The Knowledge Engineering Review, Vol. 16:1, 2001, 85–99
Printed in the United Kingdom

Combining linear programming and satisfiability solving for resource planning

STEVEN A. WOLFMAN and DANIEL S. WELD

Department of Computer Science & Engineering, University of Washington, Box 352350, Seattle, WA 98195–2350, USA
E-mail: wolf@cs.washington.edu, weld@cs.washington.edu

Abstract

Compilation to Boolean satisfiability has become a powerful paradigm for solving artificial intelligence problems. However, domains that require metric reasoning cannot be compiled efficiently to satisfiability even if they would otherwise benefit from compilation. We address this problem by combining techniques from the artificial intelligence and operations research communities. In particular, we introduce the LCNF (Linear Conjunctive Normal Form) representation that combines propositional logic with metric constraints. We present LPSAT (Linear Programming plus SATisfiability), an engine that solves LCNF problems by interleaving calls to an incremental Simplex algorithm with systematic satisfaction methods. We explore several techniques for enhancing LPSAT's efficiency and expressive power by adjusting the interaction between the satisfiability and linear programming components of LPSAT. Next, we describe a compiler that converts metric resource planning problems into LCNF for processing by LPSAT. Finally, the experimental section of the paper explores several optimisations to LPSAT, including learning from constraint failure and randomised cutoffs.

1 Introduction

Recent advances in Boolean satisfiability (SAT)-solving technology have rendered large, previously intractable problems quickly solvable (Bayardo & Schrag, 1997; Cook & Mitchell, 1997; Crawford & Auton, 1993; Gomes *et al.*, 1998; Li & Anbulagan, 1997; Selman *et al.*, 1996). SAT-solving has become so successful that many other difficult tasks are being compiled into propositional form to be solved as SAT problems. For example, SAT-encoded solutions to graph-colouring, planning and circuit verification are among the fastest approaches to these problems (Kautz & Selman, 1996; Selman *et al.*, 1997). These SAT encodings succeed because they compile to a simple yet expressive target language and take advantage of rapidly progressing solution techniques.

However, many real-world tasks have a metric aspect. For instance, resource planning, temporal planning, scheduling and analog circuit verification problems all require reasoning about real-valued quantities. Unfortunately, metric constraints are difficult and in some cases impossible to express in SAT encodings. Hence a target language and solution system that could efficiently handle both metric constraints and propositional formulae would yield a powerful substrate for handling artificial intelligence problems.

This paper introduces a new language, LCNF (Linear Conjunctive Normal Form), which combines the expressive power of propositional logic with that of linear equalities and inequalities. This combination is still less expressive than an Integer Linear Programming (ILP) encoding, but it encourages a careful division of logical and metric reasoning and combines the strengths of constraint satisfaction and Linear Programming (LP). We argue that LCNF provides an ideal target language into which a compiler might translate tasks that combine logical and metric reasoning.

Figure 1 Data flow in the demonstration resource planning system; space precludes discussion of the grey components

Moreover, we present an architecture for solving LCNF problems that takes advantage of the powerful existing technologies in both of these areas. We describe the LPSAT (Linear Programming plus SATisfiability) LCNF solver, which implements this architecture. LPSAT combines work in the artificial intelligence and operations research communities to address the problem of metric reasoning. We present a number of enhancements and alternatives to the architecture which focus on modifying the interaction of the linear programming and satisfiability components of LPSAT. These include incremental updates to the constraint set, speedup learning via conflict set (nogood) construction in the LP engine, and an optimising variant of LPSAT (i.e. a variant which searches for the best solution according to some measure rather than accepting any viable solution).

Finally, to demonstrate the utility of the LCNF approach in a concrete domain, we present a fully implemented compiler for resource planning problems. Figure 1 shows how the components fit together: a compiler translates the planning problem into LCNF, the LPSAT system solves the LCNF problem and a decoder translates the truth/real-value assignment into a plan. The performance of the system is impressive: LPSAT solves large resource planning problems (encoded in a variant of the PDDL language (McDermott, 1998) based on the metric constructs used by another metric planning system, IPP (Koehler, 1998)), including a metric version of the logistics domain (Kautz & Selman, 1996; Veloso, 1992).

2 The LCNF formalism

The LCNF representation combines a propositional logic formula with a set of metric constraints in a style similar to that proposed by Hooker *et al*. (1999). Truth assignments to the Boolean satisfiability portion of the problem define the metric constraint set, and both the satisfiability and metric portions must be solved to solve the entire LCNF problem.[1]

The key to the encoding is the simple but expressive concept of triggers – each propositional variable may *trigger* a constraint; this constraint is then enforced whenever the variable's truth assignment is `true`. A variable with an associated constraint is called a *trigger variable*. Using this framework, we can construct logical clauses that simulate constraints triggered by `false` assignments, multiple constraints triggered by a single variable and other more complex triggering schemes.

Formally, an LCNF *problem* is a five-tuple $\langle \mathcal{R}, \mathcal{V}, \Delta, \Sigma, \mathcal{T} \rangle$ in which $\mathcal{R}$ is a set of real-valued variables; $\mathcal{V}$ is a set of propositional variables; Δ is a set of linear equality and inequality constraints over variables in $\mathcal{R}$; Σ is a propositional formula in CNF over variables in $\mathcal{V}$; and $\mathcal{T}$ is a function from $\mathcal{V}$ to Δ which establishes the constraint triggered by each propositional variable. We require that Δ contain a special `null` constraint that is vacuously true, and this is used as the $\mathcal{T}$-value for a variable in $\mathcal{V}$ to denote that it triggers no constraint. Moreover, for each variable v we define $\mathcal{T}(\neg v) = \texttt{null}$.

Under this definition, an assignment to an LCNF problem is a mapping, φ, from the variables in $\mathcal{R}$ to real values and from the variables in $\mathcal{V}$ to truth values. Given an LCNF problem and an assignment, the *set of active constraints* is $\{c \in \Delta \mid \exists v \in \mathcal{V}\ \varphi(v) = \texttt{true} \wedge \mathcal{T}(v) = c\}$. We say that an assignment *satisfies* the LCNF problem if and only if it makes at least one literal true in each clause of Σ and satisfies the set of active constraints. A partial assignment to an LCNF problem is a mapping from

[1] Our current LCNF specification does not define any function for measuring the quality of the solution, so an LCNF solver need not be optimising.

