

# Preference-Based Planning via MaxSAT

Farah Juma, Eric I. Hsu, and Sheila A. McIlraith

Department of Computer Science, University of Toronto, Toronto, Canada  
{fjuma,eihsu,sheila}@cs.toronto.edu

**Abstract.** In this paper, we explore the application of partial weighted MaxSAT techniques for preference-based planning (PBP). To this end, we develop a compact partial weighted MaxSAT encoding for PBP based on the popular SAS<sup>+</sup> planning formalism. Our encoding extends a SAS<sup>+</sup> based encoding for SAT-based planning, SASE, to allow for the specification of simple preferences. To the best of our knowledge, the SAS<sup>+</sup> formalism has never been exploited in the context of PBP. Our MaxSAT-based PBP planner, MSPLAN, significantly outperformed the state-of-the-art STRIPS-based MaxSAT approach for PBP with respect to running time, solving more problems in a few cases. Interestingly, when compared to three state-of-the-art heuristic search planners for PBP, MSPLAN consistently generated plans with comparable quality, slightly outperforming at least one of these three planners in almost every case. Our results illustrate the effectiveness of our SASE based encoding and suggests that MaxSAT-based PBP is a promising area of research.

## 1 Introduction

Many real-world planning problems consist of both a set of mandatory goals and an additional set of desirable plan properties. The degree of satisfaction of these desirable properties, or plan *preferences*, determines the quality of a plan. Preference-based planning (PBP) extends the well-known classical planning problem in order to generate plans that achieve problem goals while maximizing the satisfaction of other preferred properties of the plan. In so doing, they allow a planner to generate plans of high quality, often under situations with conflicting preferences.

PBP has been the subject of substantial research in recent years. The 2006 International Planning Competition (IPC-2006) created a track on this topic which resulted in the extension of the standardized Planning Domain Description Language (PDDL) to support the specification of preferences [6]. In PDDL3, desirable properties of a plan are expressed as preference formulae. These formulae may describe properties of the final state as well as properties that hold over intermediate states visited during plan execution. The relative importance associated with not violating these preference formulae is reflected in a metric function, a weighted linear combination of preferences whose violation is minimized by the planner. At IPC-2008, this family of planning problems was extended to include action costs. The objective of these so-called *net benefit* problems

is to maximize the sum of the utilities of the goals and preferences that have been achieved, minus total costs. Action costs can be incorporated into our partial weighted MaxSAT-based PBP approach and it is something that we have explored, but do not address it in this paper.

To date, the most effective techniques for PBP have been based on heuristic search (e.g., *Yochan*<sup>PS</sup>, SGPlan<sub>5</sub>, HPLAN-P, and most recently LPRPG-P). There have also been planners that have used SAT, CSP, or Answer Set solvers (e.g., SATPLAN(P), PREFPLAN, CPP) [3]. Recently, Giunchiglia and Maratea explored a partial weighted MaxSAT-based approach to solving PBP problems, which we refer to here as GM [7]. While all of these latter systems show promise, performance has generally been inferior to heuristic search.

In this paper, we characterize the problem of PBP as a partial weighted MaxSAT problem. A major focus of our work is on how to construct an effective encoding. To this end, we propose a SAS<sup>+</sup> based (e.g., [8]) encoding of PBP as MaxSAT that is compact and correct. Our encoding builds on the success of SASE, a SAS<sup>+</sup> based encoding recently developed for SAT-based planning [11]. To the best of our knowledge, the popular SAS<sup>+</sup> formalism has never been used in the context of PBP despite its effectiveness in classical planning. Exploiting our characterization of PBP as a partial weighted MaxSAT problem, we develop a system, MSPLAN, that employs our SASE based encoding.

We experimentally evaluated our system by comparing it to GM on *Simple Preferences* problems from IPC-2006. MSPLAN consistently outperformed GM with respect to running time, in some cases by an order of magnitude. In all cases, plan quality was comparable. We also compared the performance of MSPLAN, run with two different MaxSAT solvers, to state-of-the-art heuristic search planners for PBP. MSPLAN generated plans of comparable plan quality, sometimes slightly out- or under-performing the best of the heuristic search planners we employed. However, as expected, MSPLAN was not able to solve as many problems as the heuristic search planners. While in some instances, one of the MSPLAN systems significantly outperformed HPLAN-P with respect to running time, the heuristic search planners were generally faster. Analysis showed that a significant part of MSPLAN's running time was attributed to the incremental construction of makespans, rather than the search for a solution. Given the consistent quality of MSPLAN solutions, consistent superior performance to GM, and the variability in the performance of all of these systems, we deem our SASE based encoding to be effective compared to a STRIPS encoding, and the use of MaxSAT and SAS<sup>+</sup> encodings for PBP to be promising areas of research.

## 2 PBP as Partial Weighted MaxSAT

In this section, we overview the correspondence between PBP and partial weighted MaxSAT. For the purposes of this paper, we restrict our attention to simple preferences – preferences over properties of the final state of the plan.

