GridSAT: A System for Solving Satisfiability Problems Using a Computational Grid

Wahid Chrabakh
and
Rich Wolski
Department of Computer Science
University of California Santa Barbara
{*chrabakh,rich*}@*cs.ucsb.edu*[1]

In this paper, we present GridSAT – a distributed and high performance complete satisfiability solver – and its application to a set of complex and previously unsolved satisfiability problems. Based on the sequential Chaff [32] algorithm, we combine new distributed clause "learning" techniques with an efficient and autonomous grid implementation both to speed the time to solution and to solve problems too complex for other general solvers. By automatically adapting to changes in the availability of machines and carefully distributing the clause database, we show how GridSAT has been able to use, continuously, a diverse and dynamically changing resource pool to solve previously unsolved problems from the SAT 2002 [39] and the SAT 2003 [42] competitions. We describe our enhancements to the Chaff learning algorithm that have enabled an efficient distributed implementation, and detail the technological approach we have taken to realizing this implementation. In addition, we present a web portal for GridSAT that accepts problem instances through a standard browser and returns status and results while shielding users from the complexities associated with running the application manually.

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

1 Introduction

Grid computing [23] is an emerging field in computer science which enables the aggregation of geographically distributed and federated computational resources. One important goal outlined in the original grid vision [19] is the ability to apply computing and storage capabilities aggregated from a wide variety of sources to problems that require more "power" than is available from any single site. While several applications have attempted to achieve this goal [30, 2, 36] they have done so sharing two important characteristics. First, these scientific applications exhibit fairly regular and predictable computational patterns. Secondly, while they have been able to employ widely dispersed resources, the resources they use must be dedicated exclusively to each program execution. That is, using the parlance of the grid metaphor, most applications to date have predictable power requirements and require dedicated power generators to achieve acceptable performance.

For grid computing to become a successful technique that is more generally applicable in a wider setting, it must be able to support applications which solve problems that require data dependent (e.g. unpredictable) execution patterns and also can take advantage of non-dedicated resources. One such challenging problem is propositional satisfiability – the determination of a set of assignments to boolean variables in an arbitrary expression that makes the expression's value logically "true." The ability to solve satisfiability problem instances is important from both theoretical and practical perspectives. Satisfiability (SAT) is theoretically significant because it was the first problem to be proven NP-complete [11]. Thus solving SAT instances, in general, requires large computational power for extended time periods.

In practice, many engineering disciplines require the solution of domain-specific instances of satisfiability. Such disciplines include scheduling [7], model checking [3], security [1], Artificial Intelligence [26] and software verification [24]. Satisfiability is especially important in the area of Electronic Design Automation (EDA). EDA encompasses a variety of problems such as circuit design [45], Field-Programmable Gate Arrays (FPGA) detailed routing [33], combinational equivalence checking [27, 35] and, automatic test and pattern generation [29]. Our goal is to be able to find solutions to those practical problem instances that do not require an exponential number of possible variable assignments to be tested to

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

determine whether the instance is satisfiable or not.

There are many related problems to satisfiability which are also computationally intensive, such examples include $\#SAT$ and *solution enumeration*. The $\#SAT$ problem only asks to determine the number of solutions for a given SAT problem instance. Solution enumeration problems, however, require the listing of all solutions to a SAT problem and not just the number of solutions or whether the problem is satisfiable or not. The solution enumeration problem is important because in many cases it is desirable to find all solutions to a problem or at least a representative subset of the solution set. In [24], the author presents a motivation for solution enumeration and how it can be used to improve software testing procedure. Similarly, solution enumeration could generate multiple solutions to a scheduling problem [7]. These solutions would present alternative solutions to choose from instead of being restricted to a single one. In another example, a circuit designer with access to multiple solutions can select the solution that best suits his needs. Moreover, in cases where satisfiable solutions represent design errors multiple solutions provide more information about the sources of error and may lead to quicker determination of the source of error. Solutions to the $\#SAT$ and solution enumeration problems can be derived from solutions to the original satisfiability problem.

There has been extensive research efforts focused on the development of efficient satisfiability solvers [32, 20, 22, 5]. These solvers use different techniques to navigate the entire search space. When possible, the most efficient of these solvers use optimizations which permit parts of the search space to be discarded or "pruned" during execution. However, because the general problem is NP-complete, there is no theoretical framework for comparing solvers or evaluating which solver is best suited to a particular problem or problem class. The solvers and the techniques they implement are evaluated based on empirical results by comparing the speed with which they can solve a diverse set of benchmarks and/or the number of complex or "hard" problems they can solve. Thus while the general problem remains theoretically intractable, heuristic-based approaches have yielded SAT solvers that serve as valuable verification tools in many disciplines.

Most modern solvers [32, 20, 22, 5] are sequential and employ heuristic improvements to one of a small set of fundamental search algorithms. Fewer parallel solvers such as [9, 25, 46, 17] exist, and even fewer of those parallel solvers use a heretofore sequential optimization termed *learning*. Learning (discussed in detail in Section 2.2) improves solver speed by adding propositions that the algorithm deduces to an internal database that is global to the solver. These additional "learned" propositions improve the efficiency of SAT solvers substantially, but they make the problem of parallelizing and/or distributing a solver daunting. The global clause database must be searched and updated frequently as the algorithm progresses making an efficient large-scale parallel or distributed implementation difficult. As a result, the best known solvers (in terms of speed and solution power) have until recently been sequential. One of the innovations we explore is a new technique for partially sharing learned clauses between distributed solver elements that yields improved solver speed despite the additional communication and storage overhead implied by a distributed implementation.

These sequential solvers are characterized by heavy use of compute power (CPU) as well as the memory of the host machine (which is used to hold the clause database). As a solver progresses, and more clauses are learned, it uses more of the host's memory typically until non-paged memory is exhausted. Thus, the typical sequential solver completes when one of three conditions have occurred: a variable assignment satisfying the expression is found, a user-specified timeout on the solution has expired, or the machine has run out of non-paged memory. Most sequential solvers can be configured to avoid this last termination condition by selectively deleting some of the learned clauses according to heuristics that attempt to determine which clauses are least helpful in directing the search. In this case the solver may progress slower since it has to remove some of the previously learned clauses used in pruning the search space and potentially relearn them. Both the design of these heuristics and techniques for determining the "best" memory size to use are currently active areas of research.

Another major difficulty facing the SAT solver community is the problem of predicting the time to solution for a given SAT instance. One heuristic examines internal solver state to estimate both the rate at which the solver is pruning the search space and the rate at which it is exhausting memory and then extrapolates time-to-solution from these rates. If the ratio of the speed with which the solver prunes the search space to the rate at which it is consuming memory is low, the problem is perceived as being "hard" according to this heuristic. For a hard problem, a learning solver will either run out of memory and terminate, or prematurely discard and then subsequently relearn clauses to avoid memory exhaustion and, because it must run longer, incur a user timeout.

In the rest of this paper, we use the term "hard" or "complex" to refer to problems which require a long time to solve or a great deal of memory or both. We also use the term "easy" to describe those problems which are solved quickly by the solver under consideration. Note that due to the interaction between learning heuristics and the data dependent nature of SAT problems, a SAT problem instance may be perceived differently by two SAT solvers. For example, while one solver may find a problem instance hard, another solver may find the same problem instance easy depending on the heuristics each employs. However, there are SAT problems which are con-

sidered hard by all known SAT solvers because, at present, no solver has been able to find a solution.

One approach to improving solver performance is through parallelization. Previous implementations of parallel solvers show that better performance is obtained when a large pool of computational resources is used for some problem instances. The aggregate CPU power and memory of the hosts make it possible to navigate the search space faster. However, these solvers typically do not employ learning as parallelizing access to the clause database (particularly when shared memory is not available as in a cluster setting) carries with it a potentially heavy synchronization cost. Our approach is to parallelize the solver algorithm and to have the program automatically decide if and when additional resources should be employed in parallel based on a prediction of whether the synchronization cost can be successfully amortized. Because resources will be added "on-demand," our solver must be prepared to change its resource usage dynamically in response to the way in which a particular solution is evolving. Similarly, if resource availability fluctuates causing the amount of available compute power to change, the program adjusts its resource usage to minimize its predicted time-to-solution.

Thus, the nature of our work takes the form of two contributions. First, we have developed a new, efficient method for parallelizing access to the clause database that is well suited to distributed computing environments. This methodology replicates a minimal set of clauses synchronously, and multicasts a set of "intermediate" clauses asynchronously. In this paper we present and compare three different methods for effecting distribution and sharing of the clause database in the context of parallel solvers distributed-memory solvers.

It is still an open research question to decide when using more resources increases the solver's performance. But research in parallel solvers, shows that using more resources can be (but not always is) a performance booster. Thus a computational grid populated by a a large pool of resources offers potential improvements in solver speed. These speed improvements can also enable a parallel SAT solver to solve previously unsolved problems which would have otherwise taken prohibitively longer durations to solve using a sequential solver.

The goal of our research is to develop a satisfiability solver capable of harnessing the computational power provided by a grid infrastructure. In order to enable a grid implementation of a SAT solver to use many resources simultaneously, we need to address two types of challenges.

In order to meet the first challenge, 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. In the context of a parallel solver, the individual components can share intermediate results. In this paper we present and compare three different methods for sharing partial results in a parallel solver. ¿¿¿¿¿¿¿ 1.13

The second contribution is the development of a framework for running the parallel solver in a heterogeneous and very volatile computational environment. In particular, we target a grid [19, **?**] setting in which resources can be categorized into two different types: time-shared and batch controlled. In the case of time-shared resources, the application will compete with other user applications running simultaneously on the host processors. These processors may be part of a time-shared cluster, a laboratory of time shared workstations, desktop machines, etc. Their distinguishing feature, however, is that they are available immediately if they are available at all. That is, a job or process need to wait in a queue for exclusive access to them.

The other class of resources our solver is able to use are those resources which are controlled by some form of batch scheduler. For these resources, our solver must submit requests for processing (jobs) that are queued until sufficient resources are available. These batch-controlled resources are typically space shared so a job can request multiple resources (typically processors or nodes) be allocated simultaneously, and once resource are allocated to a job, the job has them exclusively for some site-specific maximum amount of occupancy time.

Large-scale batch resources are also, typically, oversubscribed (hence the need for queuing to implement space-sharing). Thus, job turnaround time often is characterized by long queuing delays compared to relatively short maximum possible occupancy periods. To a long-running GridSAT instance, these batch systems offer periodic "bursts" of computing power separated by lengthy intervening hiatuses. Moreover, while the amount of time a GridSAT job will be able to execute is known once the job begins executing, predicting the queuing delay associated with an individual job submission remains a difficult research problem [**?**, 15, 16].