$MaxLoad \Rightarrow$ `(load` $\leq$ `30)`	; Statements
$MaxFuel \Rightarrow$ `(fuel` $\leq$ `15)`	; defining
$MinFuel \Rightarrow$ `(fuel` $\geq$ `7 + load / 2)`	; triggered
$AllLoaded \Rightarrow$ `(load = 45)`	; constraints
$MaxLoad$	; Triggers for load and
$MaxFuel$	; fuel limits are unit
$Deliver$	; The goal is unit
$\neg Move \vee MinFuel$	; Moving requires fuel
$\neg Move \vee Deliver$	; Moving implies delivery
$\neg GoodTrip \vee Deliver$	; A good trip requires
$\neg GoodTrip \vee AllLoaded$	; a full delivery

Figure 2 Portion of a tiny LCNF logistics problem (greatly simplified from compiler output). A truck with load and fuel limits makes a delivery but is too small to carry all load available (the *AllLoaded* constraint). Italicized variables are Boolean-valued; typewriter font are real

variables in $\mathcal{V}$ to values in the domain {`true`, `false`, `unassigned`}. We say that a partial assignment ϑ is *inconsistent* with respect to an LCNF problem if there is no assignment φ that satisfies the problem such that for all $v \in \mathcal{V}$ either $\varphi(v) = \vartheta(v)$ or $\vartheta(v) =$ `unassigned`. Intuitively, then, an assignment is satisfying if it makes each propositional clause true and triggers a consistent constraint set; a partial assignment is inconsistent if it cannot be extended to form a satisfying assignment.

Figure 2 shows a fragment of a very simple LCNF problem: a truck, which carries a maximum load of 30 and fuel level of 15, can make a *Delivery* by executing the *Move* action. We discuss later why it cannot have a *GoodTrip*.

3 The LPSAT solver

The LPSAT architecture uses a systematic SAT solver as the controlling component of the engine and makes calls to an LP system. The LPSAT algorithm is very similar to the Davis–Putnam–Logemann–Loveland (DPLL) algorithm for solving Boolean satisfiability problems (Davis *et al.*, 1962).

DPLL performs a depth-first search through the space of truth assignments, choosing (or revoking) an assignment for a single variable at each step of the search. The choice of which variable to set next is informed by the structure of the satisfiability problem (specifically the proportion of small constraints in which the variable participates). Certain assignments are forced by a pair of heuristics: unit propagation (which sets the values of variables that participate in unary constraints) and pure literal elimination (which sets the values of variables which are always constrained to take on the same value). The search terminates successfully when a satisfying assignment – one that gives values to all variables which are consistent with the constraints of the satisfiability problem – is found. The search backtracks when the current solution is found to be an inconsistent partial assignment – an assignment to some portion of the variables which directly contradicts some constraint.

The key difference is in the definitions of "satisfying assignment" and "inconsistent partial assignment". As described in the previous section, each of these concepts in an LCNF problem refers to the active constraint set; however, in a Boolean satisfiability problem, the solver considers only the CNF formula when checking for satisfiability. We designed LPSAT's SAT component to conform to LCNF's definitions of satisfying and inconsistent by modifying and querying the linear programming constraint set.

An assignment to an LCNF statement is satisfying only if its activated constraint set is consistent. So to handle LCNF problems, the SAT solver's check for satisfaction of its LCNF statement must be modified. The solver still checks if the Boolean portion of the problem is satisfied; however, if the

Boolean portion is satisfied, the solver constructs the active constraint set (according to the truth values of the trigger variables) and checks for consistency with the LP solver. Only if the constraint set is consistent does the solver report satisfaction; otherwise the assignment is inconsistent. In addition, when actually reporting a solution (as opposed to simply reporting satisfiability), the SAT solver must return values for its propositional variables and query the LP engine for consistent values for the metric variables.

A partial assignment to an LCNF statement is inconsistent if it activates an inconsistent constraint set. To accommodate this definition, the SAT system's check for inconsistency must – if the Boolean portion of the problem is found consistent – construct the activated constraint set, check for inconsistency with the LP solver and return the result. If a trigger variable is unassigned, its constraint should *not* be added to the active constraint set (since an extension of the partial assignment may or may not activate that constraint).

We implemented the LPSAT engine by introducing the modifications described above into the RelSAT Boolean satisfiability solver (Bayardo & Schrag, 1997) and combining it with the Cassowary incremental linear programming system (Badros & Borning, 1998; Borning *et al*., 1997). We chose RelSAT because it was the best available systematic satisfiability solver and its code was well structured. RelSAT – like all DPLL-style solvers – performs a systematic, depth-first search through the space of partial truth assignments. It also incorporates powerful learning optimisations which make it competitive with the best of the non-systematic SAT solvers. We chose Cassowary because it implements an incremental version of the Simplex Linear Programming algorithm and so supports and quickly responds to small changes in its constraint set.

In addition to modifying RelSAT's check for satisfying and inconsistent assignments, we also had to weaken the pure literal elimination rule. In a CNF problem, pure literal elimination may eliminate certain solutions but it preserves satisfiability. However, in an LCNF problem, pure literal elimination may not preserve satisfiability because a pure positive trigger variable may trigger an inconsistent constraint, so setting that variable to true would make the problem unsatisfiable. In order to maintain the satisfiability-preserving property of pure literal elimination, we never consider trigger variables to be pure positive.[2]

In general, any heuristic which, like the pure literal elimination rule, preserves satisfiability but not individual solutions may need to be modified to work with LCNF problems. However, heuristics which, like unit propagation, prune only inconsistent assignments from the search space will be applicable to LCNF problems without change.[3] Heuristics which do not prune the search space but merely guide search do not *need* to be modified, but the heuristic may benefit from considering the constraints. For example, we altered RelSAT's variable scoring heuristic for choosing splitting variables at branch points to consider whether a potential split variable is a trigger.

Figure 3 displays pseudocode for our LPSAT algorithm. The algorithm is based on DPLL (Davis *et al*., 1962); it performs a depth-first search through the space of partial truth assignments. The search backtracks if it reaches an inconsistent partial assignment and succeeds if it finds a satisfying assignment. The pure literal and unit propagation heuristics guide search. Although our discussion and this pseudocode are targeted at systematic solvers, we feel that other kinds of SAT solvers (e.g. stochastic solvers) could be similarly adapted for use in an LCNF solver.

These simple changes will result in a functional LCNF engine; however, more elaborate communication between the SAT and LP modules can exploit the strengths of each module and dramatically improve the performance and capabilities of the whole system. In the next sections, we describe incrementally updating the constraint set, constructing propositional conflict sets (nogoods) from the constraint set, and altering the satisficing SAT system to form an optimising LPSAT variant.

[2] This restriction falls in line with the pure literal elimination rule if we form implicit clauses describing the triggering action. For instance, the implicit clause for the trigger $MaxLoad \Rightarrow (\texttt{load} \leq 30)$ from Figure 2 would be $\neg MaxLoad \vee (\texttt{load} \leq 30)$ and would introduce a negated instance of the trigger variable, *MaxLoad*.

