

Better Time Constrained Search via Randomization and Postprocessing

Fan Xie and Richard Valenzano and Martin Müller

Computing Science, University of Alberta
Edmonton, Canada
{fxie2, valenzan, mmueller}@ualberta.ca

Abstract

Most of the satisficing planners which are based on heuristic search iteratively improve their solution quality through an anytime approach. Typically, the lowest-cost solution found so far is used to constrain the search. This avoids areas of the state space which cannot directly lead to lower cost solutions. However, in conjunction with a post-processing plan improvement system such as ARAS, this bounding approach can harm a planner's performance.

The new anytime search framework of *Diverse Any-Time Search* avoids this behaviour by taking advantage of the fact that post-processing can often improve a lower quality input plan to a superior final plan. The framework encourages diversity of "raw" plans by restarting and using randomization, and does not use previous solutions for bounding. This gives a post-processing system a more diverse set of plans to work on, which improves performance. When adding both *Diverse Any-Time Search* and the ARAS post-processor to LAMA-2011, the winner of the most recent IPC planning competition, and AEES, the Anytime Explicit Estimation Algorithm, the performance on the 550 IPC 2008 and IPC 2011 problems is improved by almost 60 points according to the IPC metric, from 511 to over 570 on LAMA-2011, and 73 points from 440 to over 513 on AEES.

Introduction

Since IPC-2008, the satisficing planning community has been using the IPC scoring function to evaluate planners. This function emphasizes both plan quality and coverage simultaneously. Many satisficing planners such as LAMA (Richter and Westphal 2010) and Fast Downward (Helmert 2006) use an *anytime approach*: they attempt to quickly find an initial plan of possibly low quality, then use the remaining time to improve upon this plan. *Post-processing*, as implemented in the ARAS system (Nakhost and Müller 2010), is another recent plan quality improvement technique. This approach takes an existing valid plan as input and tries to improve it by removing unnecessary actions and by finding shortcuts with a local search. Another *post-processing* technique, discussed in (Chrpa, McCluskey, and Osborne 2012), analyzes action dependencies and independencies in order to identify redundant actions or non-optimal sub-plans.

Copyright © 2013, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Since post-processing systems decouple plan improvement from plan discovery, they can easily be applied to any planner once it finds a solution. Originally, the ARAS post-processor was applied as the final step of the planning process, after the planner had completed or was considered unlikely to find a better plan. However, it is also possible to use a post-processor in anytime fashion, by running it with a relatively tight time or memory limit on *every* plan produced by an anytime planner. This approach has been used by several recent planners (Nakhost et al. 2011; Valenzano et al. 2012; Xie, Nakhost, and Müller 2012). Experience shows that there is large variance in the amount of improvement achieved by post-processing. A lower quality input plan sometimes yields a higher quality final plan.

This observation has an impact on the design of planning systems which use post-processing. Most anytime satisficing planners use the cost of the best incumbent solution to bound their future search. This avoids wasting effort in areas of the state space that cannot *directly* lead to a better solution. However, we will show below that it has a negative impact on the effectiveness of post-processing. Such bounding greatly decreases the number and variety of plans that an anytime system finds, since it makes it much more difficult to find such new plans.

The main contributions of this paper are as follows:

- We introduce the concept of *unproductive time*, which measures the amount of time after the best solution is found, to help explain this trade-off.
- We present evidence that bounding in an anytime system is detrimental when used in conjunction with a post-processing system.
- We develop the meta-algorithm *Diverse Any-time Search* (DAS), which uses restarting to generate a more diverse set of plans.
- We implement DAS on top of the Fast Downward code base (Helmert 2006), and demonstrate significant plan quality improvements for two recent planning algorithms: the state-of-the-art planner LAMA-2011 (Richter and Westphal 2010) and AEES (Thayer, Benton, and Helmert 2012).
- We show that the improvements from DAS and from post-processing are independent, and can even be synergistic: With LAMA, the improvement from using both

techniques together is slightly larger than the sum of the improvements when applied individually.