**Definition 1 (Planning domain).** *Let  $X$  be a set of state variables. A planning domain is a tuple  $D = (S, A, \gamma)$ , where  $S$  is the set of all possible states;*

$A$  is the set of actions; and  $\gamma$  is the transition function. A state  $s \in S$  is a set of assignments to all of the variables in  $X$ . An action  $a \in A$  is described by a tuple  $(\text{pre}(a), \text{eff}(a))$ , which denotes the action’s preconditions and effects. A transition  $\gamma(s, a)$  modifies the values of the state variables mentioned in  $\text{eff}(a)$ .

**Definition 2 (Simple PBP problem).** A simple PBP problem is a tuple  $P = (D, s_I, s_G, \text{Pref}, W)$ , where  $D = (S, A, \gamma)$  is the planning domain;  $s_I \in S$  is the initial state;  $s_G$  represents the goal and is a partial set of assignments to the state variables; and  $\text{Pref}$ , the preferences, is a partial set of assignments to the state variables. Optionally, each element in  $\text{Pref}$  may have a positive weight associated with it,  $W: \text{Pref} \rightarrow \mathbb{R}_{\geq 0}$ , to capture the relative importance of preferences.

Let  $\phi$  denote a CNF propositional formula over a set  $V$  of boolean variables and let  $\{C_1, \dots, C_m\}$  denote the clauses of  $\phi$ .

**Definition 3 (MaxSAT and (Partial) Weighted MaxSAT).** The MaxSAT problem is to find an assignment of values for  $V$  that maximizes the number of satisfied clauses in  $\phi$ . Given a weight  $w_i$  for each clause  $C_i$  in  $\phi$ , the weighted MaxSAT problem is to find an assignment of values for  $V$  that maximizes the total weight of the satisfied clauses in  $\phi$ . When some clauses in  $\phi$  are designated as hard clauses and other clauses in  $\phi$  are designated as soft clauses and we are given a weight  $w_i$  for each soft clause  $C_i$  in  $\phi$ , the partial weighted MaxSAT problem is to find an assignment of values for  $V$  that satisfies all designated hard clauses in  $\phi$  and maximizes the total weight of the satisfied soft clauses in  $\phi$ .

If  $\text{Pref}$ , the preferences in a simple PBP problem, are encoded as soft clauses, and  $s_I$ ,  $s_G$ , and  $D$  are encoded as hard clauses, and weights drawn from  $W$  are assigned to the soft clauses to indicate the relative importance of the preferences, then from Definition 3, it follows that finding a solution to the resulting partial weighted MaxSAT problem is equivalent to finding a plan for the original PBP problem that achieves all of the hard constraints while maximizing the weight of the satisfied preferences. How we actually encode these clauses is a challenge and is one of the contributions of this paper.

## 3 Preference-Based Planning

### 3.1 A SAS<sup>+</sup> Based Encoding

The SAS<sup>+</sup> formalism (e.g., [8]) has been increasingly exploited in the context of *classical* planning. Unlike a STRIPS-based encoding, which consists of actions and facts and represents a state using a set of facts, a SAS<sup>+</sup> based encoding consists of transitions and multi-valued state variables and represents a state using a set of assignments to all of the state variables. A transition represents a change in the value of a state variable. Consider a transportation domain in which trucks can move packages between locations with certain restrictions. To represent the possible locations of a truck, a STRIPS-based encoding might include a fact for each such location (e.g., `(at truck1 loc1)`, `(at truck1 loc2)`).

In contrast, a SAS<sup>+</sup> based encoding might include a state variable `truck1` whose domain consists of all of the possible truck locations. Moving `truck1` from `loc1` to `loc2` would be represented by an action in a STRIPS-based encoding. In a SAS<sup>+</sup> based encoding, this would be represented by a transition that changes the value of the state variable `truck1` from `loc1` to `loc2`. In general, an action might result in changes in the values of multiple state variables. In a SAS<sup>+</sup> based encoding, this would be represented by multiple transitions. Because the number of transitions in an encoding based on the SAS<sup>+</sup> formalism is normally much less than the number of actions in a STRIPS-based encoding, the search space in a transition-based encoding is smaller than the search space in an action-based encoding [11]. The compactness of the SAS<sup>+</sup> formalism and the rich structural information that can be derived from it are the major advantages of this formalism. Many state-of-the-art planners have been adopting SAS<sup>+</sup> based encodings (e.g., [8]).

The major limitation of previous STRIPS-based SAT encodings for classical planning has been the large number of clauses. The SAS<sup>+</sup> formalism was recently used for the first time in the context of a SAT-based classical planning approach with an impressive reduction in the size of the SAT encoding [11]. To the best of our knowledge, the SAS<sup>+</sup> formalism has never been used in the context of PBP despite its effectiveness in classical planning. We explore this here.