[3] This is because the set of satisfying assignments for an LCNF problem is a subset of the set of satisfying assignments to its CNF portion.

Procedure LPSAT(LCNF problem: $\langle \mathcal{R}, \mathcal{V}, \Delta, \Sigma, \mathcal{T} \rangle$)
1 If $\exists$ an empty clause in Σ or INC?(Δ), return $\{\bot\}$.
2 Else if Σ is empty, return SOLVE(Δ).
3 Else if $\exists$ a pure literal u in Σ and $\mathcal{T}(u) =$ `null`,
4 return $\{u\} \cup$ LPSAT($\langle \mathcal{R}, \mathcal{V}, \Delta, \Sigma(u), \mathcal{T} \rangle$).
5 Else if $\exists$ a unit clause $\{u\}$ in Σ,
6 return $\{u\} \cup$ LPSAT($\langle \mathcal{R}, \mathcal{V}, \Delta \cup \mathcal{T}(u), \Sigma(u), \mathcal{T} \rangle$).
7 Else
8 Choose a variable, v, mentioned in Σ.
9 Let $\mathcal{A} =$ LPSAT($\langle \mathcal{R}, \mathcal{V}, \Delta \cup \mathcal{T}(v), \Sigma(v), \mathcal{T} \rangle$).
10 If $\bot \notin \mathcal{A}$, return $\{v\} \cup \mathcal{A}$.
11 Else, return $\{\neg v\} \cup$ LPSAT($\langle \mathcal{R}, \mathcal{V}, \Delta, \Sigma(\neg v), \mathcal{T} \rangle$).

Figure 3 Core LPSAT algorithm (without learning). INC? denotes a check for constraint inconsistency; SOLVE returns constraint variable values. $\mathcal{T}(u)$ returns the constraint triggered by u (possibly `null`). $\Sigma(u)$ denotes the result of setting literal u true in Σ and simplifying. The DPLL algorithm is similar but makes no reference to $\mathcal{R}$, Δ, trigger variables, inconsistency checks, or metric constraint solves

4 Incremental updates to the constraint set

Incremental Simplex systems, including Cassowary, are capable of maintaining and incrementally modifying a constraint set, often much faster than the entire set could be constructed and solved from scratch. In order to take advantage of this behaviour in an LCNF engine, the SAT module's procedures for setting (and possibly unsetting) variable values must be modified to notify the LP system of changes in the active constraint set. In particular, when a trigger variable's value becomes `true`, its associated constraint should be added to the LP constraint set; when its value ceases to be `true`, the associated constraint should be removed from the LP constraint set.

However, for current SAT and LP systems, it is generally much faster to set or unset a single propositional variable's value than it is to add or remove a single constraint (even for incremental LP systems). Therefore it can be valuable to delay commitment of constraints to the constraint set.

With delayed commitment, rather than adding constraints directly to the LP system, the SAT system adds and removes constraints from a cache that maintains the difference between the current constraint set and the set active in the LP system. The LP system never contains any inactive constraint. However, some active constraints are not represented in the LP system and instead exist only in the cache. LPSAT maintains this invariant by actually removing from the LP system any constraint which is deactivated but was in the LP system and not just the cache. Since the LP system's constraint set is a subset of the active constraint set, it is always, at most, as constrained as the active constraint set. Therefore satisfiability checks are complete but not sound (because the extra constraints may invalidate all apparent solutions) and, conversely, inconsistency checks are incomplete but sound. When the SAT system requires a complete inconsistency check or a sound satisfiability check, all constraints in the cache are added to the LP system, making it a faithful representation of the active constraint set. Making a sound (empty cache) satisfiability check before committing to a solution allows the solver to re-establish soundness. By designating the other events that add the cache's constraints to the LP system, an LCNF engine can strike a balance between Boolean and metric reasoning time.

These techniques – both incremental updates and delayed commitment – are implemented in the LPSAT system. LPSAT dumps its cache each time it reaches a branch point in the satisfiability problem. This avoids performing any LP reasoning for those branches on which straightforward Boolean logical reasoning (e.g. unit propagation) will produce a contradiction. LPSAT also dumps its cache when checking whether a promising assignment is actually a solution. These methods provide a small but consistent efficiency benefit over using no cache at all.

5 Learning from metric constraint conflicts

LPSAT inherits methods for speedup learning from RelSAT (Bayardo & Schrag, 1997). LPSAT's depth-first search of the propositional search space creates a partial assignment to the Boolean

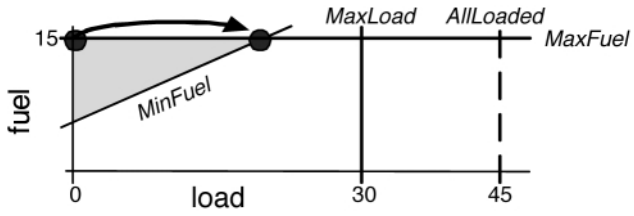

Figure 4 Graphical depiction of the constraints from Figure 2. The shaded area represents solutions to the set of solid-line constraints. The dashed *AllLoaded* constraint causes an inconsistency

variables. When the search is forced to backtrack, it is because the partial assignment is inconsistent with the LCNF problem. LPSAT identifies an inconsistent subset of the truth assignments in the partial assignment, a *conflict set*, and uses this subset to enhance its reasoning in two ways. First, since making the truth assignments represented in the conflict set leads inevitably to failure, LPSAT can learn a clause disallowing those particular assignments. For example, in the problem from Figure 2 the constraints triggered by setting *MinFuel*, *MaxFuel*, *MaxLoad* and *AllLoaded* to true are inconsistent, and *MinFuel*, *MaxFuel* and *AllLoaded* form a conflict set. So LPSAT can learn the clause $(\neg MinFuel \vee \neg MaxFuel \vee \neg AllLoaded)$. Second, because continuing the search is futile until at least one of the variables in the conflict set has its truth assignment changed, LPSAT can backjump in its search to the deepest branch point from which a conflict set variable received its assignment, ignoring any deeper branch points. Figure 5 shows a search tree in which *MinFuel*, *MaxFuel*, *MaxLoad* and *AllLoaded* have all been set to true. Using the conflict set containing *MinFuel*, *MaxFuel* and *AllLoaded*, LPSAT can backjump past the branchpoint for *MaxLoad* to the branchpoint for *MinFuel*, the deepest branchpoint at which a member of the conflict set can be changed.