The remainder of this paper is organized as follows: after introducing the concept of *unproductive time* and measuring it in LAMA for recent IPC planning problems, the new meta-algorithm DAS (Diverse Any-time Search) is introduced. DAS is tested both with and without the ARAS post-processor on LAMA-2011 and AEES. The experimental results show strong improvements in plan quality on IPC-2008 and IPC-2011 domains.

Unproductive Time in Any-time Satisficing Planning

As mentioned above, most state-of-the-art planners use an anytime strategy to improve solution quality over time. Maximizing the efficient usage of the available planning time is important for plan quality. However, there has been little investigation into how much time is actually being used for finding the final solution. Let *unproductive time* be defined as the amount of time remaining, out of the total time given, when the planner's best solution is found. For example, if an anytime planner A finds its best solution on a planning instance B at 13 minutes given a 30 minutes time limit, and A does not improve upon this plan in the remaining 17 minutes, then the unproductive time for planner A on problem B is 17 minutes. The amount of unproductive time can be used to evaluate how efficiently an anytime planner is using the given search time, since that unproductive time could be spent doing something useful, such as plan post-processing. One of the consequences of using bounding in an anytime system such as LAMA-2011 is that it often leads to huge amounts of unproductive time.

LAMA-2011 is a state-of-the-art planner that has been shown to achieve both high coverage and strong solution quality: it won the sequential satisficing track of the International Planning Competition in 2011 (IPC 2011) after, in its previous incarnation as LAMA-2008, winning IPC 2008. LAMA's high coverage is achieved through techniques such as the use of multiple heuristics (Richter, Helmert, and Westphal 2008), preferred operators (Richter and Helmert 2009), and greedy best-first search (Bonet and Geffner 2001). LAMA-2011 starts its search with two runs of greedy best-first search, first with a distance-to-go heuristic and then with cost-to-go. Next, LAMA improves the quality of its solutions through the anytime procedure of Restarting Weighted A* (RWA*) (Richter, Thayer, and Ruml 2010). This procedure starts a new WA* search with a lower weight w whenever a new best solution is found. Only cost-to-go heuristics are used in this phase.

Whenever a new best solution with cost C is found, this cost is used to *bound* the rest of the search. Only nodes with g -cost (cost of best known path to the node) less than C are added to the open list. This prunes states that cannot lead directly to an improving solution.

Figure 1 shows that this approach leads to a very large fraction of unproductive time on IPC benchmarks. Among the total of 244 problems solved in IPC-2011 with an 1800 second (30 minute) time limit, in more than 45% (111) of the

problems, LAMA-2011 is unproductive for more than 1700 seconds. Table 1 shows the amount of unproductive time separately for each IPC-2011 domain. In the four domains of 2011-barman, 2011-elevators, 2011-parcprinter and 2011-woodworking, unproductive time exceeds 90%. In these domains, the planner quickly finds an initial solution, but fails to improve it.

As a typical example, Table 2 shows the number of solutions and the amount of unproductive time for all 20 instances of the 2011-elevators domain. With the exception of elevator04, LAMA-2011 finds only a single solution to each problem. This does not at all imply that the first solution found by LAMA-2011 is optimal (**better solutions can be found after applying post-processors for these problems**). It is simply much more difficult to find a second solution when using cost-to-go heuristics and the bound from the first solution, than it was to generate the initial solution using distance-to-go heuristics and no bound.