**Translating PBP Problems to SAS<sup>+</sup>.** We consider PBP problems from the *Simple Preferences* track of IPC-2006. STRIPS problems can be compiled into the SAS<sup>+</sup> representation using the parser created for the Fast Downward planner [8]. Since IPC-2006 *Simple Preferences* problems are not represented in STRIPS, we first need to compile them into STRIPS using the compilation technique developed in [7]. The compilation proceeds as follows. For each simple preference, a so-called *dummy* action is created. The precondition of this dummy action is the simple preference property itself and the effect is a *dummy* literal that represents the simple preference. This dummy literal is then added to the definition of the goal for the problem instance and the definition of the simple preference is then removed from the problem. The execution of a dummy action indicates that the corresponding simple preference has been satisfied. For example, consider the following preference that is expressed in PDDL3:

```
(preference time (or (delivered pck1 loc3 t1) (delivered pck1 loc3 t2)))
```

The `time` preference specifies that `pck1` should be delivered to `loc3` at `t1` or `t2`. The compilation technique replaces this preference with a dummy action:

```
(:action dummy-time
:parameters ()
:precondition (or (delivered pck1 loc3 t1) (delivered pck1 loc3 t2))
:effect (dummy-goal-time))
```

The literal `dummy-goal-time` represents the satisfaction of the `time` preference. After introducing dummy actions for each simple preference, the resulting problems are compiled into STRIPS using Gerevini and Serina's ADL2STRIPS tool. Note that we must also maintain the information about which goals in the resulting STRIPS problems correspond to simple preferences.

After translating these problems into SAS<sup>+</sup>, there is a SAS<sup>+</sup> goal variable that corresponds to each dummy literal (i.e., simple preference). We will refer to these SAS<sup>+</sup> variables as preference variables. We will use  $s_P$  to denote the set of desired assignments to preference variables in the final state. As such,  $s_P(x) = p$  indicates that in the final state, the desired value of preference variable  $x$  is  $p$ .

**Encoding the Clauses.** To encode a SAS<sup>+</sup> based PBP problem as a partial weighted MaxSAT problem, we modify the SAS<sup>+</sup> based SAT encoding, SASE, first proposed in [11]. SASE is made up of transition variables  $U$ , action variables  $V$ , and eight classes of clauses. We show these classes below. Note that  $N$  is the number of time steps in the plan,  $\tau(x)$  is the set of possible transitions for  $x$ ,  $O$  is the set of actions,  $M(a)$  is the transition set of action  $a \in O$ ,  $R$  is the set of all prevailing transitions,  $\delta_{f \rightarrow f'}^x$  represents a change in the value of  $x$  from  $f$  to  $f'$ ,  $U_{x,f,f',t}$  (or  $U_{\delta,t}$ , where  $\delta$  is defined as  $\delta_{f \rightarrow f'}^x$ ) is a transition variable that indicates that variable  $x$  changes from the value of  $f$  to  $f'$  at time step  $t$ , and  $V_{a,t}$  is an action variable that indicates that the action  $a$  is executed at time  $t$ .

1. Initial state -  $(\forall x, s_I(x) = f): \bigvee_{\delta_{f \rightarrow g} \in \tau(x)} U_{x,f,g,1}$
2. Goal -  $(\forall x, s_G(x) = g): \bigvee_{\delta_{f \rightarrow g} \in \tau(x)} U_{x,f,g,N}$
3. Transition's progression -  $(\forall \delta_{f \rightarrow g}^x \in \tau \text{ and } t \in [2, N]):$   

$$U_{x,f,g,t} \rightarrow \bigvee_{\delta_{f' \rightarrow f}^x \in \tau(x)} U_{x,f',f,t-1}$$
4. Transition mutex -  $(\forall \delta_1 \forall \delta_2 \text{ such that } \delta_1 \text{ and } \delta_2 \text{ are transition mutex}):$   

$$U_{\delta_1,t} \rightarrow \neg U_{\delta_2,t}$$
5. Existence of transitions -  $(\forall x \in X): \bigvee_{\delta \in \tau(x)} U_{\delta,t}$
6. Composition of actions -  $(\forall a \in O): V_{a,t} \rightarrow \bigwedge_{\delta \in M(a)} U_{\delta,t}$
7. Action's existence -  $(\forall \delta \in \tau \setminus R): U_{\delta,t} \rightarrow \bigvee_{\forall a, \delta \in M(a)} V_{a,t}$
8. Non-interference of actions -  $(\forall a_1 \forall a_2 \text{ such that } \exists \delta, \delta \in T(a_1) \cap T(a_2) \text{ and } \delta \notin R): V_{a_1,t} \rightarrow \neg V_{a_2,t}$

These clauses encode the initial state, the goal state, the transitions that are allowed to occur at various time steps, the relationship between transitions and actions, and mutually exclusive actions. After finding a sequence of transitions that achieves the goal, a corresponding action plan is found.