However, while LPSAT inherits methods to use conflict sets from RelSAT, LPSAT must *produce* those conflict sets for both propositional and constraint failures while RelSAT produces them only for propositional failures. Given a propositional failure, LPSAT uses RelSAT's conflict set discovery mechanism unchanged, learning a set based on two of the clauses that led to the contradiction (Bayardo & Schrag, 1997). However, discovering constraint conflict sets requires a new mechanism.

While LPSAT *could* return the entire partial assignment as a conflict set upon discovering an inconsistency in the active constraint set, paring this set down to a smaller set of assignments yields greater pruning action. Since only the trigger variables with `true` values in the partial assignment add to the active constraint set, only those variables need to be included in the conflict set. We call the resulting set a *global conflict set*. To construct this conflict set, the SAT system queries the LP system to get the constraint set, then it maps the constraints back to the variables that triggered them.

Just as using the triggers of a global conflict set was an improvement over using the entire partial assignment, using any subset of the global conflict set would be an improvement over using the entire set. The logical extension of this idea is to create a conflict set which itself comprises a set of inconsistent constraints but of which every strict subset *is* consistent. We call such a set a *minimal*

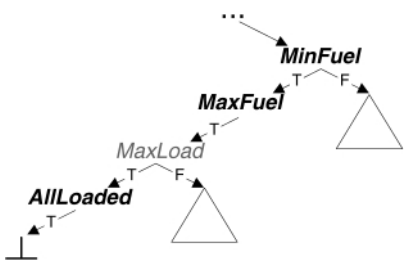

Figure 5 Possible search tree for the constraints from Figure 2. Each node is labelled with the variable set at that node; branchpoints have true (T) and false (F) branches. ⊥ indicates an inconsistent constraint set. The bold variables are members of the conflict set

conflict set. Discovery of both global and minimal conflict sets is implemented in LPSAT; Section 8.1 presents experimental results which show the effectiveness of conflict sets in speedup learning.

Informally, LPSAT finds a minimal conflict set by identifying only those constraints that are, together, in greatest conflict – causing the most error – with the new constraint. In Figure 4 the constraints *MaxLoad*, *MaxFuel* and *MinFuel*, and the implicit constraints that `fuel` and `load` be non-negative, are all consistent; however, with the introduction of the dashed constraint marked *AllLoaded* the constraint set becomes inconsistent. We now discuss how LPSAT discovers the conflicting constraints in this figure and which set it discovers.

When LPSAT adds the *AllLoaded* constraint to Cassowary's constraint set, Cassowary initially adds a "slack" version of the constraint that allows error and is thus trivially consistent with the current constraint set. This error is then minimised by the same routine used to minimise the overall objective function (Badros & Borning, 1998). In Figure 4 we show the minimization as a move from the initial solution at the upper left corner point to the solution at the upper right corner point of the shaded region. The error in the solution is the horizontal distance from the solution point to the new constraint *AllLoaded*. Since no further progress within the shaded region can be made towards *AllLoaded*, the error has been minimised; however, since the error is non-zero, the strict constraint is inconsistent.

At this point, LPSAT constructs a minimal conflict set using *marker variables* (which Cassowary adds to each original constraint). A marker variable is a variable that was added by exactly one of the original constraints and thus identifies the constraint in any derived equations. LPSAT examines the derived equation that gives the error for the new constraint, and notes that each constraint with a marker variable in this equation contributes to keeping the error non-zero. Thus all the constraints identified by this equation, plus the new constraint itself, compose a conflict set. We prove that this technique returns a minimal conflict set in a related technical report (Wolfman & Weld, 1999a).

In Figure 4 the *MinFuel* and *MaxFuel* constraints restrain the solution point from coming closer to the *AllLoaded* line. If the entire active constraint set were any two of those three constraints, the intersection of the two constraints' lines would be a valid solution; however, there is no valid solution with all three constraints.

Note that another conflict set – *AllLoaded* plus *MaxLoad* – exists. In general, there may be many minimal conflict sets, but our conflict discovery technique can discover only one of these per solve. Using careful modifications to the active constraint set and multiple incremental LP solves, we can find many or all of the minimal conflict sets. In order to differentiate between these, we may need to use a different metric from the one that led us to use minimal conflict sets over global conflict sets. Using the size of the sets is still an option, but since none of these sets is a subset of any of the others, a smaller set is no longer guaranteed to result in better (or even equally good) pruning of the search.

In order to allow the system to refine the choice of minimal conflict set, we can pass a ranking of the trigger variables (and thus the constraints) from the SAT system to the LP system. Using a linear (in the number of active constraints) number of calls to the minimal conflict set discovery mechanism described above, the LP system can construct the minimal conflict set with the highest ranked constraints. The algorithm starts from the lowest-ranked constraints and removes them one by one until the set becomes consistent, using the minimal conflict sets constructed at each resolve to guide the search. Once the lowest-ranked constraint in the minimal conflict set is established, it permanently includes that constraint and moves on to establishing the next lowest-ranked constraint. Pseudocode for this algorithm appears in Figure 6. Note that this is an anytime algorithm; it continually improves its solution but always has a solution available. For RelSAT, a DPLL-style solver, we propose ranking the trigger variables in order of their depth in the search tree with the highest ranked variables appearing nearest to the root.[4]

This algorithm has not yet been implemented in the LPSAT engine. Moreover, there is no guarantee that the "best" set produced by this algorithm will be better than the initial conflict set, and the derivation process may substantially slow down the learning process. Therefore, we envision either

[4] Thanks to Rao Kambhampati for this suggestion.

```
Procedure BEST-CONFLICT-SET(Δ: constraints {1, ..., |Δ|})
1  Let M = REV-SORT(MIN-CONFLICT-SET(Δ)).
2  Return BCS-HELPER(Δ, M, 1).

Procedure BCS-HELPER( Δ: remaining constraints in increasing order,
                      M: best conflict set so far in decreasing order,
                      i: integer)
1  If i = |M|, return M.
2  Else
3      Let Δ' = Δ − {M_{i+1}, M_{i+1} + 1, ..., M_i − 1}.
4      Let M' = REV-SORT(MIN-CONFLICT-SET(Δ')).
5      If M' = ∅,
6          return BCS-HELPER(Δ' ∪ {M_{i+1}}, M, i + 1).
7      Else
8          return BCS-HELPER(Δ, M', i).
```

Figure 6 Pseudocode for finding the "best" conflict set for the set Δ. The best set is a minimal conflict set, and its worst element is better than the worst element of all other minimum conflict sets, ties broken by comparing successively better pairs of elements. Constraints are numbered from best, 1, to worst, |Δ|. We assume that the constraint |Δ| caused the inconsistency and so is a member of every minimal conflict set. REV-SORT sorts a set into decreasing order; MIN-CONFLICT-SET finds a minimal conflict set or returns ∅ if the set is consistent; the set $\mathcal{M}$ is made up of the elements $\{\mathcal{M}_1, \ldots, \mathcal{M}_{|\mathcal{M}|}\}$

performing this check only occasionally or cutting off the search for the best set after a fixed number of steps.