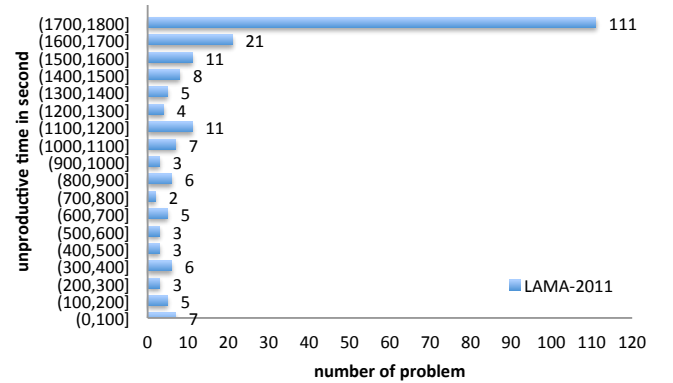


Figure 1: *Unproductive Time* of LAMA-2011 on IPC-2011

Domain	solved	total time	UT	percentage
parcprinter	20	36000	35997	99.99%
barman	20	36000	35901	99.73%
woodworking	20	36000	35357	98.21%
elevators	20	36000	33890	94.14%
visitall	20	36000	31280	86.89%
pegsol	20	36000	31107	86.41%
scanalyzer	20	36000	29844	82.90%
nomystery	10	18000	14682	81.57%
transport	15	27000	20554	76.13%
openstacks	20	36000	24902	69.17%
floortile	6	10800	7032	65.11%
tidybot	16	28800	18099	62.84%
sokoban	19	34200	20377	59.58%
parking	18	32400	13948	43.05%

Table 1: *Unproductive Time* (UT) of LAMA-2011 on different IPC-2011 domains.

Post-Processing with ARAS

Even a state-of-the-art planner such as LAMA-2011 spends the large majority of its execution time being unproductive.

Instance	#s	UT	Instance	#s	UT
01	1	1799	11	1	1760
02	1	1798	12	1	1704
03	1	1797	13	1	1679
04	3	1794	14	1	1725
05	1	1799	15	1	1659
06	1	1796	16	1	1649
07	1	1798	17	1	1576
08	1	1791	18	1	1510
09	1	1789	19	1	1152
10	1	1783	20	1	1531

Table 2: Number of solutions (#s) and *Unproductive Time* (UT) of LAMA-2011 in the 20 instances of 2011-elevators.

How can we use this time more effectively?

One possibility is to feed the solutions found by a planner into a post-processing plan improvement system such as ARAS (Nakhost and Müller 2010). ARAS consists of two components. The first, Action Elimination (AE), examines the plan for unnecessary actions. This involves scanning the plan from front to back, removing each action and its dependent actions in turn, and testing if the resulting plan is still valid and reaches a goal. If so, all the unnecessary actions are discarded and the scan continues on the resulting shorter plan.

The second and main component of ARAS is Plan Neighbourhood Graph Search (PNGS). This technique builds a neighbourhood graph in the state space around the trajectory of the input plan by performing a breadth-first search to some depth d . Once this plan neighbourhood is constructed, a lowest-cost plan from the start to a goal state is extracted from this graph. Intuitively, the algorithm identifies local shortcuts along the path.

In Anytime ARAS, AE and PNGS are alternated, starting with AE, and running PNGS with an increasing depth bound at each iteration until some time or memory limit is reached.

When ARAS was first introduced, it was used in a *process-once* fashion. This involved splitting the available planning time into two phases: plan finding and plan post-processing. In the first phase, a planner is used to find some plan (or a series of plans). When the plan finding phase is up, the best plan found so far is handed to a post-processor which is then allotted all of the remaining time.

Though ARAS was initially proposed for use in this process-once fashion, several planners which competed in IPC-2011 (Nakhost et al. 2011; Valenzano et al. 2012) used this system in an anytime fashion. This involves interleaving the plan finding and plan post-processing phase by using the post-processor in between iterations of an anytime planner. Each new plan found is immediately post-processed by ARAS. Once ARAS reaches some time or memory limit, it returns the best solution it has found and allows the anytime planner to continue looking for more plans.

