows for sharing clauses which would help prune the search space without significant communication overhead.

Unlike previous experiments there was no timeout value set for the maximum execution time. Every problem was run using different job description for the batch systems. Jobs on the different batch queues were manually re-launched at random intervals. Job re-submission could have been automated but we wanted more control over rationing our limited compute budgets to specific experiments based on their perceived progress. Experiments where GridSAT was making progress were allotted bigger jobs with longer durations and more nodes. The progress of the solver was judged by inspecting how often the checkpoints were updated. We can also inspect the internal state of a particular solver using some of the tools we developed. The VGrADS nodes were used during the entire duration of each experiment unless the hosts experienced failures.

5.1 Results

The experimental results are summarized in Table 1. The first column contains the problem file name. The second column indicates the field from which this problem instance in obtained. The third column contains the solution to the instance: satisfiable (SAT), unsatisfiable (UNSAT), or unknown. We have marked those problem instances which were previously open satisfiability problems with an asterisk (*). If a problem was originally unknown and was later solved by a solver, then we still keep it marked with an asterisk for completeness. The fourth column represents the total wall-clock time that the problem was tried. Finally, the fifth and last column represents the solution obtained by GridSAT which is represented by SAT, UNSAT or (-) if we terminated the experiment before GridSAT found an answer. Note that while we terminated these problem instances manually so that we could complete this paper, each can be continued from its last checkpoint (which we have archived).

Table 1 shows that GridSAT was able to solve three problems all of which were not previously solved. Two of the problems were found unsatisfiable and they are both from the field of FPGA routing. The

first problem *k2fix-gr-rcs-w8.cnf* was solved using the VGrADS testbed only. Batch jobs which were submitted for this experiment were canceled when the problem was solved. On the other hand the second problem *k2fix-gr-rcs-w9.cnf* took much longer to solve, it took more than two weeks. Table 2 gives a more detailed description of the resource used during this experiment. For each job a number of GridSAT solver components were launched as indicated in the last column of table 2. In table 3 a break down of the CPU-hours used on each resource are tabulated. Note that the VGrADS testbed machines were able to deliver a sizable amount of compute power because they were available in a shared mode for the duration of the experiment.

The last problem *cnt10.cnf* was also solved using the VGrADS testbed only under similar circumstances to *k2fix-gr-rcs-w8.cnf*. We previously tried solving this problem in [5] using the same testbed for four days in addition to Blue Horizon for 12 hours but were not successful. We believe the improvements made to the solver and especially the new clause sharing method have helped achieve this result.

In order to illustrate further GridSAT’s success in using all the above variety of resources mentioned earlier we present a section of a run using instance *hanoi6.cnf*. This problem is a SAT representation of the *Hanoi Towers* problem using six disks. A six day snapshot from a 23 day run is shown in figure 3(a) using logarithmic scale. The figure shows several jobs from Blue Horizon, Data Star and TeraGrid sites participating in the execution. This figure shows that GridSAT was able to make use of the available resource when some of their nodes became available and then continued to run after the nodes were taken away to serve other users. GridSAT processes continue to run on the batch controlled resources until the scheduler decides to terminate them. This abrupt termination has no effect on the application which deals with these events as (scheduled) resource failures. GridSAT was able to manage up to 350 processes running on different resources as show in this figure.

The satisfiability solver performs mostly integer, branching and load/store operations. The number of floating point operations is very low (less than .1 FLOPS). We present in figure 3(b) an estimate of the total number of instructions per second during the same six day period. Since instrumenting GridSAT can cause significant slow down, we conducted some benchmarking on some machines at UTK to determine the average efficiency of the solver. Since the solver code is mostly sequential, we assume that at the maximum only one instruction per cycle can be finished by the processor. The determined efficiency is 70%. We estimated that other hardware and OS combinations will exhibit equal efficiencies. The number of operations provided by a resource is estimated to be the product of its peak performance and the estimated efficiency. The total number of instructions in figure 3(b) is the sum of operations of all active resources. We notice that the VGrADS testbed is able to deliver about 20 Billion instructions per second (IPS). In the middle of the graph, there is a batch job from Blue Horizon

Compute resource	Job count	Job dur.(hr)	Node count	procs /node
BlueHorizon	2	10	100	3
Blue Horizon	1	12	100	3
DataStar	2	10	8	11
TG@SDSC	1	10	40	2
TGd@SDSC	1	12	40	2
TG@SDSC	3	10	4	2
TG@SDSC	4	5	4	2
TG@NCSA	3	10	4	2
TG@NCSA	4	5	4	2

in addition to 40 machines from VGrADS testbed for 14 days 7 hours and 44 minutes

Table 2. Batch jobs used to solve the k2fix-gr-rcs-w9.cnf instance from SAT 2003 benchmark

which failed suddenly while joining the GridSAT execution. This might have happened because the Blue Horizon machine became unavailable for scheduled maintenance. The total number of IPS was multiplied by more than five times when some batch jobs became active. It reached up to 110 Billion IPS.