6 LPSAT for optimisation problems

Since LP systems provide a clear notion of optimality – minimise (or maximise) the value of an objective function over the variables in the problem – it is natural to extend this notion to an LCNF system. Given an objective function over the metric variables, we define an optimal solution to an LCNF problem to be that satisfying solution which yields the minimal value for the objective function. However, choosing the optimal values for the LP variables in a given active constraint set will not necessarily minimise the objective function over all possible active constraint sets. There may be another satisfying assignment to the Boolean variables in the problem that activates a constraint set with a better value for the objective function. Therefore, in order to construct an optimising LCNF solver, the satisficing SAT system must be modified.

LPSAT's systematic SAT engine has the capacity to find every possible solution to a SAT problem by continuing its depth-first search even after a solution is found. Using this capacity and keeping track of the optimal solution so far, we can construct an optimising version of LPSAT. Of course, not every solution needs to be visited; optimising LPSAT can use a branch-and-bound strategy to eliminate unpromising search branches. This is because the objective value of a partial assignment is always at least as good as the value of its extensions since, given a constraint set with some value for the objective function, more constraints can never improve that value.

Implementing this modification to the LCNF system requires enhancing the communication between the SAT and LP components; in this case, the SAT component must query the LP component for objective values of each partial (and total) assignment. Also, the LP system must have an objective function on which to base its evaluations. In Cassowary, there is a default objective function; however, in general, an optimising LCNF system would require support for an objective function in the LCNF language.

Neither these optimising extensions nor the enhancements to LCNF have yet been implemented for the LPSAT system.

7 The resource planning application

Planning with resources or "metric" planning is the problem of suggesting a course of action for an agent in a domain which includes scarce resources (usually measurable as real- or integer-valued quantities). These domains take the form of descriptions of available actions specifying the conditions under which the actions can be performed and the effects of those actions on the agent and the world.

In order to demonstrate LPSAT's utility, we implemented a compiler for resource-planning domains which translates metric planning problems into LCNF form. After LPSAT solves the LCNF problem, a small decoding unit maps the resulting Boolean and real-valued assignments into a solution plan (Figure 1). We believe that this translate/solve/decode architecture will prove effective for a wide variety of problems.

7.1 Action language

Our planning problems are specified in an extension of the PDDL language (McDermott, 1998). In the specification of a planning problem, we allow metric values with two new built-in types: *float* and *fluent*. A float's value may not change over the course of a plan, whereas a fluent's value may change from time step to time step. Moreover, we support fluent- and float-valued functions, such as `?distance[Seattle,Durham]`.

Floats and fluents are manipulated with three special built-in predicates: `test`, `set`, and `influence`. `Test` statements are used as testing predicates in the preconditions of actions; `set` and `influence` are used to indicate changes to metric values in action effects. As its argument, `test` takes a constraint (an equality or inequality between two expressions composed of floats, fluents and basic arithmetic operations); it evaluates to true if and only if the constraint holds. `Set` and `influence` each take two arguments: the object (a float or fluent) and an expression. If an action causes a `set` to be asserted, the object's value after the action is constrained to equal the expression's value before the action. An asserted `influence` changes an object's value by the value of the expression, as in the equation *object* := *object* + *expression*; multiple simultaneous `influences` are cumulative in their effect (Falkenhainer & Forbus, 1988).

7.2 Plan encoding

As the basis for our compiler, we use a technique described in Ernst *et al*. (1997).[5] We adopt a model in which time takes nonnegative integer values. Predicates describing the state of the world occur at even-numbered times, and predicates describing actions occur at odd times. The initial state is completely specified at time zero, including all properties not explicitly specified in the initial state of the world (which are presumed false by the closed-world assumption).

Each `test`, `set`, and `influence` statement compiles to a propositional variable that triggers the associated constraint. These propositions are logically implied by their associated actions, so a statement restricting a metric variable will become active if the action that causes it occurs in the plan.

The compiler must generate frame axioms for constraint variables – statements which ensure that the value of a resource does not change from step to step unless a specific action causes that change. We use two steps to create these axioms. First, we create a constraint which, if activated, will set the value of the variable at the next step equal to its current value plus all the influences that might act on it (untriggered influences are set to zero). Next, we construct a clause which activates this constraint unless some action actually sets the variable's value.

For an encoding that allows actions to occur simultaneously, the compiler must create clauses that make certain `set` and `influence` statements mutually exclusive. For simplicity, we adopt the

[5] In particular, in the terms described by Ernst *et al*. (1997), our compiler uses a regular action representation with explanatory frame axioms and conflict exclusion.

```
Action loop_a                Action loop_b
 pre: test fluent1 = 0        pre: test fluent2 = 0
 eff: set fluent2 = 1         eff: set fluent1 = 1
```

Figure 7 Two actions which can execute in parallel, but which cannot be serialised

following convention: two actions are mutually exclusive due to their metric effects if at least one `sets` a variable which the other either `influences` or `sets`.

This exclusivity policy results in a plan which is correct if actions at each step are executed strictly in parallel; however, the actions may not be serialisable, as demonstrated in Figure 7. In order to make parallel actions arbitrarily serialisable, we would have to adopt more restrictive exclusivity conditions and a less expressive format for our `test` statements.

8 Experimental results

There are currently few available metric planners with which to compare LPSAT. The ZENO system (Penberthy & Weld, 1994) is more expressive than our system, but Zeno is unable even to complete *easy-1*, our simplest metric logistics problem. There are only a few results available for Koehler's metric IPP system (Koehler, 1998), and code is not yet available for direct comparisons.

In light of this, this section concentrates on displaying results for LPSAT in an interesting domain and on displaying our heuristics and the benefits of communication between the LP and SAT components. We report LPSAT solve time, running on a Pentium II 450 MHz processor with 128 MB of RAM, averaged over 20 runs per problem, and showing 95 per cent confidence intervals. We do not include compile time for the (unoptimised) compiler since the paper's focus is the design and optimisation of LPSAT; however, compile time can be substantial (e.g. twenty minutes on *m-log-c*) since it is heavily memory-bound.