The key challenge is determining how to encode the preferences. We modify SASE to handle preference variables (i.e., variables that we would *like* to achieve a certain value in the final state) in addition to goal variables (i.e., variables that *must* achieve a certain value in the final state). To this end, we create a new class of clauses that represents the preferences. Specifically, for each preference variable  $x$ , we create a clause that is the disjunction of all transitions which result in the desired value for  $x$  in the final state. We define this as follows:

$$\text{Preferences - } (\forall x, s_P(x) = p): \bigvee_{\delta_{f \rightarrow p} \in \tau(x)} U_{x,f,p,N}$$

It is important to note that once a dummy action is executed, the dummy goal literal it produces persists indefinitely, i.e., no action deletes a dummy goal

literal. Because the execution of a dummy action is meant to indicate that a simple preference is satisfied in the *final* state of the plan, we must ensure that once a dummy action has been executed, no other action that can invalidate the precondition of this dummy action can occur after it. Thus, we restrict the execution of dummy actions,  $O_{dummy}$ , to the final time step of the plan by only defining dummy action variables that correspond to this time step, i.e.,

Dummy action variables -  $V_{a,N}, \forall a \in O_{dummy}$

Because preference variables do not necessarily have to achieve their desired value, we treat all preference clauses as soft clauses. We treat the clauses that encode the initial state, goal state, and the planning domain, as hard clauses.

**Weighting the Clauses.** We assign weights to all of the soft clauses using the PDDL3 metric function. This function is a linear function defined over simple preferences and is used to determine the quality of a plan. An example of a PDDL3 metric function over two preferences named `time` and `loc` is shown below:

```
(:metric minimize(+ (* 10 (is-violated time)) (* 5 (is-violated loc))))
```

Note that the `is-violated` function returns 1 if the preference with the given name does not hold in the final state of the plan and returns 0 otherwise. This metric function indicates that satisfying the `time` preference is twice as important as satisfying the `loc` preference. An IPC-2006 *Simple Preferences* task can thus be viewed as the problem of finding a plan that satisfies all of the hard goals while minimizing the total weight of the preferences that are not satisfied. Since each simple preference is represented by one soft clause in our encoding, we assign a weight to this soft clause based on the weight assigned to this preference in the PDDL3 metric function from the original IPC-2006 problem instance.

### 3.2 Planning with MaxSAT

With our SAS<sup>+</sup> based partial weighted MaxSAT encoding for PBP in hand, the next step is to determine how to find a plan using such an encoding. As in SAT-based planning, an incremental construction of makespans is required, i.e., for a specific value  $n$  for the makespan (the number of time steps in the plan), we must encode a given PBP problem into our SAS<sup>+</sup> based partial weighted MaxSAT encoding with a makespan of  $n$ , attempt to solve the problem using a partial weighted MaxSAT solver, and continue on in this manner, trying increasing values of  $n$ , until a solution is found.

**Stopping Conditions.** Because the task of PBP involves finding a high-quality plan, and not just a plan with the minimum number of time steps (as in SAT-based planning), determining when to stop trying increasing values for the makespan is more difficult. One possibility we consider is to stop trying increasing values for the makespan after the first satisfiable partial weighted MaxSAT formula is found. The solution to this formula corresponds to a plan with optimal makespan and optimal plan quality for that particular makespan. However, this

plan is not guaranteed to be globally optimal. We also consider the case where a time limit is given and an incremental construction of makespans is carried out until the allotted time expires, at which point the plan returned is the plan with the best quality that was found during the given amount of time. We will refer to this algorithm as the BEST algorithm.

**Properties of the Plan.** We can show that for any fixed makespan, our approach will always return a solution with optimal plan quality with respect to that particular makespan and that when restricted to plans with makespan bounded by  $k$ , our approach will always return a solution that is  $k$ -optimal.

**Lemma 1.** *For any fixed makespan  $n$ , the solution to the PBP problem, encoded as a partial weighted MaxSAT problem with makespan  $n$ , yields a plan with optimal quality with respect to the set of all plans with makespan  $n$ , if such a plan exists.*

From Lemma 1, which follows directly from Definition 3, we can conclude that any plan  $P$  returned by our approach has optimal quality with respect to the set of all plans with the same makespan as  $P$ . In certain cases, we may want to restrict our attention to plans with a makespan bounded by some value.

**Definition 4 ( $k$ -Optimality).** *We can say that a partial weighted MaxSAT-based PBP algorithm is  $k$ -optimal if it is always able to find a plan that is optimal, in terms of quality, with respect to the set of all plans with makespan  $i \leq k$ .*

**Theorem 1.** *If the search is restricted to plans with makespan bounded by  $k$  and the BEST algorithm is run long enough for the PBP problem to be encoded into a partial weighted MaxSAT formula using each makespan  $i \leq k$ , then the BEST algorithm is  $k$ -optimal.*

*Proof.* For each makespan  $i \leq k$ , the PBP problem is encoded into a partial weighted MaxSAT formula with makespan  $i$ . If a solution is found, the quality of the plan is optimal w.r.t. the set of all plans with makespan  $i$  (by Lemma 1). The result follows since the BEST algorithm returns the plan with the best quality among those found.