Thus, the execution model for GridSAT uses immediately available interactive resources continuously and batch-controlled resources opportunistically. Each instance begins execution on a single interactive resource. As the problem execution progresses, GridSAT examines the rate of its progress and issues requests for additional resources in response the the problem's perceived difficulty. After a specified number of interactive resources have been consumed (typically those local to the launching site), GridSAT begins issuing requests for large-scale batch resources. While waiting for batch resources to become available, GridSAT maintains a work backlog (in the form of stored checkpoints) for work it would have initiated had interactive resources been available. This backlog is mined whenever a GridSAT batch request is eventually scheduled on a set of processors. If the request is scheduled

before a backlog forms, GridSAT checkpoints the "hardest" of its currently executing subproblems and migrates it one the newly available processors (to take advantage of the exclusive availability). Finally, GridSAT checkpoints all processes – on interactive and batch resources – so that a lost process or a terminated occupancy period does not cause the application to halt. It is the seamless "blending" of cheap, readily available interactive resources and large-scale, heavily shared batch resources, combined with GridSAT's ability to exploit them effectively, that make it a foremost example of the grid computing paradigm.

Our initial work with GridSAT [9] shows that dynamically acquiring resources in the manner described can provide large reductions in time to solution for particular problem instances. Indeed, GridSAT outperforms the best-known solver on all problems that this leading solver can complete. At the same time, GridSAT uniquely has been able to solve several previously unsolved problems using non-dedicated, wide-area Grid resources. Thus, by using Grid resources effectively, GridSAT constitutes a speed improvement over the fastest-known technique and has achieved new scientific results that have not previously been possible. Few grid applications or infrastructures have been able to realize the grid metaphor as completely, with such effective results.

Having established GridSAT's performance benefit over the state-of-the-art, in this paper we focus exclusively on the use of GridSAT to attack additional previously unsolved problem instances. New optimizations and a better integration of batch and interactive resource usage makes it possible to solve previously unsolved satisfiability problems from the field of FPGA [33] routing as well as artificially generated instances. In this work we also present a version of GridSAT that is capable of solution enumeration. GridSAT solves the solution enumeration problem by extending the sequential solver. To the best of our knowledge, most sequential and all parallel solvers lack this feature.

We also make the GridSAT solver available to potential users through a web portal `http://orca.cs.ucsb.edu/sat_portal/`. This portal allows users to easily use a parallel solver deployed over a pool of powerful resources which are not readily accessible to most users. By hiding all the complexity of the application and resources used, the portal makes it accessible for users to experiment with their domain related problems.

The paper's remaining sections are organized as follows. Section 2 introduces the basic SAT solver algorithm and some of the more advanced techniques used in modern solvers. In section 3 we present GridSAT's parallel version of the algorithm and the techniques used to increase its efficiency. In section 4, we present how solution enumeration is implemented in GridSAT. The architecture and resource scheduler of GridSAT is presented in section 5. We present experimental setup and results in section 6. Finally, we conclude in section 7.

2 Sequential Solvers

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). An expression in CNF is a conjunction (logical AND) of *clauses* each an injunction (logical OR) of *literals*. A literal is either an instance of a variable (V) or its complement($\sim V$). A problem is called satisfiable if there exists a set of variable assignments that makes the formula evaluate to *true*. If such an assignment does not exist the the problem is declared *unsatisfiable*. The CNF has two important properties: any boolean formula can be algebraically converted to CNF, and for the original formula to be satisfiable all constituent clauses must be satisfiable.

GridSAT is based on Chaff [32], a sequential SAT solver. In Chaff, as well as other solvers, the performance of the algorithm is enhanced by using techniques for adding new deduced clauses. In this section we explain the basic algorithm and how new clauses are generated.

2.1 The Basic Algorithm

The basis of Chaff and many modern SAT solvers is the Davis-Putnam-Logeman-Loveland (DPLL) [14] algorithm. Figure 1 shows a simplified flow chart describing the algorithm. This algorithm and its derivatives belong to the family of "complete" solvers that are guaranteed to find an instance of satisfiability if the problem is satisfiable, or to terminate once a sufficient set of all possible variable assignments have been examined proving that the problem is unsatisfiable. Variables can be assigned the values *true* or *false* but they are all marked as *unknown* initially. The algorithm uses heuristics to assign values to variables speculatively, but in an order that is likely to yield a truth assignment quickly if one exists. The speculative assignment of values to variables is called a *decision*. Because decisions are speculative (and may be undone) and because decisions have deductive implications, they are maintained as a stack. Each decision has a unique *level* in the *decision stack* with the first level in the decision stack containing variable assignments necessary for the problem instance to be satisfiable. For example, variables in clauses composed of a single literal will be added to this level. Other variables will be deduced to have a specific value and will be added to the first level as the algorithm progresses.

After each new decision, the algorithm searches for *unit clauses* – ones with a single unassigned literal and all other truth values false. In a unit clause, the last remaining literal

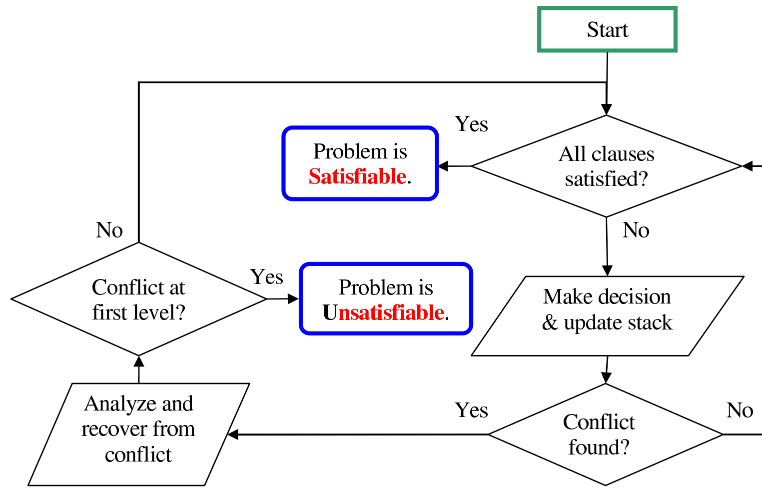

Figure 1: Flow chart for the DPLL algorithm

must have the value *true* for the clause to be *true*. When the algorithm encounters a unit clause, it sets the previously unknown literal to *true*. When a literal is set to *true* because of a unit clause, this is called an *implication*. The corresponding variable is assigned the value that makes the literal *true* and is pushed onto the current decision level. Therefore each new unit clause results in a new variable being assigned a truth value. Even though an implication is a direct result of the previous assignment, it is also predicated on some subset of the previous variable assignments.

In DPLL a variable assignment occurs when a decision is made or a variable is implied. *Boolean Constant Propagation* (BCP) is the procedure where the algorithm inspects the clause *database* in search of unit clauses, after each variable assignment. We use the term *database* in the rest of this paper to refer to the set of clauses used by the solver. Actually we can think of the solver as performing only a very specific *query* but very often. The query is executed after every decision or implication. The query matches all clauses which contain a specific literal. However since efficiency is of utmost importance, all clauses with a particular literal occurrence are indexed using a list of pointers. BCP which occurs after every new variable assignment is the most costly operation and accounts for up to 90% of the runtime [32].

When a decision is made, resulting implications are added to the current decision level. More implications might be added in a cascade because of earlier implications. This process continues until no more implications are found or contradicting assignments to the same variable are detected. In the case when there are no more implications and if not all clauses are satisfied, a new decision is made as shown in figure 1. When a new decision is made an additional decision level is added to the decision stack. In the other case where a contradiction happens, the algorithm has encountered a *conflict*. A conflict occurs when a variable is deduced to have two conflicting values (i.e. being *true* and *false* simultaneously). When a conflict happens the algorithm resolves the contradiction when possible before proceeding. In order to remedy a conflict, a simple approach is to undo the last decision that caused this conflict. The solver can proceed by flipping the value of the previous decision and then try again. In case when a decision has been tried both ways, the first previous decision which can be flipped is tried. If the algorithm cannot find a previous decision which was not tried both ways then the problem is found to be unsatisfiable. This method is slow and may require trying all 2^N combinations of variable assignments when the problem is unsatisfiable, where N is the number of variables. More sophisticated conflict analysis techniques are presented in the next section. These conflict analysis procedures select a level in the decision stack to which the algorithm can back-jump (as opposed to back tracking a single level). Non-chronological back-jumping [60] occurs if the algorithm jumps by more than one decision level. After back-jumping the algorithm continues by making a new decisions or deducing new implications.

Eventually the algorithm terminates under one of two possible conditions. If the problem is satisfiable, a set of variable assignments which result in all clauses evaluating to *true* is found. This termination condition occurs when all clauses are satisfied because of the current set of variable to assignments. Note that not all variables need to be assigned a truth value for this to happen. The problem is deemed unsatisfiable if the algorithm backtracks or back-jumps completely to the first decision level and there is a conflict due to deduced variable assignments at this level. Since the variable assignments at this level are necessary for the problem to be satisfiable then

this is a conflict that the algorithm cannot resolve. Therefore, the algorithm concludes that the formula is unsatisfiable.

2.2 Conflict Analysis and Learning

One of the more sophisticated and effective methods of conflict analysis is *Learning*. Learning [43, 28, 44] is the augmentation of the initial formula with additional implicate clauses that are deduced during the search procedure. The addition of these clauses restricts the search space and prevents the solver from retrying those parts of the search tree. Learned clauses represent redundant information because they are deduced from the initial set of clauses – that is, they carry no additional logical information that bears upon the satisfiability of the original formula. Thus learned clauses can be discarded without changing the solution set of the initial problem.

In DPLL with learning new implicate clauses are deduced due to a conflict. Conflict analysis is based on implication graphs. An *implication graph* is a DAG which expresses the implication relationships of variable assignments. The vertices of the implication graph represent assigned variables. The incident edges on a vertex originate from those variables that triggered the implication of the represented variable assignment. The implication graph is not maintained explicitly in memory. Instead each implied variable points to the clause that caused its implication. That is, the clause that has previously become a unit clause and caused this variable to be implied (i.e. assume some truth value). This clause is called the *antecedent* of this variable. Note that decision variables have no antecedents because they are not implied. In practice decision variables are given a fictitious antecedent clause. Initial and learned clauses are given indexes greater than 1, thus we use clause 0 (which does not exist) as antecedent for decision variables.

A *learned clause* is obtained by partitioning the implication graph into two sides using a *cut*. One partition is called the *reason side* and contains all the decision variables. The other partition which contains the conflict is called the *conflict* side. The cut is used to generate a new *learned* clause using the literals on the reason side with edges intersecting the cut. Different learning schemes are generated from different partitioning methods. However not all cuts generate clauses which lead to a more efficient algorithm. A cut must be selected in order to make learning effective [60] in improving the algorithm's performance. A trivial partition would result in a clause which includes all the previous decision variables made before reaching the current conflict. However, in many cases not all previous decisions have contributed to the current conflict. Also a more carefully selected cut would have fewer intersections and therefore will produce a smaller clause. Smaller clauses are more effective in pruning the search space than longer ones [9].

The purpose of the new clause is to prevent, in the future, the set of simultaneous assignments which led to the current conflict. The new learned clause is obtained by using the complement of the variables on the reason side with edges intersecting the cut. In addition the conflict clauses cause the solver to perform a non-chronological back-jump. After back-jumping, the new decision level is the highest decision level among all the decision levels of the variables in the new learned clause. Chaff [32] uses a method called FirstUIP. This method is based on finding a *dominant node* to the conflict nodes defined as a node where all paths from the current decision to the conflict pass through. The variable corresponding to the selected dominant node is the only variable added to the learned clause which is not a decision variable. Since there might be many such nodes, the FirstUIP method uses the node closest to the conflict. In this case the cut is made such that all implications between the dominant node and the conflict site are on the conflict side. For a more detailed explanation of the algorithm please refer to [9, 8, 32, 44].