We report on a sequence of problems in the metric logistics domain, which includes all the features of the ATT logistics domain (Kautz & Selman, 1996; Veloso, 1992): airplanes and trucks moving packages among cities and sites within cities. However, our metric version adds fuel and distances between cities; airplanes and trucks both have individual maximum fuel capacities, consume fuel to move (the amount is per trip for trucks and based on distance between cities for airplanes), and can refuel at depots. *m-log-a* to *m-log-d* are the same as the ATT problems *log-a* to *log-d* except for the metric component. *easy-1* to *easy-4* are simplifications of *m-log-a* with more elements retained in the higher numbered problems. We report on experiments with learning as well as two other interesting optimisations.

8.1 Learning

The results in Figure 8 demonstrate the improvement in solving times resulting from activating the learning and backjumping facilities which were described in Section 5. Runs were cut off after one hour of solve time (the minimal conflict set technique ran over an hour only twice on *m-log-c* and not at all on easier problems). Without learning or backjumping LPSAT quickly exceeds the maximum time allotted to it. With learning and backjumping activated using global conflict sets, the solver handles larger problems and runs faster. Our best method, minimal conflict sets, quickly solves even some of the harder problems in the metric logistics domain. Figure 9 shows that the minimal conflict technique runs particularly well when verifying the absence of a solution.

8.2 Splitting heuristic

Line 7 of the LPSAT pseudocode (Figure 3) makes a choice of variable v – called the *splitting variable* – before the recursive call; although we do not need to backtrack over this choice (i.e. the choice is not non-deterministic), a good choice of splitting variable can speed search. We expected the existing

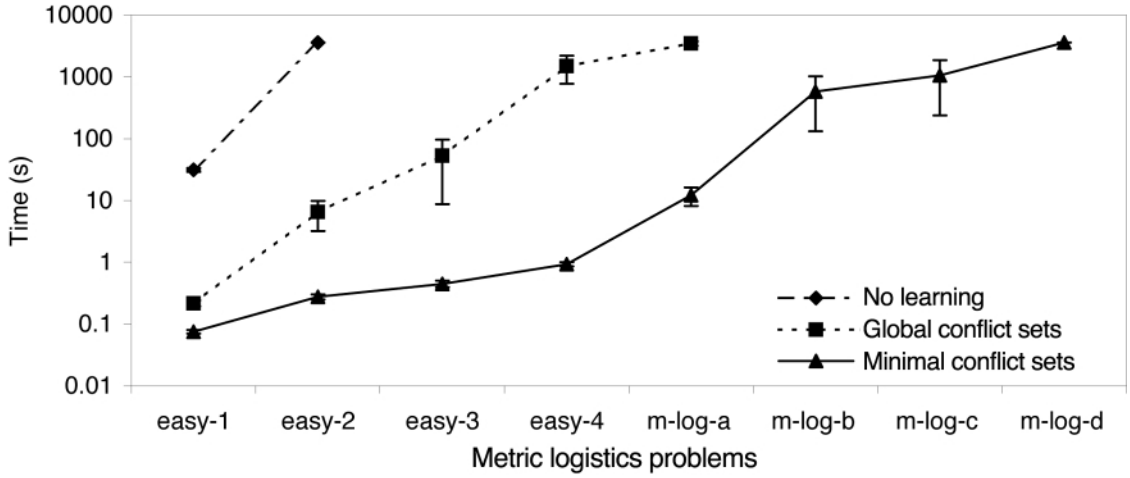

Figure 8 Solution times for three versions of LPSAT in the metric logistics domain. No learning or backjumping is performed in the line marked "No learning". Global conflict sets and minimal conflict sets use progressively better learning algorithms. Note that the final point on each curve reaches the resource cutoff (one hour)

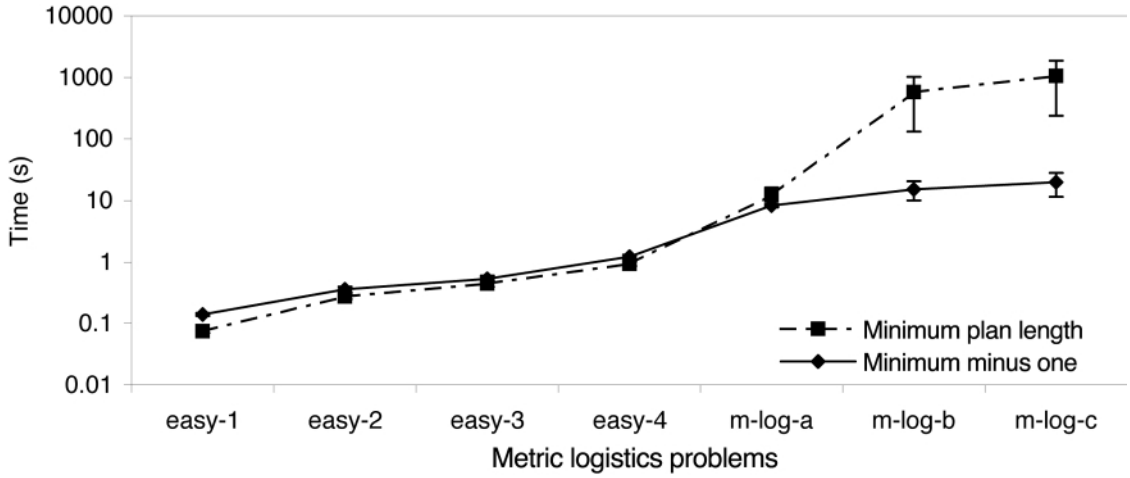

Figure 9 Solution times for LPSAT with minimal conflict sets. The dashed line is the time to find a solution for each problem compiled with the minimum number of steps. The solid line is the time to find that no solution exists for each problem compiled with one fewer steps. LPSAT quickly finds that no solution exists with fewer than the minimum number of steps

RelSAT splitting heuristic to perform poorly because it could not take into account whether a variable is a trigger. This blindness is particularly important since each time the solver modifies a trigger variable it may call the LP system, and these calls often dominate runtime. We tried several methods of including information about the trigger variables in the splitting heuristic, including adding and multiplying the score of trigger variables by a user-settable preference value. To our surprise, modifying the score of trigger variables resulted in no significant improvement. These results are inconclusive but may indicate that our compilation of metric planning domains already encodes some information about trigger variables in the structure of the problem which the current heuristic already uses. Further experiments will decide the issue.