## 4 Implementation and Evaluation

To implement our planner, MSPLAN, we augmented the SASE planner to operate with our PBP encoding and various partial weighted MaxSAT solvers. We employed both MiniMaxSAT v1.0 [9] and SAT4J v2.1 [4]. We attempted to use MSUncore [15] but the developers were unable to provide a system that ran on our hardware. We also modified the PDDL3 metric functions used to specify PBP problems so that all weights were positive integers, as required by partial weighted MaxSAT solvers. This was done via a uniform order of magnitude increase of all weights in a function until weights were integer rather than real. Finally, many partial weighted MaxSAT solvers require a special weight to be specified for hard clauses in addition to weights for soft clauses. Following the convention used to specify this special weight, we assigned a weight to each hard clause that exceeded the sum of the weights of all of the soft clauses.

Our evaluation of MSPLAN was motivated by two objectives. We wanted to: (1) compare the performance of our planner to the only previous, and therefore state-of-the-art, partial weighted MaxSAT-based approach, GM; and (2) compare our planner to state-of-the-art heuristic search planners for PBP. In doing so, we also hoped to gain some insight into the impact of the underlying MaxSAT solver on the performance of MSPLAN.

To support comparison with GM, the domains we evaluated were limited to those used in [14]. Four domains from the IPC-2006 *Simple Preferences* track were used in our evaluation: `trucks`, `storage`, `pathways`, and `openstacks`. The `trucks` and `openstacks` domains contain both simple preferences and hard goals. However, the `pathways` and `storage` domains contain only simple preferences and it is generally not possible to satisfy all of these preferences. MSPLAN requires STRIPS-encodings of these problem instances as input. Since such encodings were also used in [14], we were able to obtain the problems from the *Simple Preferences* track which Giunchiglia and Maratea have been able to compile into STRIPS. The `trucks` and `storage` domains consist of 7 problems, the `pathways` domain consists of 20 problems, and the `openstacks` domain consists of 1 problem.

We compared the performance of MSPLAN to GM. GM uses a STRIPS-based partial weighted MaxSAT encoding as opposed to a SAS<sup>+</sup> based encoding. We also compared the performance of MSPLAN to SGPlan-W [5], HPLAN-P [2], and LPRPG-P [5]. SGPlan-W is a recent re-implementation of SGPlan<sub>5</sub> [10], winner of the IPC-2006 *Simple Preferences* track. HPLAN-P did not compete in this track and came in 2nd to SGPlan<sub>5</sub> in the *Qualitative Preferences* track. Nevertheless, experiments from [2] show that HPLAN-P would have outperformed all entrants in the *Simple Preferences* track, other than SGPlan<sub>5</sub>. Finally, LPRPG-P is a recently developed heuristic search planner, shown to outperform HPLAN-P.

Our experiments were performed on an AMD Opteron 1GHz processor with a timeout of 60 minutes. The memory usage did not exceed 1GB. We were not able to obtain a copy of GM. The results for these experiments are taken from [14]. Different machines have thus been used in our comparison. In [14], experiments were performed on a Pentium IV 3.2GHz processor with 1GB of RAM, a faster machine than ours. Note that [14] gives only the time required for solvers to find a solution to the first satisfiable partial weighted MaxSAT formula. The total amount of time required to find a plan does not appear in [14].

Table 1 shows the performance of MSPLAN compared to GM when evaluating the time required for the solver to find a solution to the first satisfiable partial weighted MaxSAT formula when using MiniMaxSAT v1.0 and SAT4J v2.1 as the underlying solvers. The results show that for all problem instances in the `trucks` and `storage` domains which could be solved using both GM and MSPLAN, the time required to find a solution to the first satisfiable partial weighted MaxSAT formula was less, by an order of magnitude in many cases, when using MSPLAN than when using GM, regardless of the solver used. The value of the plan quality metric does not appear in [14] and we were unable to obtain the corresponding quality values from the authors. However, from a different encoding given to us by the authors, we were able to generate an upper bound on the quality values.

This information indicated that the plan quality was comparable in all cases. This information also indicated that the number of clauses in our encoding was often significantly smaller, by a large constant factor. While this comparison is not precise, it is the best that could be done with the available data and gives an indication of the positive properties of our approach. There were problems in both the `trucks` and `storage` domains which could be solved by MSPLAN with a particular solver but could not be solved by GM when using the same solver. In fact, MSPLAN was able to optimally solve three problems in the `trucks` domain that GM could not solve using any of the partial weighted MaxSAT solvers evaluated in [14]. Neither MSPLAN nor GM could generate a plan for the one problem instance in the `openstacks` domain. Finally, we were not able to do a direct comparison of our results for the `pathways` domain because the running time of GM on these problems was unavailable.