During execution, the number of learned clauses is potentially very large, thereby consuming and ultimately exhausting the memory capacity of any given host. Since all learned clauses represent redundant information, the algorithm can discard them (at a potential increase in execution cost) without affecting correctness. Chaff implements specific heuristics [32] (the details of which are beyond the scope of our work) to select which learned clauses are deleted depending on their size and other properties. Deleting some of the learned clauses periodically alleviates memory use and allows the addition of new learned clauses which are currently more relevant.

3 GridSAT: SAT Solver for the grid

GridSAT's distributed solver addresses three significant challenges. First, GridSAT parallelizes the search algorithm that is navigating the space of possible truth assignments. Second, certain learned clauses from the various solvers are distributed and shared across grid resources. Finally, the GridSAT application components are dynamically scheduled so that they may take advantage of the best possible resources available at the time and they can be used profitably by the algorithm.

SAT problems vary in terms of their resource requirements. The two main resources which affect solver performance are CPU speed and memory size. Greater CPU speed makes execution faster and available memory is used to store learned clauses which may help prune the search space. In practice, uncontrived industrial example problems benefit considerably from clause learning As such, a fast CPU with little memory will result in extremely slow progress.

To apply a parallel search technique to SAT, we split the problem at hand 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. A subproblem represents part of the search space. Clause sharing is facilitated by identifying the important clauses relevant to each side of a split, and by eliminating clauses from the clause database pertaining to each side.

The goal of GridSAT is to keep the execution as sequential as possible and to use parallelism only when it is needed. Because problem difficulty is unpredictable and parallelism overhead could be high, GridSAT attempts to add resources (machines with sizable memory) only when the current resource set (which starts with one machine) becomes overloaded.

3.1 Parallelizing SAT

GridSAT acquires new resources when existing sub-problems are split into two sub-problems covering disjoint, but complementary, parts of the original search space. For GridSAT the split process modifies the current problem and spawns a new one as shown in Figure 2. The left part of figure 2 shows the old decision stack of process A before splitting. This process (also called *client* in GridSAT parlance) was assigned a subproblem and is now splitting its search space with client B. The right part of figure 2 shows the modified problem stack for client A and the newly created problem stack for client B after splitting. The first decision variable in the second decision level of Client A's original stack is the pivotal point in the split. Clients A and B assume two different values for this variable. Since this variable is given a specific value in both clients, then it becomes part of the first decision level in both cases. For client A, all implications which were previously in the second decision level are now also part of the first decision level of the modified decision stack. Therefore, Client A's new decision stack is created by making all variables on the second decision level of the assignment stack part of the first decision level. The newly generated problem stack for client B consists of a set of variable assignments and a set of clauses. The variable assignments include all assignments from the first decision level and the complement of the first assignment in the second decision level of Client A's original stack. Thus insuring the splitting of the search space.

After splitting, each process maintains its own separate clause database. In order to alleviate memory usage, inconsequential clauses are removed. A clause is removed from a client's database when it evaluates to *true* because of the assignments made at the first level of its decision stack as a result of the split. In addition, inconsequential clauses are removed every time the first decision level is augmented.

A notable risk in parallelizing a SAT solver comes from the possibility of excess overhead introduced by parallel execution. In particular, because the duration of execution time that will be spent to solve a subproblem cannot be predicted easily beforehand, it is possible for subproblems to be investigated in such a short amount of time that the overhead associated with spawning them cannot be amortized. As a result a solver spends more time communicating the necessary subproblem descriptions, thinning the database, and collecting the results than it does actually investigating assignment values. Even though the solver is advancing, the execution time may be slower than if it were executed sequentially. This problem is occasionally referred to as the "ping-pong" effect [25].

In the following sections we will describe optimizations introduced to the splitting procedure and clause sharing to help improve the overall solver performance. These optimizations include several aspects:

- Different methods for merging shared clauses
- Adaptive clause sharing
- Reduction of communication overhead during problem transfer

3.2 Sharing and Distributing the Clause Database

Each GridSAT process is assigned a part of the search space disjoint from the search space of all other processes. This is insured by giving each process a unique top decision level in the stack. This level may be augmented but is never reduced. Because of the uniqueness of the stack, solvers will tend to make different decisions which in turn results in varying implications. Therefore, the learned clauses, which are dependent on the decision stack, as well as previous learned clauses, will most probably differ for various processes. Thus when these learned clauses produced by one client are shared with other clients they help prune parts of their search space which they have not yet investigated. The overall effect is improved solver performance.

Allowing clause sharing, however, limits the kind of simplifications that can be made. For example, variables (and their complements) which have known truth assignments (i.e. in the first decision level) can be removed since they will not influence future decisions made by the solver. Removing such variables can be accomplished by deleting the occurrence of all literals with known values from all clauses. This deletion results in shorter clauses and more efficient use of the memory. However, variables of known values in one process might still be unknown in another process. Thus in order for a clause to be still valid when shared with another process it must contain complete variable information. Therefore simplifications such as removing known variables are not possible when clauses are shared because they make learned clauses only valid in the context of the current solver.

When new learned clauses are received from other clients, they are merged with the local clause database. Next we

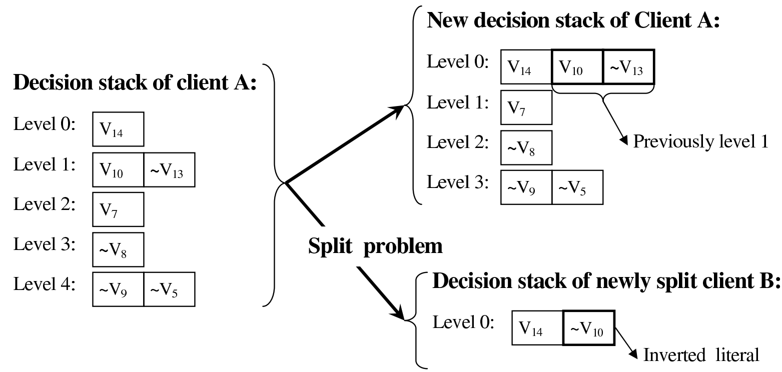

Figure 2: Example of stack transformation when a problem is split into two clients

present and analyze three different methods GridSAT uses for sharing learned clauses. The three methods are: the *lazy method*, the *immediate method* and the *periodic method*. To our knowledge there has been no previous empirical evaluation of the performance of different sharing methods. In the next sections, we explain the characteristics of each method and give motivation for using the two last methods.

3.2.1 The "Lazy Method" for Sharing Clauses: This method was the first we implemented because of its simplicity. In this method, the newly obtained clauses are only merged into the clause database after the algorithm has backtracked to the first decision level. We call this method the *Lazy method* because it might take a long time before the solver backtracks to the first decision level. When using this method, merging the new clause does not involve any stack manipulation because the stack in this case contains one level and no speculative decisions. The only variables to take into consideration are in the first level of the stack. The truth values of these variables will not be altered by subsequent decisions.

Under the conditions outlined previously, merging a received clause is straightforward. The literals of the received clause are examined for their truth values which can be either *true*, *false* or *unknown*. For a given clause there are four possibilities:

- If the clause contains at least one *true* literal, then the entire clause is *true*. Since the decision stack contains no speculative decisions, then the variable corresponding to the *true* literal could only come from the first decision level. Since this variable will always be *true*, then the clause will always be satisfied. Therefore the clause is of no value to the solver since it does not help restrict the search space and is discarded. In the rest of the cases we assume that no literal is *true*.
- If the clause has only one *unknown* literal, and the remaining literals are *false* then an implication is generated. The newly implied variable assignment is therefore predicated only on variables on the first decision level. Thus the implied variable is added to the first level of the decision stack. The clause under consideration is marked as the antecedent for the newly implied variable.
- If the clause has more than one *unknown* literal, then the clause can be used to restrict the search space. In this case the clause is added to the set of learned clauses and the decision stack is not altered.
- If the clause has all literals set to *false*, then this clause is not satisfied by the existing variable assignments and a conflict exists. Since the decision stack contains no speculative decisions, then all the variables in the new clause must be in the first decision level. Therefore we have a conflict because of variable assignments which should be *correct* if the subproblem were satisfiable. Thus the subproblem is unsatisfiable.

The clauses are processed in batches where no BCP is performed until all clauses in the same batch are inspected. During the batch processing, some clauses may be added to the database while new implications are saved to a temporary queue. If there is no conflict after all new clauses are processed, the solver continues by retrieving the queued implications one at a time, adding them to the first decision level and performing BCP as described earlier. If a clause in the batch causes a conflict then the solver terminates immediately.

3.2.2 The Immediate Method for Sharing Clauses: We made an important observation while running different experiments especially with problems that were hard and took a very long time without making progress. We realized that when the problem was hard, all processes were not able to use the clauses received from other processes because none

of them were able to achieve enough progress to backtrack to the first level of the decision stack. Therefore, all the shared clauses which were saved by local solvers wasted valuable memory space but were never used. Thus sharing clauses did not have the desired effect of helping to prune the search space of the local solver. Instead performance was degraded because of wasted memory space. The solution we implemented allows immediate integration of received clauses into the solver's clause database. This method is called the *immediate method* because clauses are merged immediately after they are received. The implementation of this solution is more complex compared with the algorithm above because the decision stack may contain multiple levels of speculative decisions.

The algorithm for merging clauses starts by inspecting the newly obtained clause. The algorithm determines how many literals in the clause have values *true*, *false* or *unknown*. Also the algorithm determines for clauses with a single literal being *true*, the decision level *true_lit_dl* of such a literal. For the given clause it determines the maximum decision level (*false_lit_max_dl*) amongst the decision levels of the literals set to *false*. After determining these value there are only five possible outcomes:

- If the clause is satisfied because of a variable assignment at the first decision level, then this clause is useless for the local solver and is discarded. This case is similar to the first case in the old merging algorithm.

- If the clause has only one *unknown* literal and no *true* literals, then the clause results in an implication. Actually if the clause was available when the solver was still generating implications for *false_lit_max_dl* decision level, then this clause would have become a unit clause and generated an implication. Because generating implications as early as possible is very important for directing the search, we allow the solver to backtrack in order to make use of this implication. In this case, the solver backtracks to decision level *false_lit_max_dl* and the clause is inserted to the clause database. After the solver backtracks to *false_lit_max_dl* decision level, the same previous speculative decision at this level is put in temporary queue.

- If the clause has only one *true* literal and no *unknown* literals, then if *false_lit_max_dl* is smaller than *true_lit_dl* then this is indeed an implication. This restriction is necessary because there might be cases where the clause has only one *true* variable but it does not represent an implication. In such cases the *true* variable was set at a level while some of the remaining literals were *unknown* but are now set to *false*. The solver proceeds by backtracking to *false_lit_max_dl* and queuing an implication in the same way as the previous case.