8.3 Random restarts

Because LPSAT uses a randomised backtracking algorithm and because early experimental results showed a small percentage of runs far exceeded the median runtime, we experimented with random

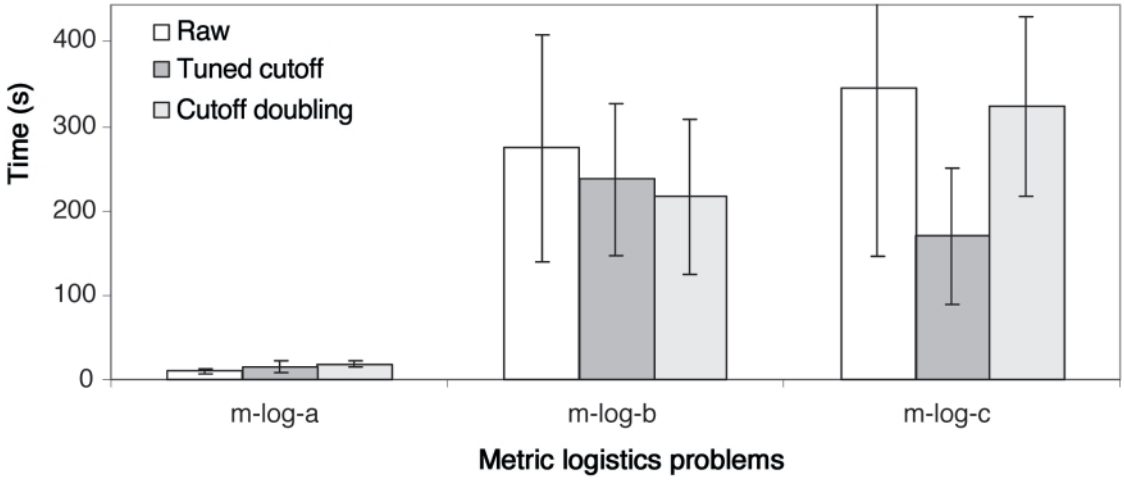

Figure 10 Solution times for two types of random restarts. Tuned cutoff uses raw experimental data to select a constant cutoff. Cutoff doubling starts with a cutoff of one second and doubles it on each run

restarts using a process similar to the one described in Gomes *et al*. (1998). We cut off solving at a deadline – which can be either fixed beforehand or geometrically increasing – and restart the solver with a new random seed.

Figure 10 shows the results of these experiments. We first ran the algorithm twenty times on each problem to produce the "raw" entries.[6] Then we calculated the cutoff time that minimised the expected runtime of the system based on these twenty runs. Finally we reran the problems with this tuned cutoff time to produce the "tuned cutoff" data.

While this technique provides some speedup on *m-log-b* and impressive speedup on *m-log-c*, it requires substantial, preliminary research into the difficulty of the problem (in order to determine the appropriate cutoff time). Unless LPSAT is being used repeatedly to solve a single problem or several very similar problems, the process of finding good restart times will dominate overall runtime.

Therefore we also experimented with a restart system which requires no prior analysis. This "cutoff doubling" approach sets an initial restart limit of one second and increases that limit by a factor of two on each restart until reaching a solution. We have not yet performed any theoretical analysis of the effectiveness of this technique, but Figure 10 demonstrates a small improvement. More interesting than the average improvement, however, is the fact that this method improved the consistency of the runtimes on the harder problems; indeed, on *m-log-c* five of the twenty "raw" runs lasted longer than the longest "cutoff doubling" run.

9 Related work

Limited space precludes a survey of Boolean satisfiability algorithms and linear programming methods in this paper. See Cook & Mitchell (1997) for a survey of satisfiability and Karloff (1991) for a survey of linear programming.

Several other papers in this journal issue also address applying combined artificial intelligence and operations research approaches to planning (Smith *et al*., 2000).

Our work was inspired by the idea of compiling probabilistic planning problems to MAJSAT (Majercik & Littman, 1998). It seemed that if one could extend the SAT "virtual machine" to support probabilistic reasoning, then it would be useful to consider the orthogonal extension to handle metric constraints. Hooker *et al*. (1999) argue convincingly that operations research techniques (such as LP)

[6] All three sets of runs use minimal conflict sets, learning and backjumping.

and artificial intelligence techniques (such as SAT solving) could be combined to their mutual benefit, and our system bears this notion out.

Other researchers have combined logical and constraint reasoning, usually in the context of programming languages. CLPR may be thought of as an integration of Prolog and linear programming, and this work introduced the notion of incremental Simplex (Jaffar *et al*., 1992). Saraswat's thesis (1989) formulates a family of programming languages which operate through the incremental construction of a constraint framework.

A variety of recent systems have addressed the issue of integrating metric reasoning into planning. ILPPLAN (Kautz & Walser, 1999) solves planning problems that have been manually encoded as integer linear programs. While integer linear programming (ILP) is more expressive than LCNF, solvers for ILP problems tend to perform poorly on problems which are primarily propositional (Kautz & Walser, 1999); therefore, LPSAT has the advantage over ILPPLAN on many planning problems, and future solvers for LCNF can continue to exploit advances in SAT engines for solving purely propositional problems. Alternatively, the LPSAT compiler could be used to construct ILPs through a straightforward transformation of the propositional portion to an integer program. Vossen *et al*. (1999) investigate a variety of new encodings and techniques to use ILP to solve planning problems.

CPLAN (van Beek & Chen, 1999) is similar to ILPPLAN but solves planning problems that have been manually encoded as constraint satisfaction problems. While CPLAN's results are promising, no automatic compiler from a planning language to a CPLAN-style CSP exists, and it is unclear how to take advantage of the flexibility of general CSPs.

BLACKBOX uses a translate/solve/decode scheme to solve planning problems as satisfiability problems (Kautz & Selman, 1998). ZENO is a causal link temporal planner which handles resources by calling an incremental Simplex algorithm within the plan-refinement loop (Penberthy & Weld, 1994). The GRAPHPLAN (Blum & Furst, 1995) descendant IPP has also been extended to handle metric reasoning in its plan graph (Koehler, 1998).

10 Conclusions and future work

LPSAT is a promising new technique that combines the strengths of fast Boolean satisfiability solving methods with an incremental Simplex algorithm to efficiently handle problems involving both propositional and metric reasoning. This paper describes several main contributions. We defined the LCNF formalism for combining Boolean satisfiability with linear (in)equalities. We presented the implemented LPSAT solver for LCNF which combines the RelSAT satisfiability solver (Bayardo & Schrag, 1997) with the Cassowary constraint reasoner (Badros & Borning, 1998). We developed a variety of optimisations and enhancements for LPSAT: incrementally updating the constraint set, caching constraints, constructing and improving constraint conflict sets, and an optimising variant of LPSAT. We experimented with three optimisations for LPSAT: adapting the splitting heuristic to trigger variables, adding random restarts, and incorporating learning. Using *minimal conflict sets* to guide learning provided four orders of magnitude of speedup. Finally we described a compiler for resource planning problems.