**Table 1.** Performance of MSPLAN compared to GM. Entries indicate the time required by MiniMaxSAT and SAT4J to find a solution to the first satisfiable partial weighted MaxSAT formula, in seconds. Dashes indicate timeout. Stars indicate optimal quality.

Instance	MiniMaxSAT		SAT4J	
	GM	MSP <sub>PLAN</sub>	GM	MSP <sub>PLAN</sub>
trucks1	7.7	1.18	359.17	1.95
trucks2	308.92	44.803	-	24.868
trucks3	-	89.15*	-	446.578*
trucks4	-	128.904*	-	-
trucks5	-	652.877*	-	-
trucks6	-	-	-	-
trucks7	-	-	-	-
storage1	0.21	0.008	0.32	0.171
storage2	0.44	0.032	0.65	0.21
storage3	0.59	0.032	1.45	0.503
storage4	0.71	0.084	2.8	0.667
storage5	58.79	13.721	16.35	1.228
storage6	-	43.059	70.6	2.564
storage7	-	-	365.53	6.232

Table 2 shows the performance of MSPLAN compared to the recent LPRPG-P planner and to the two top heuristic search PBP planners from IPC-2006, the last IPC with a PBP track. In this set of experiments, we ran MSPLAN with the BEST algorithm from Sect. 3.2, i.e., MSPLAN returned the highest quality plan it found during a 60 minute period. We evaluated the total running time of all four planners and the quality of the plans found. The results show that for all problems that all planners were able to solve, the quality of the plan generated by MSPLAN was equal or superior to the quality of the plan returned by at least one of SGPlan-W, HPLAN-P, and LPRPG-P, regardless of the partial weighted MaxSAT solver that was used. As expected, SGPlan-W and LPRPG-P generated plans significantly faster than MSPLAN in all cases. This is consistent with current findings for classical planning, given the superior performance of heuristic search *classical* planners relative to their SAT-based counterparts. However, MSPLAN solved some problems more quickly than HPLAN-P, while returning a plan of equal or better quality. There were problems in the `storage` and `pathways` domains that could be solved by MSPLAN but not by SGPlan-W.

As in SAT-based planning, a bottleneck in partial weighted MaxSAT-based planning seems to be the iteration required to determine the makespan for the solution. In our experiments, we encoded the problem using each possible makespan until the plan with the best quality was found. This is clearly a worst-case scenario. Many SAT-based planners first generate an estimate for the appropriate makespan and use this as a starting point for incremental search. As seen in

**Table 2.** Performance of MSPLAN when run with the BEST algorithm compared to the recent LPRPG-P planner and to the top PBP heuristic search planners from IPC-2006. Inst denotes the problem instance:  $t$  stands for trucks,  $s$  stands for storage,  $p$  stands for pathways,  $o$  stands for openstacks. Let  $B$  be the makespan that results in the plan with the best quality within the time limit. Soln denotes the time required for MiniMaxSAT and SAT4J to find a solution to the partial weighted MaxSAT formula with makespan  $B$ . This includes only the solver time to find the solution to the formula. Oracle denotes the total running time of MSPLAN when given, *a priori*, the best makespan  $B$ . This time includes the time for the translation to SAS<sup>+</sup>, the time for encoding the partial weighted MaxSAT formula with makespan  $B$ , and the solver time. Total denotes the total running time when this best makespan  $B$  is not known ahead of time and must be determined through an incremental construction of makespans. This time includes the time for the translation to SAS<sup>+</sup>, the time for repeatedly encoding a partial weighted MaxSAT formula with increasing makespans, and the total solver times. The total running time of HPLAN-P, SGPlan-W, and LPRPG-P is denoted by Time. M is the value of the plan metric, the total value of the violated preferences in the plan found. (Low is good.) A dash indicates timeout or a crash.