There are two motivations for using ARAS in this way. First, for most anytime satisficing planners, it is difficult to determine if any new plans will be found during the remaining time or if the planner will be unproductive. Sec-

ond, even if the anytime planner can find a better plan with more search, there is no guarantee that post-processing the best plan found by an anytime planner will have a lower cost than the post-processed version of a weaker plan found earlier. The following experiment illustrates this. There are 151 problems from IPC-2011 for which LAMA-2011 generates more than one plan (without any assistance from a post-processor). If ARAS is run with a 2 GB memory limit on all these plans, then on 44 (29%) of these problems, the lowest cost solution is generated by post-processing an earlier plan, not from the final and best input plan found by LAMA-2011. The instance 2011-transport07 is typical: here, LAMA-2011 generates two plans of cost 8396 and 7222. Post-processing with ARAS improves the weaker input plan to 6379, and the stronger one to 6402. The weaker initial plan leads to a better final result.

For post-processing, though it is important to get high quality input plans, when they are becoming very expensive to find, we should change to find more lower quality input plans, which might also contribute to the final quality improvement. If the anytime planner is only able to generate a small number of similar plans, then ARAS can only search within a very restricted set of neighbourhoods. This likely means that all its output plans will be similar, and of similar quality. When ARAS is handed a larger and more diverse set of input plans, it has a greater chance of finding significant improvements for at least one of them.

The Diverse Any-time Search Meta-Algorithm

The *Diverse Any-time Search (DAS)* meta-algorithm, shown in Algorithm 1, uses restarts with no bounds and post-processing in order to improve a given anytime planner. DAS divides the given total planning time into N equal time segments, where N is a user-supplied constant. In the first segment, the anytime planner P is run normally, except that ARAS is used to post-process each plan generated by P . At any time during this first segment, P can use its best plan found so far, *excluding* post-processing, to bound its search.

When time runs out on the first time segment, the best plan found so far, *including* post-processing, is saved. The anytime planner is restarted without any knowledge of previous solutions, and with a planner-specific randomization which will vary the planner’s search. At the end of each search segment, the best overall plan is updated, and P restarts from scratch with a new random seed.

By restarting from scratch with the early, greedier iterations of P , which typically find plans much more frequently, DAS increases the number of plans available to the post-processor, thereby increasing the number of opportunities for finding a strong improvement.

If P cannot find any solution by the end of the first time segment, then it is kept running, without restarts. In this case, finding even a single solution is deemed to be hard, and restarts may hurt the planner¹. Algorithm 2 shows details.

¹By taking this approach, we are ensuring that the coverage of the restarting planners is the same as the baseline planners. As our focus is quality, it is easier to analyze the difference in plan quality between planners by taking this approach.

The algorithm uses two time limits: a soft time limit t for each segment, used for restarting if at least one solution was found, and a hard time limit T for the whole search. If in the first segment, P cannot find a solution within the soft time limit t , then restarts are disabled. In Algorithm 2, in the assignment $\langle p1, t1 \rangle \leftarrow P.search(conf, bound, rand, time)$, $conf$ is the current search configuration (such as the weight for RWA*), $bound$ is the plan quality bound, and $rand$ is the random seed. If the search finds a solution within time limit $time$, it returns the plan $p1$ and time used $t1$. Otherwise, the search terminates with $p1 = NULL$ and $t1 = time$.

The expression $\langle p2, t2 \rangle \leftarrow PP.process(p1, time)$ indicates a call to the post-processor PP with $p1$ as input plan. The post-processor PP returns when it either reaches the time limit $time$ or a pre-set memory limit. When PP terminates because of the time limit, it returns the best plan found (could be the input plan $p1$) and $t2 == time$. When it terminates because of the memory limit, it returns the best plan found (also could be the input plan $p1$) and the time used.