Much remains to be done. We would like to run experiments comparing LPSAT to systems based purely on integer linear programming. We wish to investigate the issue of tuning restarts to problems, including a thorough investigation of exponentially growing resource limits. It would also be interesting to implement an LCNF solver based on a stochastic engine. It would be interesting to investigate precomputing all minimal conflict sets in an LCNF problem and add appropriate learned clauses to the CNF portion of the problem, entirely removing the need for the metric constraints.[7] Finally, we would like to add dynamic backtracking (Ginsberg & McAllester, 1994) to LPSAT in the hope that it will reduce the number of unnecessary constraint adds and deletes incurred during normal backjumping.

[7] Thanks to Dave Smith for this suggestion.

Acknowledgements

We thank people who provided code, help and discussion: Corin Anderson, Greg Badros, Alan Borning, Mike Ernst, Carla Gomes, Zack Ives, Subbarao Kambhampati, Henry Kautz, Jana Koehler, Tessa Lau, Denise Pinnel, Rachel Pottinger, Bart Selman, Dave Smith, and blind reviewers. This research was funded in part by Office of Naval Research Grant N00014-98-1-0147, by the ARCS Foundation Barbara and Tom Cable fellowship, by National Science Foundation Grants IRI-9303461 and IIS-9872128 and by a National Science Foundation Graduate Fellowship. This paper extends a previous paper appearing in the proceedings of IJCAI-99 (Wolfman & Weld, 1999b).

References

Badros, GJ and Borning, A, 1998, "The cassowary linear arithmetic constraint solving algorithm: interface and implementation" Technical Report 98-06-04, University of Washington, Department of Computer Science and Engineering.

Bayardo, R and Schrag, R, 1997, "Using CSP look-back techniques to solve real-world SAT instances" *Proceedings of the Fourteenth National Conference on Artificial Intelligence* 203–208.

Blum, A and Furst, M, 1995, "Fast planning through planning graph analysis" *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence* 1636–1642.

Borning, A, Marriott, K, Stuckey, P and Xiao, Y, 1997, "Solving linear arithmetic constraints for user interface applications" *Proceedings of the 1997 ACM Symposium on User Interface Software and Technology* 87–96.

Cook, S and Mitchell, D, 1997, "Finding hard instances of the satisfiability problem: a survey" *Proceedings of the DIMACS Workshop on Satisfiability Problems* 11–13.

Crawford, J and Auton, L, 1993, "Experimental results on the cross-over point in satisfiability problems" *Proceedings of the Eleventh National Conference on Artificial Intelligence* 21–27.

Davis, M, Logemann, G and Loveland, D, 1962, "A machine program for theorem proving" *C. ACM* **5** 394–397.

Ernst, M, Millstein, T and Weld, D, 1997, "Automatic SAT-compilation of planning problems" *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence* 1169–1176.

Falkenhainer, B and Forbus, K, 1988, "Setting up large scale qualitative models" *Proceedings of the Seventh National Conference on Artificial Intelligence* 301–306.

Ginsberg, ML and McAllester, DA, 1994, "Gsat and dynamic backtracking" *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning* 226–237.

Gomes, CP, Selman, B and Kautz, H, 1998, "Boosting combinatorial search through randomization" *Proceedings of the Fifteenth National Conference on Artificial Intelligence* 431–437.

Hooker, JN, Ottosson, G, Thorsteinsson, ES and Kim, H, 1999, "On integrating constraint propagation and linear programming for combinatorial optimization" *Proceedings of the Sixteenth National Conference on Artificial Intelligence* 136–141.

Jaffar, J, Michaylov, S, Stuckey, P and Yap, R, 1992, "The CLP ($\mathcal{R}$) language and system" *ACM Transactions on Programming Languages and Systems* **14**(3) 339–395.

Karloff, H, 1991, *Linear Programming* Birkhäuser.

Kautz, H and Selman, B, 1996, "Pushing the envelope: planning, propositional logic, and stochastic search" *Proceedings of the Thirteenth National Conference on Artificial Intelligence* 1194–1201.

Kautz, H and Selman, B, 1998, "Blackbox: a new approach to the application of theorem proving to problem solving" *AIPS98 Workshop on Planning as Combinatorial Search* 58–60.

Kautz, H and Walser, JP, 1999, "State-space planning by integer optimization" *Proceedings of the Sixteenth National Conference on Artificial Intelligence* 526–533.

Kautz, H and Walser, JP, 2000, "Integer optimisation models of AI planning problems" *Knowledge Engineering Review* **15**(1) 101–117.

Koehler, J, 1998, "Planning under resource constraints" *Proceedings of the Thirteenth European Conference on Artificial Intelligence* 489–493.

Li, C and Anbulagan, 1997, "Heuristics based on unit propagation for satisfiability problems" *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence* 366–371.

Majercik, SM and Littman, ML, 1998, "MAXPLAN: a new approach to probabilistic planning" *Proceedings of the Fourth International Conference on Artificial Intelligence Planning Systems* 86–93.

McDermott, D, 1998, "PDDL – the planning domain definition language" AIPS-98 Competition Committee, draft 1.6 edition, June 1998.

Penberthy, JS and Weld, D, 1994, "Temporal planning with continuous change" *Proceedings of the Twelfth National Conference on Artificial Intelligence* 1010–1015.

Saraswat, VA, 1989, "Concurrent constraint programming languages" Ph.D. thesis, Carnegie-Mellon University, Computer Science Department.

Selman, B, Kautz, H and Cohen, B, 1996, "Local search strategies for satisfiability testing" *DIMACS Series in Discrete Mathematics and Theoretical Computer Science* **26** 521–532.

Selman, B, Kautz, H and McAllester, D, 1997, "Computational challenges in propositional reasoning and search" *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence* 50–54.

Smith, DE, Frank, J and Jonsson, AK, 2000, "Bridging the gap between planning and scheduling" *Knowledge Engineering Review* **15**(1) 47–83.

van Beek, P and Chen, X, 1999, "CPlan: A constraint programming approach to planning" *Proceedings of the Sixteenth National Conference on Artificial Intelligence* 585–590.

Veloso, M, 1992, "Learning by analogical reasoning in general problem solving" Ph.D. thesis, Carnegie Mellon University; available as technical report CMU-CS-92–174.

Vossen, T, Ball, M, Lotem, A and Nau, D, 1999, "On the use of integer programming models in AI planning" *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence* 304–309.

Vossen, T, Ball, M, Lotem, A and Nau, D, 2000, "Applying integer programming to AI planning" *Knowledge Engineering Review* **15**(1) 85–100.

Wolfman, S and Weld, D, 1999a, "The LPSAT system and its application to resource planning" Technical Report 99-04-04, University of Washington, Department of Computer Science and Engineering.

Wolfman, S and Weld, D, 1999b, "The LPSAT system and its application to resource planning" *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence* 310–317.