Inst	MSPLAN: MiniMaxSAT				MSPLAN: SAT4J				HPLAN-P		SGPlan-W		LPRPG-P	
	Soln	Time (s)			Soln	Time (s)			Time (s)	M	Time (s)	M	Time (s)	M
t1	2.16	5.08	14.01	0	1.78	5.18	21.22	0	5.58	0	0.03	1	0.03	0
t2	36.51	41.3	165.91	0	36.04	37.66	173.63	0	92.28	0	0.07	3	0.04	0
t3	89.15	111.54	304.1	0	527.21	553.62	861.07	0	660.1	0	0.09	6	0.12	0
t4	128.9	426.26	596.18	0	-	-	-	-	563.93	3	0.16	0	0.22	0
t5	652.88	675.59	2553.01	0	-	-	-	-	1015.36	0	0.27	1	0.16	0
t6	-	-	-	-	-	-	-	-	-	-	0.84	0	1.64	0
t7	-	-	-	-	-	-	-	-	-	-	1.38	59	2.56	0
s1	0.01	1.31	1.66	3	0.17	2.15	2.37	3	0.07	3	0.02	5	0.0	0
s2	4.9	27.62	34.18	5	4.3	27.15	43.58	5	2.93	5	2.26	11	0.01	1
s3	708.84	730.15	1192.33	18	1061.36	1083.39	2769.18	7	38.06	6	552.68	38	0.04	0
s4	23.71	34.51	47.34	38	532.2	637.1	1063.3	24	183.37	9	-	-	0.09	3
s5	70.08	123.29	214.77	107	677.02	705.96	937.9	76	76.57	94	-	-	0.23	0
s6	43.06	71.49	73.19	173	1573.75	1605.23	1978.41	150	459.93	141	-	-	30.7	59
s7	-	-	-	-	127.5	168.81	183.76	277	1280.43	160	-	-	0.41	203
p1	0.05	1.48	1.6	2	0.18	2.56	3.09	2	3.88	2	-	-	0.02	5
p2	0.07	0.99	1.29	3	0.46	1.99	3.34	3	186.34	4	-	-	0.01	6
p3	0.36	3.12	4.95	3	0.59	5.0	7.42	3	115.6	3.7	-	-	0.01	5.7
p4	0.36	5.33	5.85	2	0.78	6.25	8.66	2	324.82	2	-	-	0.02	6.7
p5	6.47	36.59	45.26	6	2.08	35.39	42.73	6	413.42	9	-	-	0.02	10.2
p6	434.02	797.73	1468.57	6.4	114.92	103.81	296.1	6.4	0.04	12.9	-	-	0.04	12.9
p7	1251.16	1466.24	1710.57	11.5	1378.79	1406.88	2325.46	10.3	0.05	12.5	-	-	0.04	12.5
p8	201.349	229.51	823.23	18.2	993.28	1020.99	1370.69	18	0.05	20.2	-	-	0.05	20.2
p9	-	-	-	-	2.33	31.50	32.31	15.7	0.07	15.7	-	-	0.06	15.7
p10	1182.02	1209.2	1744.95	12.9	775.77	803.84	1692.19	10.1	0.06	16.8	-	-	0.05	16.8
p11	5.56	32.88	37.85	11.8	1302.28	1332.93	2062.62	9.6	0.0	12.5	-	-	0.05	12.5
p12	37.29	47.72	62.02	18.8	2.81	28.77	29.88	18.8	0.04	18.8	-	-	0.07	18.8
p13	-	-	-	-	-	-	-	-	0.03	22	-	-	0.07	22
p14	55.28	62.81	80.5	20.7	942.9	1008.96	1092.89	20	0.03	20.7	-	-	0.07	20.7
p15	-	-	-	-	-	-	-	-	0.06	20.9	-	-	0.07	20.9
p16	-	-	-	-	660.60	693.78	1060.91	25.7	0.11	25.7	-	-	0.12	25.7
p17	-	-	-	-	582.82	603.99	615.12	22.3	0.1	22.3	-	-	0.13	22.3
p18	-	-	-	-	2.26	30.35	30.75	22.8	0.1	22.8	-	-	0.08	22.8
p19	-	-	-	-	-	-	-	-	0.05	26.5	-	-	0.13	26.5
p20	-	-	-	-	8.31	40.08	41.64	24.7	0.08	24.7	-	-	0.1	24.7
o1	-	-	-	-	-	-	-	-	128.14	6	0.03	70	N/A	N/A

Table 2, the time required for MSPLAN to generate a plan when given, by an oracle, the makespan that yields the plan with the best quality was typically much smaller than the time required by MSPLAN to generate a plan when this makespan was not known *a priori*. If an incremental approach could be created, whereby the partial weighted MaxSAT encoding of a problem with makespan  $k$  is extended, instead of being encoded from scratch, to generate the partial weighted MaxSAT encoding with makespan  $k + 1$ , the running time of MSPLAN could likely be improved. A similar incremental approach has recently been investigated for compiling PDDL planning problems into answer set programs [13].

Our planner appeared to be sensitive to the underlying MaxSAT solver. With MiniMaxSAT, MSPLAN was able to solve 12 of the 20 problems in the `pathways` domain but when SAT4J was used instead, 17 of the 20 problems were solved. Most of the `pathways` problem instances which could not be solved by MiniMaxSAT had a relatively large number of simple preferences (usually between 35 to 50 simple preferences). Since partial weighted MaxSAT solvers are believed to be very sensitive to the ratio of soft constraints versus hard constraints [1], the relatively larger proportion of soft constraints in these problems might explain the superior performance of SAT4J in the `pathways` domain.