- If the clause has all its literals set to *false*, then the clause has resulted in a conflict. In fact if this clause was available when decision level *false_lit_max_dl* was still being populated by implications then this clause would have caused a conflict at this level. This conflict would have helped direct the search, if detected. Thus the solver backtracks to make use of this conflict. However, if the conflict is at the first decision level, then this situation is the same as the fourth case in the previous merging algorithm mentioned above. Therefore the sub-problem is unsatisfiable. If the conflict is at a higher level, then the solver backtracks to *false_lit_max_dl*. Also previous decision at this level is saved in a temporary queue in the same way as the previous two cases.

- If none of the above cases apply then the clause is added immediately to the clause database without altering the decision stack.

When a new clause is merged, the decision stack is modified and a backtrack is performed in three of the five cases presented above. In addition, every backtrack reduces the stack depth, unless the top level is reached. When the stack depth is reduced the implication queue is cleared before any new implications are added. Also the decision level from which the solver will start is also cleared so that the solver can reconstruct the resulting implications while taking the new clauses into consideration. When the solver backtrack to the first level in the decision stack, the new merging method becomes the same as the simpler previous method.

The effect of backtracking to a higher level in the decision stack helps the solver investigate a more relevant part of the search space due to the newly found implication or conflict. The merging of shared clauses from other solvers restricts the search space and prevents the solver from wastefully revisiting some parts of the search space. Merging new clauses has an effect similar to randomization. Randomization [32] is a process where the decision stack is cleared after a timeout period and then starts at another random location in the search space. The hope is that the restart will lead to another location in the search space which will result in solving the problem faster. Randomization is used by most solvers and has been shown to improve solver performance. By merging new clauses, more relevant search spaces are chosen based on new implications and conflicts and not by random chance.

As described in [46], the exact effect of sharing clauses is not yet known. In addition, when a large number of clients are sharing even a small number of clauses the total communication overhead becomes significant. Shared clauses could be

streaming into the solver at sometimes high rates, especially if the number of processes used is high. Therefore merging the clauses immediately will cause frequent preemption of the solver. When the solver is preempted it stops until the received clauses are merged. The next method we present is designed to mitigate this problem.

There is a chance that some of the newly merged clauses which are added to the clause database can be duplicates of other previously existing clauses. Only clauses which do not result in implications or conflicts can be duplicates. Duplicate clauses will waste valuable memory space. Checking each new clause received by a solver to insure that it is not a duplicate before adding it to the database is computationally expensive. It requires scanning the entire database and comparing the new clause with every clause in the database. However, since GridSAT broadcasts clauses immediately after they are learned then all solvers are aware of the new clause quickly. Once a solver has a copy of the clause in its database it will not re-learn it. Therefore there is a slim chance that duplicate clauses will become an overwhelming problem. In future work, we will instrument GridSAT to find out how much duplication really occurs for a given set of problems.

3.2.3 The Periodic Method for Sharing Clauses:

The two previous methods represent two opposite extremes in how long a shared clause is delayed before it is merged with the local clause database. On the one hand, the lazy method delays delays the merging for a long time which potentially hinders the effectiveness of clause sharing and wastes valuable memory space. On the other hand, the immediate method merges clauses as soon as they are received which leads to frequent interruption of the local solver.

The *periodic method* is designed to tune the periodicity of clause merging. When using this method the local solver merges the clauses at periodic intervals. The periodicity is determined by a user specified parameter. The periodic method allows the solver to merge received clauses more frequently than the lazy method. Also the periodic method merges clauses in batches and interrupts the local solver less frequently than the immediate method of sharing clauses. In this paper we set the periodicity to 60 seconds.

3.2.4 Dynamically Adjusting Size of Shared Clauses:

GridSAT clients only share "short" clauses in order to minimize communication cost. Short clauses are expected to have a higher impact on pruning the search space and are more probable to generate implications. In fact the pruning effect of a clause is inversely proportional to its size (i.e. number of its literals). Previous GridSAT implementations take the maximum length of shared clauses as a static parameter.

Using a static value for determining the maximal size of shared clauses, may lead to one of two possible bad scenarios. First, if the value is too small the processes will not generate clauses smaller than the suggested value and no clause sharing will happen. In the second scenario, the used maximal clause size is low causing a large number of clauses to be shared. As a result, an influx of learned clauses may overwhelm the solvers with unnecessary communication and computational overhead. In addition, it is hard to determine *a priory* what the maximal clause size should be for a given SAT instance. In order to avoid both of these scenarios, the maximal clause size can be varied during the application execution using a given problem instance.

In the current implementation of GridSAT, the maximal size of shared clauses is determined dynamically. We set the absolute minimum for the maximal size to two. The maximal size of learned clauses is adjusted depending on a user supplied maximum rate of communication overhead due to clause sharing. The user can supply a maximal rate for shared clauses or use the default (set to 3). A process monitors the rate of shared clauses and calculates it periodically every five minutes. When this process notices that the maximal rate was exceeded, it broadcasts immediately an incremental decrease of the maximal clause size. This step insures that communication overhead resulting from shared clauses will only exceed its maximum for a short period of time. If the rate is below the maximal rate, then the monitoring process waits for half an hour before increasing the maximal rate and broadcasting the new value to the rest of the solvers. This allows the communication overhead to remain under its maximum value for a long time period. The user can also set an absolute maximal size for shared clauses.

4 GridSAT with Solution Enumeration

When a satisfiability problem is satisfiable, it may have more than one solution. However, the basic DPLL algorithm is designed to terminate after the first solution is determined. There are multiple ways to alter a DPLL-based sequential solver in-order to enable solution enumeration. For example, a simple approach would be to augment the initial set of clauses with a clause for every solution encountered. The addition of each solution clause would prevent the solver from generating the same solution in later steps. Such clauses are usually long because satisfiable instances often include most of the variables. In addition, these clauses will not be deleted in the future unlike learned clauses which are dispensable. A major drawback of such a solution is the need to use more memory to store an additional clause for each new solution. As the number of solutions is usually high, the memory needed to store the clauses produced by solutions becomes very large. This makes to solver less efficient as less memory

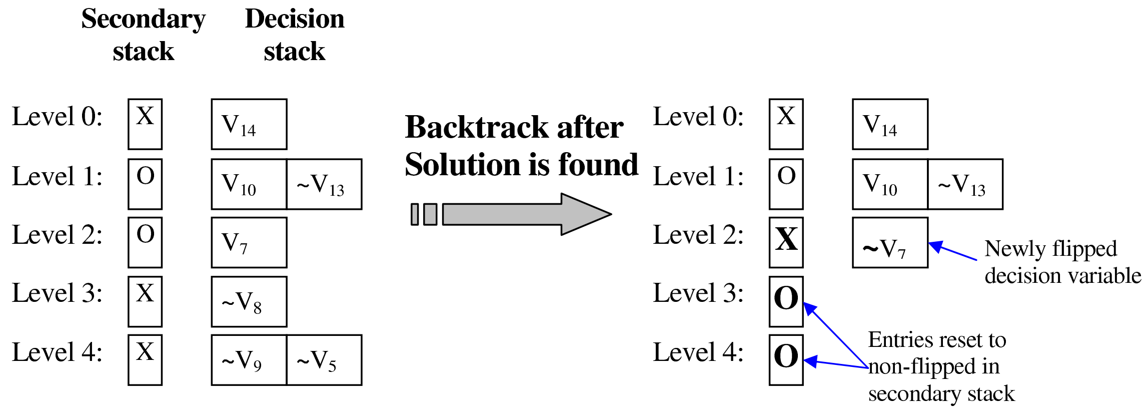

Figure 3: Example decision and secondary stack modifications after a solution is encountered. In the secondary stack *X* and *O* stand for flipped and not flipped respectively.

is available for clauses obtained through learning.

GridSAT with enumeration uses a different approach which uses little memory overhead which does not increase with the number of solutions found. This method could easily be integrated to other DPLL-based solvers and is not specific to zChaff. In the next section, we describe this method as it applies to a sequential solver. We will also present the ramifications of this method for parallel solvers such as GridSAT.

4.1 Solution Enumeration with Modified DPLL

The DPLL-based solver with solution enumeration we present, uses an additional stack in order to prevent itself from reproducing identical solutions. After each solution is found, the solver continues search for satisfiable solutions by moving to a different part of the search space. The solver terminates when the search space is exhausted and the last portion of the search space is determined to have no solutions (i.e. unsatisfiable).

The enumeration solver in GridSAT uses an additional stack called, the *secondary stack*, to track the state of each decision variable on each level of the decision stack. Thus the size of the secondary stack is equal to the size of the decision stack and is at most the number of variables in the original SAT problem. A variable on the secondary stack could be either *flipped* or *not flipped*. In the following description we also use the equivalent pair of terms *inverted* and *non-inverted*

Initially all entries on the secondary stack are marked as not flipped. When a solution is found the current variable assignments are saved in a repository external to the solver process. After the solution is saved then both stacks are modified. First, the decision variable in the highest level of the decision stack is inverted. Second, the corresponding entry in the secondary stack is marked as *flipped*. The solver continues by clearing the highest decision level and assuming the inverted value of the same decision variable. After both stacks are updated, the sequential solver proceeds as usual by making more speculative decisions, augmenting the decision stack and backtracking when a conflict is encountered leading to a reduction in the decision stack. When the decision stack shrinks because of backtracking all states in the secondary stack above the current decision level are cleared and marked as non-inverted.

When a solution is encountered and the current decision level in the secondary stack is marked as inverted, the solver proceeds by removing the highest decision levels and backtracking to a level where the decision variable on the secondary stack is not flipped. When such a level is found before reaching the second decision level, the solver marks that level as flipped and continues by assuming the inverted value at the same decision level. For example, in figure 3 we suppose that a solution was encountered at level four as shown on the left side of the figure. The right hand side shows how the decision and secondary stacks are modified. The solver backtracks to level 2 since it is the first non-flipped entry encountered on the secondary stack. This entry is marked with an **O** (not-flipped) before the solution is found. The same entry is marked with an **X** (flipped) after the solution is found. Notice that entries in the secondary stack (level 3 and 4) below the new decision level (level 2) are cleared and marked as not-flipped after the solution is found. Also the variable at the new decision level V_7 is flipped in the modified decision stack to $\sim V_7$. After updating both stacks, the solver then proceeds to explore the rest of the search space.

If the solver backtracks to the second decision level, then the solver has finished sweeping the branch of the search space which assumes the current value of the decision variable at the this level. Therefore, the solver can assume the opposite value of this variable for the remaining search space. Thus, the solver backtracks to the first decision level. and augments

this level with the inverted value of the decision variable previously found at the second decision level. The solver than proceeds by searching for implications produced by the newly assumed values.

The secondary stack is used as an additional mechanism to restrict the search space after a solution is found. In addition, the new extension to the basic algorithm does not restrict the efficient sequential algorithm in any fashion. The solver continues to navigate and prune the search space as before. The role of the secondary stack is to prevent the solver from reproducing the same solutions unnecessarily.

4.2 Parallel Solver with Solution Enumeration

Deploying parallel solvers with solution enumeration requires only modifications to the sequential solver. The role of the parallel solver infrastructure is to collect all the solutions in a repository.