Algorithm 1 Diverse Any-time Search

Input Initial State I , goal condition G given search time T , planner P , post-processor PP

Parameter N

Output A solution plan

```

 $t \leftarrow T/N$ 
 $\langle plan_{best}, cost_{best} \rangle \leftarrow \langle [], \infty \rangle$ 
for  $i$  from 1 to  $N$  do
   $rand \leftarrow generate\_random\_seed()$ 
   $isSolved \leftarrow i == 1$ 
   $plan \leftarrow AnytimeSearchWithPostprocessing$ 
    ( $I, G, T, t, rand, P, PP, isSolved$ )
  if  $cost(plan) < cost_{best}$  then
     $\langle plan_{best}, cost_{best} \rangle \leftarrow \langle plan, cost(plan) \rangle$ 
  end if
end for
return  $plan_{best}$ 

```

Experiments

Experiments include tests on all 550 problems from IPC-2008 and IPC-2011, and were run on an 8 core 2.8 GHz machine with a time limit of 30 minutes and memory limit of 4 GB per problem. Results for planners which use randomization are averaged over 5 runs. All planners are implemented on the same version of the Fast Downward code base (Helmert 2006). The translation from PDDL to SAS+ was preprocessed by Fast Downward, and that time is not included here. The scores shown use the IPC metric with LAMA-2011 as a baseline. If L is the plan computed by LAMA-2011, then the score of a given plan P is calculated by $cost(L)/cost(P)$. We use the ARAS system as the post-processor in our experiments. ARAS is given a memory limit of 2 GB, and the time limit of running ARAS is controlled by the DAS meta-algorithm as explained in last section.

Algorithm 2 Anytime Search with Post-processing

Input Initial State I , goal condition G , hard timelimit T , soft timelimit t , random seed $rand$, planner P , post-processor PP , first solution found flag $solved$

Parameter N

Output A solution plan

```

 $conf \leftarrow P.GetFirstConf()$ 
 $bound \leftarrow \infty$ 
 $plan_{best} \leftarrow []$ 
 $isSolved \leftarrow solved$ 
 $totalTime \leftarrow 0$ 
 $restart \leftarrow true$ 
while have time do
  if not  $isSolved$  or not  $restart$  then
     $time \leftarrow T - totalTime$ 
     $\langle p1, t1 \rangle \leftarrow P.search(conf, bound, rand, time)$ 
    if not  $isSolved$  and  $t1 > t$  then
       $restart \leftarrow false$ 
    end if
  else
     $time \leftarrow t - totalTime$ 
     $\langle p1, t1 \rangle \leftarrow P.search(conf, bound, rand, time)$ 
  end if
  if  $p1 == []$  then
    return  $[]$ 
  end if
   $isSolved \leftarrow true$ 
   $totalTime \leftarrow totalTime + t1$ 
   $conf \leftarrow P.nextConf(conf)$ 
   $bound \leftarrow cost(p1)$ 
  if  $restart$  then
     $time \leftarrow t - totalTime$ 
  else
     $time \leftarrow T - totalTime$ 
  end if
   $\langle p2, t2 \rangle \leftarrow PP.process(p1, time)$ 
   $totalTime \leftarrow totalTime + t2$ 
  if  $cost(p2) < cost(plan_{best})$  then
     $plan_{best} \leftarrow p2$ 
  end if
end while
return  $plan_{best}$ 

```

Experiment 1: Evaluating Anytime ARAS and Bounding with LAMA

Table 3 compares three planners on IPC-2011 domains:

- **LAMA-2011** is the IPC-2011 version of LAMA.
- **LAMA-Aras** is an implementation of *Diverse Any-time Search (DAS)* with the input planner *LAMA-2011*, the input post-processor *Any-time ARAS* and Parameter $N=1$. The improved plans by Any-time ARAS are *not* used for bounding the following search, but LAMA-2011 *still* does its own bounding internally. Any-time ARAS is given 2 GB of memory.
- **LAMA-Aras-B** is like LAMA-Aras. However, all plans,

Domain	LAMA-2011	LAMA-Aras	LAMA-Aras-B
barman	20	23.99	23.99
elevators	20	26.01	25.87
floortile	6	6.77	6.77
nomystery	10	10.00	9.83
openstacks	20	19.98	19.89
parcprinter	20	20.10	19.78
parking	18	18.93	17.46
pegsol	20	20.00	20.00
scanalyzer	20	23.35	21.26
sokoban	19	20.23	19.05
tidybot	16	16.77	16.77
transport	15	17.70	16.38
visitall	20	20.45	20.37
woodworking	20	20.96	20.85
Total	244	265.24	258.27

Table 3: Plan Quality of LAMA-2011, LAMA-Aras and LAMA-Aras-B on IPC-2011.

including the improved plans found by Any-time ARAS , are used to bound subsequent LAMA searches.