## 5 Concluding Remarks

In this paper, we characterized the PBP problem as a partial weighted MaxSAT problem. We developed a compact encoding of PBP as partial weighted MaxSAT by building on the success of a SAS<sup>+</sup> based SAT encoding. To the best of our knowledge, this is the first time that the SAS<sup>+</sup> formalism has been used in the context of PBP. Our experimental evaluation showed that our MSPLAN system (with our SAS<sup>+</sup> encoding), consistently outperformed an existing MaxSAT-based planner (with a STRIPS encoding) with respect to running time, while generating plans of comparable quality. Remarkably, when run with two different MaxSAT solvers, MSPLAN generated plans with comparable quality to those generated by state-of-the-art heuristic search planners for PBP. Although the heuristic search planners were generally faster, there were problem instances for which at least one of the two MSPLAN systems ran significantly faster than HPLAN-P. There were also problem instances that could be solved by MSPLAN but not by SGPlan-W. While the horse race between heuristic search planners and SAT-based planners ever continues, the strong performance of MSPLAN serves to illustrate the effectiveness of our SAS<sup>+</sup> encoding and suggests that both MaxSAT and SAS<sup>+</sup> encodings for PBP are worthy of continued exploration.

In this paper, we focused on simple preferences to support comparison with GM. Our approach is extendable to temporally extended preferences following the compilation technique described in [2]. Our approach is also extendable to net benefit problems by a simple extension to the encoding. These are topics of current investigation. A current limitation of MSPLAN is that it is  $k$ -optimal rather than optimal. (SGPlan<sub>5</sub> is neither optimal nor  $k$ -optimal whereas HPLAN-P has the capacity to be  $k$ -optimal with respect to plan length and optimal if

run to completion with certain restrictions on metric functions [3].) Optimality may be achievable with a MaxSAT-based approach by exploiting a translation of soft goals into action costs [12] and a cost-optimal planner based on partial weighted MaxSAT (e.g., [16]).

**Acknowledgements.** We thank Enrico Giunchiglia and Marco Maratea for providing us with their STRIPS-encodings of the IPC-2006 *Simple Preferences* problems. We gratefully acknowledge funding from Microsoft Research, the Natural Sciences and Engineering Research Council of Canada (NSERC), and MITACS.

## References

- [1] Ansótegui, C., Bonet, M.L., Levy, J.: Solving (Weighted) Partial MaxSAT through Satisfiability Testing. In: Kullmann, O. (ed.) SAT 2009. LNCS, vol. 5584, pp. 427–440. Springer, Heidelberg (2009)
- [2] Baier, J.A., Bacchus, F., McIlraith, S.A.: A heuristic search approach to planning with temporally extended preferences. *AIJ* 173(5-6), 593–618 (2009)
- [3] Baier, J.A., McIlraith, S.A.: Planning with preferences. *AI Magazine* 29(4), 25–36 (2008)
- [4] Berre, D.L., Parrain, A.: The SAT4J library, release 2.2. *Journal on Satisfiability, Boolean Modeling and Computation* 7, 59–64 (2010)
- [5] Coles, A.J., Coles, A.: LPRPG-P: Relaxed plan heuristics for planning with preferences. In: ICAPS (2011)
- [6] Gerevini, A., Haslum, P., Long, D., Saetti, A., Dimopoulos, Y.: Deterministic planning in the fifth international planning competition: PDDL3 and experimental evaluation of the planners. *Artificial Intelligence* 173(5-6), 619–668 (2009)
- [7] Giunchiglia, E., Maratea, M.: A pseudo-boolean approach for solving planning problems with IPC simple preferences. In: COPLAS, pp. 23–31 (2010)
- [8] Helmert, M.: The Fast Downward planning system. *JAIR* 26, 191–246 (2006)
- [9] Heras, F., Larrosa, J., Oliveras, A.: MiniMaxSat: A New Weighted Max-SAT Solver. In: Marques-Silva, J., Sakallah, K.A. (eds.) SAT 2007. LNCS, vol. 4501, pp. 41–55. Springer, Heidelberg (2007)
- [10] Hsu, C.-W., Wah, B., Huang, R., Chen, Y.: Constraint partitioning for solving planning problems with trajectory constraints and goal preferences. In: IJCAI, pp. 1924–1929 (2007)
- [11] Huang, R., Chen, Y., Zhang, W.: A novel transition based encoding scheme for planning as satisfiability. In: AAAI, Atlanta, GA, pp. 89–94 (2010)
- [12] Keyder, E., Geffner, H.: Soft goals can be compiled away. *JAIR* 36, 547–556 (2009)
- [13] Knecht, M.: Efficient Domain-Independent Planning using Declarative Programming. Master’s thesis, University of Potsdam (December 2009)
- [14] Maratea, M.: An experimental evaluation of Max-SAT and PB solvers on over-subscription planning problems. In: RCRA, vol. 616 (2010)
- [15] Marques-Silva, J.: The MSUncore MaxSAT solver (2009)
- [16] Robinson, N., Gretton, C., Pham, D.N., Sattar, A.: Partial weighted MaxSAT for optimal planning. In: PRICAI, Daegu, Korea, pp. 231–243 (2010)