In the parallel version, described in section 3.1, which uses the basic DPLL algorithm, each solver is given an initial decision stack and a clause database. In the parallel solver which uses solution enumeration, each client is given the same decision stack and an additional secondary stack. In order to illustrate how the secondary stack is split in the case of a parallel solver with solution enumeration, we use the same example as in figure 2. Before splitting, client A has both a secondary stack and a decision stack. After splitting, both client A and B receive the same decision stack and clause databases as described in section 3.1. Also, the old client (A) will receive the same original secondary stack except that level one is deleted. The new client (B), however, receives a totally blank secondary stack.

Now we show that this method of solution enumeration will not produce redundant solutions. In a parallel solver, each client starts from a distinct initial decision stack as described earlier. If any client finds a solution the initial decision stack will be a subset of that solution. Therefore, it is guaranteed that no two clients will produce the same solution since all clients start from distinct initial decision stacks. Furthermore, no client will produce the same solution more than once because the decision stack is different for each iteration of the DPLL algorithm. Therefore, the above algorithm will produce the a set of distinct solutions.

5 GridSAT Architecture and Resource Scheduling

The design of the GridSAT application has three main goals. The first goal is to allow for an efficient parallel SAT solver which adjusts to the variable resource usage of the problem being solved. The second goal is to use the available resources efficiently. The final goal is to make GridSAT adapt to variations in the availability and composition of the resource pool.

GridSAT is implemented as a special form of the master/client model where individual clients communicate directly and share clauses. The master consists of four main components: the resource manager, the client manager, the scheduler and the checkpoint server. A general architecture of the master process is shown is figure 4. External components with which the master interacts are shown as "clouds."

The resource manager loads resource information from one or more Grid information systems such as Globus MDS [12] and the NWS [58, 47, 59]. The scheduler as shown in figure 4 is responsible for coordinating the interactions between all the components. In addition it handles interactions with external resources and monitors them to detect failures. The resource manager is aware of the different types of resources. For shared resources only one GridSAT process per host is launched. For batch systems, the resource manager launches one job at the start of the execution. Additional, jobs could be manually submitted and GridSAT will use their resources when they become available. Actually the client manager will accept any additional clients launched from newly available resources or previously submitted batch jobs. It is the role of the client manager to maintain the list of active clients and monitor their progress.

The GridSAT scheduler is the focal point and is responsible for coordinating the rest of the components. It is also responsible for launching the clients. The scheduler uses a progressive scheme for acquiring resources and adding them to the resource pool. Also 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 two 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 and a smaller resource set will be used, and excess resources will be released.

A typical execution will start by launching the master. The master will examine the problem to find any obvious variable assignments and remove any inconsequential clauses. Some problems might be solved at this stage because of an easily detectable conflict. After this stage, the master requests the

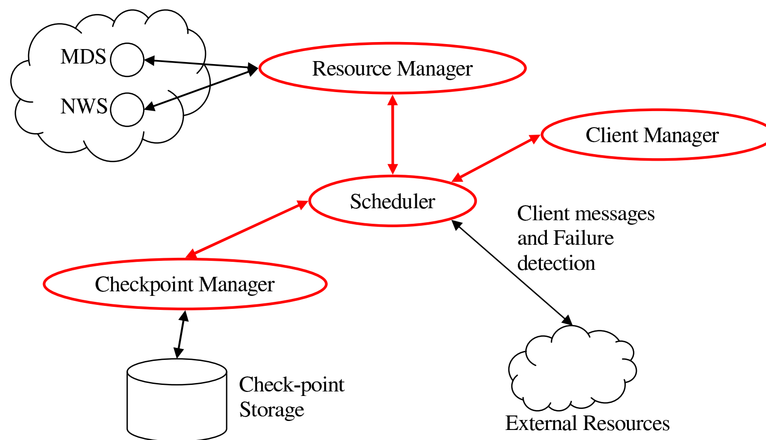

Figure 4: GridSAT components and their internal and external interactions. The external components and systems which GridSAT uses, such as the Globus MDS and the NWS, are shown in clouds.

resource list available from deployed Grid services such as the MDS [12] and NWS [58] or simply a configuration file. The scheduler immediately submits any batch jobs to their respective queues. When a remote client starts running it contacts the client manager and registers with it. The scheduler ranks the set of available clients based on their processing power and available memory as provided by the NWS [58, 47]. Static values for these resource parameters can be used when GridSAT is configured without NWS or the Globus MDS.

The GridSAT scheduler uses the first available client to immediately 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 when it wants to split its assigned subproblem with another client because its memory usage exceeded a certain limit or after running for a specified period of time. This time period is determined as two times 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. 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 three components: the master, the splitting client and an *idle* client. The *idle* client is a process which was not previously assigned a sub-problem to investigate. Figure 5 shows the steps taken during the splitting process. Client A which has presumably been solving a sub-problem, has detected that it needs to split its search space. Client A, then notifies the master using message (1). Upon receiving this message the master selects the highest ranked client and includes it in message (2) which it sends to client A. Using the information in message (2) client A determines which of its peers it will split the problem with. Client A then proceeds to communicate directly with client B by sending it message (3). In previous GridSAT implementations, message (3) is very large and varies in size from 10 KB to 500 MB. By using direct peer-to-peer communication the overall communication overhead is reduced. When the splitting is successfully completed, both clients alert the master using messages (4) and (5). In Message (4), client A sends new stacks for both clients A and B. Each stack is used as a checkpoint for its respective client. Both messages are used so that GridSAT can recover gracefully if one or both clients fail during the splitting procedure. Also if only one of the clients fail, then only that client will be restarted because the acknowledgements (4) and (5) are received separately.

Message (3) above allows the transfer of a newly created subproblem to the idle client. This message is the largest message and contains three different parts:

- The assignment stack: It is the smallest part and is in the order of the number of variables.
- The set of original problem clauses: This could be as large as the initial problem file
- The database of learned clause: It is the largest component and is 100s of Mega-Bytes in size.

5.1 Reducing Communication Overhead

GridSAT reduces the communication overhead of the solver in two ways. First, problem files are copied only once where several hosts share a common file system. Therefore split messages to the same set of hosts will be smaller since it will not include the second part of message (3) mentioned above. The second modification makes it possible for the new client to proceed with its computations immediately after it receives

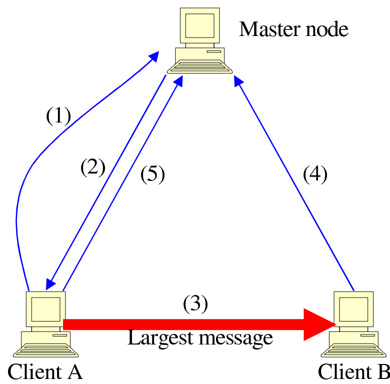

Figure 5: Communication scenario of splitting the subproblem assigned to client A with client B

the assignment stack and load the problem file from the shared file system. Since learned clauses contain redundant information, then they are not required to start solving the new subproblem. Therefore they are sent in a separate message. This message takes a long time to transfer, and the new clauses will be merged as they are received using the clause merging algorithms mentioned above. Using this methods the new client (client B in figure 5) will not have to idly wait for the entire message to arrive before starting solving the newly assigned sub-problem. The old client (client A in figure 5) still waits because the clause database is locked until the transfer is completed. Making an additional copy of the clause database in order to prevent the old client from stalling is not practical because the size of the clause database is very large and there is not sufficient memory to hold a separate copy. The old client waits and does not proceed until the clause database destined for the new client is transferred. Transferring these clauses to the new client is essential to the efficiency of the solver. Eliminating this transfer would slow the solver significantly.

Finally the GridSAT solver terminates when all sub-problems have been solved or one the clients finds a satisfying assignment. In the latter case the client which finds the satisfying assignment sends its stack to the master. The master verifies that the set of truth assignments it received does indeed satisfy all clauses in the initial problem. Most solvers in the literature are evaluated based on the time the first satisfiable instance is found. But there are cases [24] where knowing all satisfiable instances is helpful. GridSAT can also enumerate all the instances where a problem is satisfiable. In all cases, when the master determines that the problem is solved, it sends a message to all clients requesting them to terminate.

5.2 Failure Recovery and Check-pointing System

In a computational grid environment resource failures are frequent. Therefore a grid application has to be able to recover from such failures. There are two components to failure recovery. First, the grid application should establish a mechanism for detecting the failure of remote components. GridSAT uses heartbeat messages to decide when a remote solver has failed. Second, the grid application should be able to restart with minimal work loss when failures occur. The current version of GridSAT uses check-pointing to recover from such failures. Each checkpoint belongs to one remote solver and represents a SAT sub-problem that can be restarted when combined with the initial SAT problem. GridSAT can use two types of checkpoints:

- Light checkpoints: This method requires little storage space and communication overhead. Only the top level of the assignment is recorded for each client. In this case checkpoints for a client will be updated only when more variables are added to the first decision level.
- Heavy checkpoints: In addition to the light checkpoint data, we save all newly learned clauses. It is also possible to save the top levels of the decision stack in order to reconstruct the exact decision levels after restart. This type of checkpoints can be saved at regular time intervals in addition to the instances when the top level is augmented.

The master stores and updates the checkpoints as they are received from the clients. The checkpoints can be stored either on a local file system or in a distributed fashion using IBP [34]. Idle clients are assigned new sub-problems either through splitting or from saved checkpoints. Sometimes the number of checkpoints exceeds the number of active clients. This happens when a large number of previously active clients terminate leaving behind their checkpoints. In this case the scheduler keeps a list of checkpoints and assigns them to newly created clients or those that have just finished solving their own sub-problem. Idle clients are assigned problems through splitting only after all checkpointed sub-problems are assigned to active clients.

When the master fails, GridSAT can recover by simply re-instantiating the master process on another machine if necessary. If checkpoints are available, the new master process can use them to recover pre-failure state. Also a user could cause an *intentional* failure by halting the master while it is solving a problem in order to start another problem for example. The user can later resume solving the previous problem using the saved set of checkpoints.

5.3 Efficient Use of Batch Jobs

Batch schedulers are usually used to control super-computing facilities [6, 49] and collections of grid resources such as Condor [10, 48]. Users in these environments are given a budget (i.e. a quota of CPU-hours) to use. Since this is valuable time, it is important from the user's perspective to use it effectively.

The scheduler bills the user and deducts from his budget the total time the nodes in the batch job are assigned to his job. The user is billed for the time used and not the time he initially requested. Thus if a job terminates early the user is only billed for the time during which his job actually ran. From a user's perspective, the goal is to minimize the cumulative idle time for all nodes during a batch job execution.

In traditional parallel applications, which mostly use MPI [18], the number of processes spawned is sufficient to insure that all nodes have a slice of the work assigned to them during the entire duration of the execution. All nodes start and stop execution simultaneously. This scenario leads to an efficient use of the batch jobs. GridSAT is not a traditional parallel application. In the case of GridSAT, the number of jobs (i.e. sub-problems) varies during execution. Actually, when a new large batch job becomes available the number of workers might be much larger than the number of available sub-problems. The goal of GridSAT is to make good use of the newly available and valuable processing power. It is possible to immediately split a sufficient number of sub-problems. This will lead to more efficient use of batch jobs but may affect negatively the solver's performance. If GridSAT, however, waits till enough problems split to populate all the batch nodes, it may lead to an inefficient use of super-computing nodes.