Another measure of performance, is how much of the batch job maximum computational power is actually used by GridSAT processes. Most other parallel jobs run on all the processes from start to finish with little overhead. In this case, batch jobs are efficiently used. In the of case GridSAT, however, there are two main sources of inefficiency. First, some jobs might wait ideally at the start. Batch jobs usually include a large number of processes. Some of these processes have to wait until a sufficient number of splits occur to generate new sub-problems for all the newly created solvers. Second, some batch processes may contain idle solvers for a period of time after they solve the previously assigned sub-problem. The solver in this case, waits until it is assigned a new sub-problem by the master. For the first job in figure 3(a), which is a large 100-node job, the efficiency is 98.9%. Thus GridSAT was able to use batch jobs efficiently. The main reason is that batch jobs usually wait in the batch queue for a long time before executing. Thus by the time the job is executed, GridSAT was unable to solve the problem because it is hard. This means that batch jobs are only used when the problem is in deed hard. It is possible that for certain problems, the efficiency of batch jobs might be low. In this case, future versions of GridSAT might monitor the batch job efficiency to determine whether and when a job is to be terminated.

During our experiments, the Blue Horizon super-computer was being decommissioned. GridSAT was able to continue running experiments on the set of available resources through this transition. The scheduler would try to submit jobs but it would notice that the Blue Horizon resource was not responding. The failure of this single (but important) resource which did not affect the already running experiments shows the robustness of GridSAT.

6 Conclusion

This paper presents a new version of GridSAT which implements a parallel, distributed and complete satisfiability solver. In order to solve harder problems, new improvements to both the algorithm and architecture of GridSAT were introduced. GridSAT is capable to dynamically selecting resources to enable improved overall performance.

The experiments we presented show GridSAT’s ability to manage and use a diverse set of dynamic computational grid resources. The experiments lasted for weeks as a testament to the robustness of the application. During these experiments new previously unsolved problems from practical and theoretical fields were solved.

Compute resource	node-hours	CPUs/node	CPU -hours
BlueHorizon	3200	8	25600
DataStar	160	11	1760
TG@SDSC	1080	2	2160
TG@NCSA	200	2	400
GrADS(*)	13750	1	13750

(*) machines were shared with other users

Table 3. CPU-hours per resource used to solve the k2fix-gr-rcs-w9.cnf instance from SAT 2003 benchmark




Figure 3. A six day snapshot representing (a) processor count usage in logarithmic scale and (b) estimated instructions per second (IPS) usage for all resources during a six day period.

References

- [1] G. Allen, D. Angulo, I. Foster, G. Lanfermann, C. Liu, T. Radke, E. Seidel, and J. Shalf. The Cactus Worm: Experiments with dynamic resource discovery and allocation in a Grid environment. *The International Journal of High Performance Computing Applications*, 15(4):345–358, 2001.
- [2] F. Berman, A. Chien, K. Cooper, J. Dongarra, I. Foster, L. J. Dennis Gannon, K. Kennedy, C. Kesselman, D. Reed, L. Torczon, , and R. Wolski. The GrADS project: Software support for high-level grid application development. *International Journal of High Performance Computing Applications*, 15(4), Winter 2001. available from "http://hipersoft.cs.rice.edu/grads/publications_reports.htm".
- [3] A. Biere. Limmat. <http://www.inf.ethz.ch/personal/biere/projects/limmat/>.
- [4] W. Chrabakh and R. Wolski. GrADSAT: A Parallel SAT Solver for the Grid. Technical Report 2003-05, UCSB, March 2003.
- [5] W. Chrabakh and R. Wolski. GridSAT: A chaff-based Distributed SAT solver for the Grid. In *Supercomputing Conference, Phoenix, AZ*, November 2003.
- [6] K. Czajkowski, S. Fitzgerald, I. Foster, and C. Kesselman. Grid information services for distributed resource sharing. In *Proc. 10th IEEE Symp. on High Performance Distributed Computing*, 2001.
- [7] M. Davis, G. Logeman, and D. Loveland. A machine program for theory proving. *Communications of the ACM*, 1962.
- [8] S. L. Forman and A. M. Segre. Nagsat: A randomized, complete, parallel solver for 3-sat. SAT2002, 2002.
- [9] M. P. I. Forum. Mpi: A message-passing interface standard. Technical Report CS-94-230, University of Tennessee, Knoxville, 1994.
- [10] I. Foster and C. Kesselman. Globus: A metacomputing infrastructure toolkit. *International Journal of Supercomputer Applications*, 1997.
- [11] I. Foster and C. Kesselman, editors. *The Grid – Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1998.
- [12] E. Goldberg and Y. Novikov. BerkMin: A fast and robust SAT-solver. In *Design, Automation, and Test in Europe (DATE '02)*, pages 142–149, March 2002.
- [13] GrADS. <http://hipersoft.cs.rice.edu/grads>.
- [14] S. Graham and B. Murray. <http://docs.oasis-open.org/wsn/2004/06/wsn-WS-BaseNotification-1.2-draft-03.pdf>.