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 than it most probably is a hard problem. Thus batch jobs are only used when the problem is hard. When a batch job starts execution, GridSAT uses problem migration 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. Actually the number of active nodes (i.e. those with sub-problems) will increase exponentially. This happens because the number of new sub-problems is increased in proportion to the number of existing active solvers. Therefore, problem migration leads to a more efficient use of batch jobs.

5.4 Multiple Site Scheduling and Migration

GridSAT processes communicate as peers during problem splitting. Even after the implementation of the optimization presented above which reduce communication overhead, peer-to-peer messages are still the largest. Therefore, more efficient problem splitting will improve the overall solver's efficiency. More efficient problem splitting could be accomplished when clients belong to a pool of well connected resources. Such pools of resources are usually presented when new batch jobs reach the head of their waiting queue and start running. GridSAT migrates problems from dispersed nodes to processes which are part of a batch job.

The scheduler identifies batch processes in a static fashion using their host names. Instead of creating a new sub-problems through splitting with a remote node, the scheduler requests the remote node to migrate to one of the clients in the batch job. Migration allows future splitting to happen between peers belonging to the same batch jobs. This leads to reduction in overall communication overhead. In future versions of GridSAT, determining when migration happens will be achieved through a more dynamic approach.

6 Experimental Apparatus and Results

In the experimental section we present three experimental sets. The first set of experiments allows the selection of the best strategy for merging learned clauses by comparing the effectiveness of those strategies as used in a parallel SAT solver. The second set of experiments compares the performance of GridSAT to that of the sequential solver zChaff. The third set of result show how GridSAT is used to coordinate a large pool of resource for extended time periods to solve "hard" satisfiability problems.

6.1 Learning Methods Experiments

In these set of experiments we study the effectiveness of the three different learning methods: the lazy method, the immediate method and the periodic method. The experiments are conducted using the set of 33 benchmark problems used by the different satisfiability competitions [37, 40] and previous evaluations of GridSAT [9]. The experiments were conducted on set of dedicated nodes on a cluster available at UCSB. Each experiment uses ten nodes and one of the three methods. The experiments are grouped into three sets where the maximal size of a shared clause is varied between 5, 10 and 15. The cluster nodes are Pentium V CPUs with 2.66 GHz frequency and 2 GB of memory. During these experiments automatic adjustment of shared clause size was disabled.

6.1.1 Results for Comparing Different Learning Methods: The experimental results are shown in table 1. This table shows experimental results for using a maximal size of shared clauses of 5, 10 and 15 respectively. The table contains three section, one for each of size of shared clauses used. Each section shows to total time for each of the three methods and the relative speed-up compared to the lazy method.

In order to save space the runtimes for the individual problems were not included. But from inspecting each of the three

Method	**Lazy**	**Immediate**	**Periodic**
Maximum size of shared clause = 5			
Total	76776	68620	64675
% Speedup	*(base)*	10.6	15.8
Maximum size of shared clause = 10			
Total	71860	67292	63400
% Speedup	*(base)*	6.4	11.8
Maximum size of shared clause = 15			
Total	69527	67292	63400
% Speedup	*(base)*	3.2	8.8

Table 1: GridSAT results comparing all three learning methods with maximal learn clause size equal to 5, 10 and 15.

experimental sets, we learned that no particular method outperformed the other two methods all the time. Instead each method outperforms the other two methods for a subset of the problems. We use the total runtime of all the problems to compare the efficiency of the methods. Using the total runtime for all problems in a benchmark is the standard method for comparing solvers.

We notice that in each case both the immediate and periodic methods outperform the lazy method. The immediate method outperforms the lazy method by an average of about 7%. The periodic method was the most efficient and showed a speedup of about 12% on average compared to the lazy method. We also notice that the speedup decreased as the size of maximal shared clause increased. These experiments show that using the periodic method gives the best overall performance.

6.2 Comparing GridSAT and the Sequential Solver

In the first set of experiments, we used 34 machines from the GrADS testbed and an additional machine (that we could completely instrument) as a master node. The machines were distributed among three sites: two clusters (separated by campus networking) at the University of TN, Knoxville (UTK), two clusters at the University of Illinois, Urbana-Champaign (UIUC) and 8 desktop machines at the University of San Diego (UCSD). The master node was also at UCSD. The machines had varying hardware and software configurations, with one of the UTK clusters having the best hardware configuration. For each zChaff (single machine) test we used a dedicated node from this cluster.

As a set of test applications, we chose a suite of challenge problems used to judge the performance of automatic SAT solvers at the SAT2002 conference [37]. 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* [38]. The solvable category contains problem instances that SAT solvers have been known to solve correctly. They are useful for comparing the speed of competitive solvers since it is likely that each solver in the competition will be able to generate an answer when the competition is held. 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. Of these problems, many have solutions that are known through analytical methods, but several are open questions in the field of satisfiability research.

In these experiments the maximum size of learned clauses shared is 10. Learned clauses bigger than 10 are not shared. This size allows for sharing of important clauses which would have maximal effect without increasing significantly the overhead of clause sharing. Also the time out for clients to request that their problems be partitioned is set to 100 seconds. For the solvable problems we set an overall maximum execution time out to a total of 6000 seconds for GridSAT. That is, if the entire problem is not solved in 6000 seconds, the application gives up and terminates without a definitive answer. For the challenging benchmarks, we double the overall time out to 12000 seconds.

In all of the experiments, we compare GridSAT to zChaff running in dedicated mode on the fastest processor to which we have access with an 18000 second total time out. For the challenging set we used 12000 seconds as the timeout value. Note that in the actual 2002 competition, using faster machines than the fastest we had available, zChaff was only able to complete a few instances from this set using a six-hour (21600 second) time out. Thus we believe that the comparison between the two using the machines in the GrADS testbed offer useful insight into the additional capability provided by GridSAT.

6.2.1 Results: The first set of results are presented in Table 2. The second column contains the solution to the instance: satisfiable(SAT), unsatisfiable(UNSAT), or unknown. We have marked those problem instances that 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 last column shows the maximum number of active clients during the execution of an instance. For all instances this number starts at one and varies during the run. The maximum it could reach is 34, the number of hosts in the testbed, but the scheduler may choose to use only a subset. This column records the maximum that the scheduler chose during each particular run. When a problem is solved the number of active clients collapses to zero. Speedup is measured as the ratio of the fastest sequential execution time of zChaff (on the fastest, dedicated machine) to the time recorded by GridSAT.

The problem instances in Table 1 are split into three categories. The first section represents the set of instances which were solved by both zChaff and GridSAT (taken from both the solvable and challenging categories of the SAT2002 benchmark suite since zChaff was able to solve some of the latter). On the small instances (ones that complete in less than 300 seconds) where communication costs are significant we notice that zChaff running on a single machine outperforms GridSAT. The slowdown however is not very significant because the actual time is short. For instances with long running times GridSAT shows a wide range of speed-ups ranging from almost none to almost 20 for *dp12s12*. Because GridSAT was using more machines it was capable of covering much more of the search space even when the run times were comparable. In only one relatively long running instance, *grid_10_20*, did GridSAT show a slowdown. The maximum number of active clients for the entire problem only reached a maximum of twelve during its execution. With this little sharing, parallelism did not seem to improve performance. This particular problem comes from a non-realizable circuit design illustrating the data-dependent nature of SAT solver performance results.

The second category of problems represent those that GridSAT was able to solve while zChaff either timed-out or ran out of memory. In addition, only three out of the ten problems in this category were solved by another solver during the SAT2002 competition [39]. Note that zChaff was crowned the overall winner because of its cumulative performance across benchmarks. Individual instances may have been better solved by particular solvers, but because the competition attempts to identify the best general method, aggregate time is used, and zChaff is the best on aggregate.

The rest of the seven instances in this second category have only been solved by GridSAT to the best of our knowledge. Three of the solved instances were part of the challenging benchmark for which results were originally unknown constituting new domain science in the field of satisfiability. The other four had known analytical answers, but no automatic generalized solver had been able to correctly generate them indicating the additional solution power that a Grid implementation brings to the field.

These results show that GridSAT provides a speedup compared to existing sequential solvers. This speed up is not linear with respect to the number of resources used because the DPLL algorithm used to solve SAT instances is a branch-and-bound algorithm. In such search based algorithms the time to solution is not always proportionally related to the number of times the search space is divided. For example, dividing the search space in half may not cause a two-fold speedup in time to solution. In fact, the two sub-problems may have very different times to solution. In this case, even though there will be an overall improvement by using two resources, the speedup may be less than two. Actually there is no theoretical guarantee that dividing the search space will result in speedup because of all the heuristics involved. In practice, however, partitioning the search space causes performance improvements most of the time. The contribution of GridSAT is not only to provide speedup over sequential solvers but also to enable the solution of problems that were previously unsolved as shown by the next set of experiments.

6.3 Solving Hard Problems

Since GridSAT is a true grid application 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 experiment we simultaneously use computational resources which belong to collections of individual machines, small size research clusters and super-computing scale clusters. The computational resources we used are composed from four main sources:

- VGrADS [55] testbed with additional machines from the University of California, Santa Barbara (UCSB)
- Blue Horizon [6] located at the San Diego Supercomputer Center (SDSC)
- TeraGrid [54] site at the San Diego Super Computing Center (SDSC)
- TeraGrid [53] site at the National Center for Supercomputing Applications (NCSA)
- Data Star [13] supercomputer at SDSC

The TeraGrid [49] project is a national scale project which is aimed at building the worlds largest distributed infrastructure for open scientific research. It includes five sites at SDSC [54],NCSA [53], Argonne National Laboratory [50], Pittsburgh Super Computing center [52] and Caltech CACR [51]. Additional sites and resources are planned at Oak Ridge National Lab (ORNL); Purdue University, Indiana; Indiana University, Bloomington; and the Texas Advanced Computing Center (TACC) at The University of Texas at Austin.

The **V**irtual **Gr**id **A**pplication **D**evelopment **S**oftware (GrADS) project [55], a continuation of GrADS [4, 21] is a comprehensive research effort studying Grid programming tools and application development. GrADS includes a set of programming tools for managing grid applications using performance models. Scheduling applications in GrADS uses compiler pre-processing of the programs and introduced instrumentation combined with NWS data. The tools GrADS uses are included in a software package termed *GrADSoft*. To facilitate experimental application research and testing, the

File name	SAT/UNSAT/ UNKNOWN	zChaff (sec)	GridSAT (sec)	Speed-Up	Max # of clients
Problem solved by zChaff and GridSAT					
6pipe.cnf	UNSAT	6322	4877	1.23	34
avg-checker-5-34.cnf	UNSAT	1222	1107	1.10	9
bart15.cnf	SAT	5507	673	8.18	34
cache_05.cnf	SAT	1730	1565	1.11	34
cnt09.cnf	SAT	3651	1610	2.27	12
dp12s12.cnf	SAT	10587	532	19.90	8
homer11.cnf	UNSAT	2545	1794	1.42	10
homer12.cnf	UNSAT	14250	4400	3.24	33
ip38.cnf	UNSAT	4794	1278	3.75	11
rand_net50-60-5.cnf	UNSAT	16242	1725	9.42	20
vda_gr_rcs_w8.cnf	SAT	1427	681	2.10	15
w08_14.cnf	SAT	14449	1906	7.58	34
w10_75.cnf	SAT	506	252	2.01	2
Urquhart-s3-b1.cnf	UNSAT	529	526	1.01	4
ezfact48_5.cnf	UNSAT	127	196	0.65	1
glassy-sat-sel_N210_n.cnf	SAT	7	68	0.10	1
grid_10_20.cnf	UNSAT	967	3165	0.31	12
hanoi5.cnf	SAT	2961	1852	1.60	33
hanoi6_fast.cnf	SAT	1116	831	1.34	4
lisa20_1_a.cnf	SAT	181	243	0.75	2
lisa21_3_a.cnf	SAT	1792	337	5.32	4
pyhala-braun-sat-30-4-02.cnf	SAT	18	84	0.21	1
qg2-8.cnf	SAT	180	224	0.80	2
Problems solved by GridSAT only					
7pipe_bug.cnf	SAT	TIME_OUT	5058	–	34
dp10u09.cnf	UNSAT	TIME_OUT	2566	–	26
rand_net40-60-10.cnf	UNSAT	TIME_OUT	1690	–	30
f2clk_40.cnf	UNSAT(*)	TIME_OUT	3304	–	23
Mat26.cnf	UNSAT	MEM_OUT	1886	–	21
7pipe.cnf	UNSAT	MEM_OUT	6673	–	34
comb2.cnf	UNSAT(*)	MEM_OUT	9951	–	34
pyhala-braun-unsat-40-4-01.cnf	UNSAT	MEM_OUT	2425	–	34
pyhala-braun-unsat-40-4-02.cnf	UNSAT	MEM_OUT	2564	–	34
w08_15.cnf	SAT(*)	MEM_OUT	3141	–	34
Remaining problems					
comb1.cnf	*	TIME_OUT	TIME_OUT	–	34
par32-1-c.cnf	SAT	TIME_OUT	TIME_OUT	–	34
rand_net70-25-5.cnf	UNSAT	TIME_OUT	TIME_OUT	–	34
sha1.cnf	SAT	TIME_OUT	TIME_OUT	–	34
3bitadd_31.cnf	UNSAT	TIME_OUT	TIME_OUT	–	34
cnt10.cnf	SAT	TIME_OUT	TIME_OUT	–	34
glassybp-v399-s499089820.cnf	SAT	TIME_OUT	TIME_OUT	–	34
hgen3-v300-s1766565160.cnf	*	TIME_OUT	TIME_OUT	–	34
hanoi6.cnf	SAT	TIME_OUT	TIME_OUT	–	34

(*): problem solution is unknown

Table 2: GridSAT and zChaff SAT2002 Benchmark Results on GrADS testbed

project maintains a nationally distributed grid of resources for use as a production testbed. The baseline Grid infrastructure is provided by Globus and the NWS, upon which is layered a set of programming abstractions. In this work we extend GridDSAT to use all resources that do not currently benefit from these sophisticated Grid programming tools. GridSAT components (i.e. master and client) use the EveryWare [56, 57] messaging system for communication.

During our experiments, none of the resources we used were dedicated to our use. The VGrADS testbed, the UCSB machines, and the super-computing resources were all in continuous use by various researchers and application scientists at the time of the experiment. 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. Sometimes we were requested to temporarily vacate some specific resources because some users wanted to run experiments without interference from other applications. We consider this to be a realistic scenario for Computational Grid computing, but it makes repeatable timings of similar problems (particularly those we ran for long periods) difficult. In particular, in batch controlled system such as Blue Horizon, Data Star and the TeraGrid, a user presents a request for a number of nodes and a maximum duration. After waiting in the job queue, the user's job runs with exclusive access to the nodes during execution, but the queue wait time incurred before execution begins is highly variable. However, the effect of resource contention is almost assuredly a performance-retarding one. 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.

In previous experiments [9] we showed how GridSAT can simultaneously use small clusters and a collection of lab machines in conjunction with high end supercomputers such as Blue Horizon. The experiments used a single job request on the Blue Horizon with a maximum timeout of 12 hours.

The set of experiments we present in this paper 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 conference [37] and SAT2003 benchmarks [40]. These benchmarks are used to judge and compare the performance of automatic SAT solvers at the SAT2002 [39] and SAT2003 [42] conferences. 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* [38, 41]. 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. We only chose problems which are hard so that we can demonstrate the ability of the GridSAT system to solve such challenging problems. These problems were deemed hard by all participating solvers.

We investigate seven previously unsolved problems divided as follows:

- 3 instances from the SAT 2003 benchmark category,
- 4 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 satisfiability 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 this set of experiments, the resource pool included 40 machines from the VGrADS testbed and an additional machine (that we could completely instrument) as a master node. The machines were distributed among three sites: three clusters (separated by campus networking) at the University of TN, Knoxville (UTK), five desktop machines at the University of San Diego (UCSD) and ten machines from the MAYHEM [31] lab at the University of California, Santa Barbara. An additional node, designated the master node, was at UCSB. The machines had varying hardware and software configurations.

In these experiments 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.

6.3.1 Results: Solving Hard Satisfiability Problems: The experimental results are summarized in Table 3. 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. In such cases, experiments could be continued using the saved checkpoints.

Table 3 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 still waiting in the queue. Thus when the problem got solved before they got to the head of the queue the batch jobs were canceled. On the other hand the second problem *k2fix-gr-rcs-w9.cnf* took much longer to solve, it took more than two weeks. We expect that some Grid applications will require running for such extended periods of time. Table 4 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 4. The number of processes per node is determined so that each process gets a minimum of 1/2 GByte or 1 GByte of memory. In table 5 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 [9] 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 6. The figure shows several jobs from Blue Horizon, Data Star and TeraGrid sites participating in the execution. Note that the processor count is represented in logarithmic scale. 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. In figure 7 we show the total number of processes used by GridSAT during the same period. 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). Floating point operations are only used to handle time related events. We present in figure 8 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 8 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 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 idly 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

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

Computational resource	Job count	Job duration(hours)	Number of nodes	processes/node
Blue Horizon	2	10	100	3
Blue Horizon	1	12	100	3
DataStar	2	10	8	11
TeraGrid @ SDSC	1	10	40	2
TeraGrid @ SDSC	1	12	40	2
TeraGrid @ SDSC	3	10	4	2
TeraGrid @ SDSC	4	5	4	2
TeraGrid @ NCSA	3	10	4	2
TeraGrid @ NCSA	4	5	4	2

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

Table 4: Batch jobs used to solve the k2fix-gr-rcs-w9.cnf instance from SAT 2003 benchmark

Computational resource	node-hours	CPUs/node	CPU-hours
Blue Horizon	3200	8	25600
Data Star	160	11	1760
TeraGrid @ SDSC	1080	2	2160
TeraGrid @ NCSA	200	2	400
GrADS testbed(*)	13750	1	13750

(*) machines were shared with other users

Table 5: CPU-hours per resource used to solve the k2fix-gr-rcs-w9.cnf instance from SAT 2003 benchmark

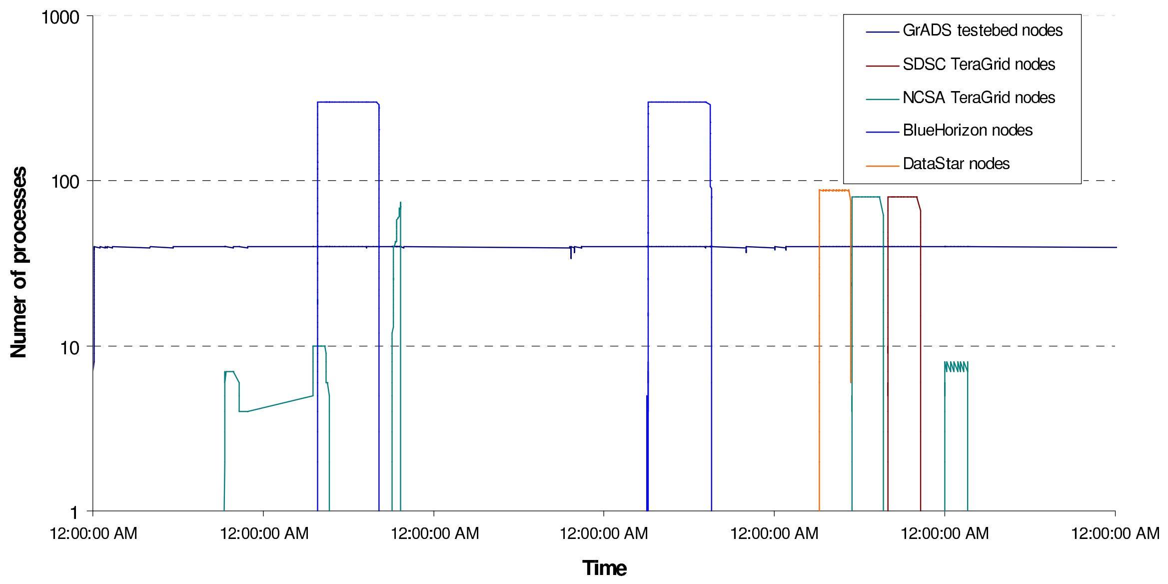

Figure 6: A six day snapshot representing GridSAT processor count usage from the different resources in logarithmic scale.

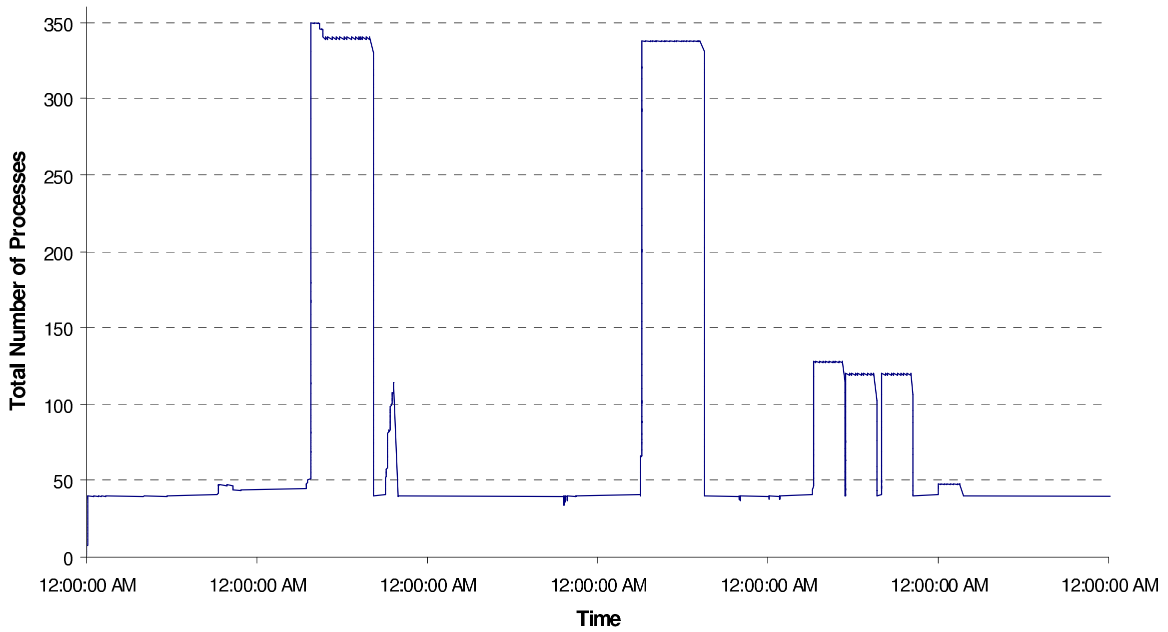

Figure 7: The total processor count usage from all the resources for the same six day snapshot shown in figure 6.