- [15] E. A. Hirsch and A. Kojevnikov. UnitWalk: A new SAT solver that uses local search guided by unit clause elimination. In *PDMI preprint 9/2001, Steklov Institute of Mathematics at St.Petersburg*, 2001.
- [16] D. Jackson and M. Vaziri. Finding bugs with a constraint solver. *International Symposium on Software Testing and Analysis*, 2000.
- [17] B. Jurkowiak, C. M. Li, and G. Utard. Parallelizing Satz Using Dynamic Workload Balancing. In *Proceedings of Workshop on Theory and Applications of Satisfiability Testing (SAT'2001)*, pages 205–211, June 2001.
- [18] W. Kunz and D. Stoffel. *Reasoning in Boolean Networks: Logic Synthesis and Verification Using Techniques*. Kluwer Academic Publishers, Boston, 1997.
- [19] T. Larrabee. Efficient generation of test patterns using boolean difference. pages 795–802.
- [20] T. Larrabee. Test pattern generation using boolean satisfiability. In *IEEE Transactions on Computer-Aided Design*, pages 4–15, January 1992.
- [21] W. W. Li, R. W. Byrnes, J. Hayes, A. Birnbaum, V. M. Reyes, A. Shahab, C. Mosley, D. Pekurovsky, G. B. Quinn, I. N. Shindyalov, H. Casanova, L. Ang, F. Berman, P. W. Arzberger, M. A. Miller, and P. E. Bourne. The encyclopedia of life project: grid software and deployment. *New Gen. Comput.*, 22(2):127–136, 2004.
- [22] M. Moskewicz, C. Madigan, Y. Zhao, L. Zhang, and S. Malik. ”Chaff: Engineering an Efficient SAT Solver. 38th Design Automation Conference (DAC2001), Las Vegas, June 2001.
- [23] G. Nam, F. Aloul, K. Sakallah, and R. Rutenbar. A Comparative Study of Two Boolean Formulations of FPGA Detailed Routing Constraints. *International Symposium on Physical Design (ISPD), Sonoma Wine County, California*, pages 222–227, 2001.
- [24] OASIS Web Services Resource Framework (WSRF) TC. http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsrf.
- [25] SAT 2002 Competition. <http://www.satlive.org/satcompetition/>.
- [26] SAT 2003 Competition. <http://satlive.org/SATCompetition/2003/>.
- [27] M. H. Schulz and E. Auth. Improved Deterministic Test Pattern Generation with Applications to Redundancy Identification. *IEEE Transactions on ComputerAided Design*, 8(7):811816, July 1989.
- [28] J. M. Silva and K. Sakallah. Grasp - a new search algorithm for satisfiability. ICCAD. IEEE Computer Society Press, 1996.
- [29] J. P. M. Silva. Search Algorithms for Satisfiability Problems in Combinational Switching Circuits. Ph.D. Thesis, The University of Michigan, 1995.
- [30] C. Sinz, W. Blochinger, and W. Kuchlin. PaSAT - Parallel SAT-Checking with Lemma Exchange: Implementation and Applications. In *Proceedings of SAT2001*, pages 212–217, 2001.
- [31] T. Tannenbaum and M. Litzkow. The condor distributed processing system. *Dr. Dobbs Journal*, February 1995.
- [32] TeraGrid. <http://www.teragrid.org/>.
- [33] The W3C Web Services Architecture working group, public draft, August 2003. <http://www.w3.org/TR/2003/WD-ws-arch-20030808/>.
- [34] S. Tuecke, L. Liu, and S. Meder. <http://docs.oasis-open.org/wsrp/2005/03/wsrp-WS-BaseFaults-1.2-draft-04.pdf>.
- [35] VGrADS. <http://hipersoft.cs.rice.edu/vgrads>.
- [36] R. Wolski, J. Brevik, C. Krintz, G. Obertelli, N. Spring, and A. Su. Running EveryWare on the Computational Grid. In *Proceedings of SC99*,, November 1999.
- [37] R. Wolski, N. Spring, and J. Hayes. The network weather service: A distributed resource performance forecasting service for metacomputing. *Future Generation Computer Systems*, 1999.