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 6, 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.

7 Conclusion

We have described GridSAT a distributed satisfiability solver for the computational grid. GridSAT is shown capable of running on a dynamic and heterogeneous set of resources. GridSAT was capable of solving previously unsolved problems. In order to solve even harder problems, new optimizations to both the algorithm and architecture of GridSAT were introduced. GridSAT is capable of merging newly received shared clauses immediately to the clause database to improve the solver's efficiency. Also communication overhead is reduced by selectively sending important information first and avoiding redundancy when possible. 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.

We also present a version of GridSAT capable of enumerating all solutions of a given satisfiability problem. Finally, a grid portal was developed to enable users to submit their specific problems to GridSAT running transparently on a set of computational resources.

References

[1] L. C. Alessandro Armando. Abstraction-driven sat-based analysis of security protocols. In *Theory and Applications of Satisfiability Testing, 6th International Conference, SAT 2003*, pages 257–271, May 2003.

[2] 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.

[3] C. W. B. Li and F. Somenzi. Abstraction refinement in symbolic model checking using satisfiability as the only decision procedure. December 2003.

[4] 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".

[5] A. Biere. http://www.inf.ethz.ch/personal/biere/projects/limmat/.

[6] BlueHorizon. http://www.npaci.edu/BlueHorizon/.

[7] R. Bjar and F. Many. Solving the Round Robin Problem Using Propositional Logic. AAAI/IAAI, 2000.

[8] W. Chrabakh and R. Wolski. GrADSAT: A Parallel SAT Solver for the Grid. Technical Report 2003-05, UCSB, March 2003.

[9] W. Chrabakh and R. Wolski. GridSAT: A chaff-based Distributed SAT solver for the Grid. In *Supercomputing Conference, Phoenix, AZ*, November 2003.

[10] Condor home page – http://www.cs.wisc.edu/condor/.

[11] S. A. Cook. The complexity of theorem-proving procedures. *Third Annual ACM Symposium on Theory of Computing*, 1971.

[12] 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.

[13] Data Star. http://www.npaci.edu/DataStar/.

[14] M. Davis, G. Logeman, and D. Loveland. A machine program for theory proving. Communications of the ACM, 1962.

[15] A. Downey. Predicting queue times on space-sharing parallel computers. In *Proceedings of the 11th International Parallel Processing Symposium*, April 1997.

[16] D. G. Feitelson and L. Rudolph. *Parallel Job Scheduling: Issues and Approaches*. Springer-Verlag, 1995.

[17] S. L. Forman and A. M. Segre. Nagsat: A randomized, complete, parallel solver for 3-sat. SAT2002, 2002.

[18] M. P. I. Forum. Mpi: A message-passing interface standard. Technical Report CS-94-230, University of Tennessee, Knoxville, 1994.

[19] I. Foster and C. Kesselman. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann Publishers, Inc., 1998.

[20] 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.

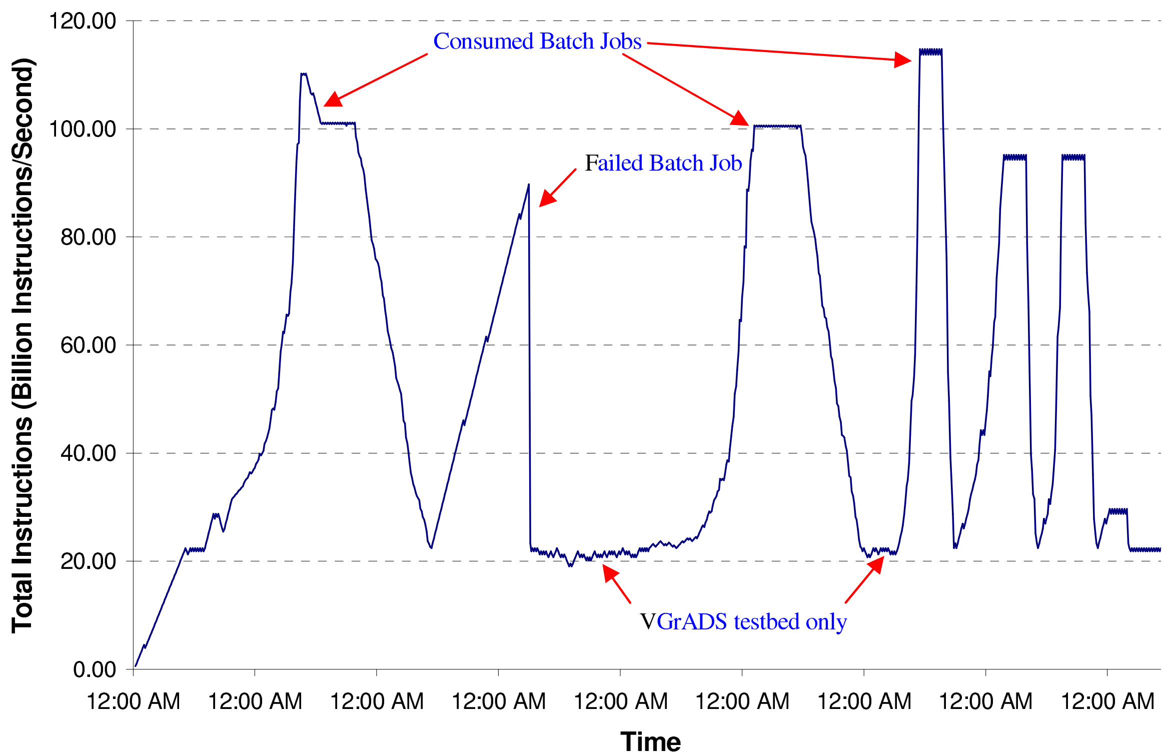

Figure 8: Estimation of Integer Operations per second usage for all resources during the same six day snapshot shown in figure 6.

[21] GrADS. `http://hipersoft.cs.rice.edu/grads`.

[22] 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.

[23] C. K. I. Foster and S. Tuecke. The anatomy of the grid. 2001.

[24] D. Jackson and M. Vaziri. Finding bugs with a constraint solver. *International Symposium on Software Testing and Analysis*, 2000.

[25] 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.

[26] H. Kautz and B. Selman. Planning as satisfiability. In *Proceedings of the Tenth European Conference on Artificial Intelligence*, pages 359–379, August 1992.

[27] W. Kunz and D. Stoffel. *Reasoning in Boolean Networks: Logic Synthesis and Verification Using Techniques*. Kluwer Academic Publishers, Boston, 1997.

[28] T. Larrabee. Efficient generation of test patterns using boolean difference. pages 795–802.

[29] T. Larrabee. Test pattern generation using boolean satisfiability. In *IEEE Transactions on Computer-Aided Design*, pages 4–15, January 1992.

[30] 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.

[31] MAYHEM home page – `http://pompone.cs.ucsb.edu//`.

[32] 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.

[33] 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.

[34] J. Plank, M. Beck, and W. Elwasif. IBP: The internet backplane protocol. Technical Report UT-CS-99-426, University of Tennessee, 1999.

[35] S. Reda and A. Salem. Combinational equivalence checking using boolean satisfiability and binary decision diagrams. In *Proceedings of the conference on Design, automation and test in Europe*, pages 122–126. IEEE Press, 2001.

[36] M. Russell, G. Allen, G. Daues, I. Foster, E. Seidel, J. Novotny, J. Shalf, and G. von Laszewski. The astrophysics simulation collaboratory: A science portal enabling community software development. *Cluster Computing*, 5(3):297–304, 2002.

[37] SAT 2002 benchmarks. `http://www.satlive.org/SATCompetition/2002/submittedbenchs.html`.

[38] SAT 2002 challenge benchmark. `http://www.ececs.uc.edu/sat2002/sat2002-challenges.tar.gz`.

[39] SAT 2002 Competition. `http://www.satlive.org/SATCompetition/`.

[40] SAT 2003 benchmarks. `http://satlive.org/SATCompetition/2003/`.

[41] SAT 2003 challenge benchmark. http://satlive.org/SATCompetition/2003/.
[42] SAT 2003 Competition. http://satlive.org/SATCompetition/2003/.
[43] 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.
[44] J. M. Silva and K. Sakallah. Grasp - a new search algorithm for satisfiability. ICCAD. IEEE Computer Society Press, 1996.
[45] J. P. M. Silva. Search Algorithms for Satisfiability Problems in Combinational Switching Circuits. Ph.D. Thesis, The University of Michigan, 1995.
[46] 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.
[47] M. Swany and R. Wolski. Building performance topologies for computational grids. In *Proceedings of Los Alamos Computer Science Institute (LACSI) Symposium, 2002*, October 2002.
[48] T. Tannenbaum and M. Litzkow. The condor distributed processing system. *Dr. Dobbs Journal*, February 1995.
[49] TeraGrid. http://www.teragrid.org/.
[50] TeraGrid at Argonne National Laboratory. http://www-unix.mcs.anl.gov/teragrid-anl/.
[51] TeraGrid at Caltech Center for Advanced Computing Research. http://www.cacr.caltech.edu/resources/teragrid/.
[52] TeraGrid at Pittsburgh Supercomuting Center. http://teragrid.psc.edu/.
[53] TeraGrid at the National Center for Scientific Applications. http://teragrid.ncsa.uiuc.edu/.
[54] TeraGrid at the San Diego Supercomputing Center. http://www.sdsc.edu/teragrid.
[55] VGrADS. http://hipersoft.cs.rice.edu/vgrads.
[56] R. Wolski, J. Brevik, C. Krintz, G. Obertelli, N. Spring, and A. Su. Running EveryWare on the computational grid. In *SC99 Conference on High-performance Computing Proceedings*, 1999.
[57] R. Wolski, J. Brevik, G. Obertelli, N. Spring, and A. Su. Writing programs that run everyware on the computational grid. *IEEE Transactions on Parallel and Distributed Systems*, 12(10), 2001. available from http://www.cs.ucsb.edu/~rich/publications/ev-results.ps.gz.
[58] R. Wolski, N. Spring, and J. Hayes. The network weather service: A distributed resource performance forecasting service for metacomputing. *Future Generation Computer Systems*, 1999.
[59] R. Wolski, N. Spring, and J. Hayes. Predicting the cpu availability of time-shared unix systems on the computational grid. In *Proc. 8th IEEE Symp. on High Performance Distributed Computing*, 1999.
[60] L. Zhang, C. F. Madigan, M. W. Moskewicz, and S. Malik. Efficient conflict driven learning in boolean satisfiability solver. In *ICCAD*, pages 279–285, 2001.