Table 3 demonstrates that not using Aras-improved plans for bounding LAMA is clearly the better choice, since LAMA-Aras dominates LAMA-Aras-B in almost all domains. Among the three planners, LAMA-Aras almost always gets the best score, except by a small margin in openstacks and tidybot. ARAS is not useful in the openstacks domain (Nakhost and Müller 2010), and the time wasted running it leads causes a slight decrease in plan quality there.

Experiment 2: Evaluating Diverse Any-Time Search with no Postprocessing

This section examines the impact of using Diverse Any-time Search in terms of unproductive time and the number of solutions it generates when added to LAMA-2011. These tests do not use any post-processing system. They show that the new meta-algorithm increases the number of plans found and even improves solution quality in a number of domains.

When Diverse Any-time Search is added to LAMA-2011, RWA* is being used within each time segment (as is bounding), though the algorithm begins from scratch with a greedy best-first search iteration at the beginning of each time segment, without any bounding from previous time segments. The source of diversity is *random operator ordering* (Valenzano et al. 2012). This involves randomly shuffling the order of the generated set of successors of an expanded node before they are added to the open list. Random operator ordering affects the search by changing the way that ties are broken. However, so as to maintain the same coverage between competing algorithms, we use the default operator ordering during the first time segment.

We refer to the new planner as **Diverse-LAMA(N)** where N is the parameter which affects the length of the time segments. Table 4 shows that this planner generates many more plans² and has much less unproductive time than the regu-

²here we only show one run result instead of the average results

lar version of LAMA-2011 on the 2011-elevators domain. In particular:

- The average number of plans increases from 1.1 to 4.3. On 17 of the problems, we see an increase from 1 plan with standard LAMA-2011 to 4 plans with Diverse-LAMA(4), since Diverse-LAMA(4) finds one plan per segment. In the cases of elev02 and elev03, Diverse-LAMA(4) sometimes finds more than 1 plan per segment, depending on the random seed, whereas there is a segment in which no plans are found on elev05.
- The amount of unproductive time decreases in all but one instance (elev11), often drastically. The problem elev11 is the only exception as **Diverse-LAMA(4)** is unable to improve the plan it finds during the first segment when it finds the same plan as LAMA-2011. However, in all other problems there was at least one plan found in a later segment that was better than the initial plan found in the first plan.

Due to LAMA-2011’s high amount of unproductive time in this domain, Diverse-LAMA(4) is also able to find better solutions in this domain. This is because LAMA-2011 rarely finds a new solution after the first $1800/4 = 450$ seconds. In contrast, Diverse-LAMA(4) continues to find solutions by restarting and returning to a greedier search. Sometimes, these are better than the best solution found in the first 450 seconds.

Table 5 compares the plan quality of LAMA-2011 and Diverse-LAMA(4) on IPC-2011 domains and shows that this behaviour is also not limited to the elevators domain. In total, Diverse-LAMA(4) improves by a score of 6.1 though this improvement is not uniform over all domains. Instead, Diverse-LAMA(4) improves its solution quality over LAMA-2011 in 8 domains, while it is worse in 5 domains. These improvements are mainly made in domains in which LAMA-2011 has a high percentage of *unproductive time* as was explained for the elevators domain. In contrast, in those domains in which LAMA-2011 is more productive later in the search, the restarts are preventing Diverse-LAMA(4) following through on one search long enough in order to find the best solutions. This is more apparent in Table 6, which shows the number of problems on which each of LAMA-2011 and Diverse-LAMA(4) found the best plan. In those domains in which LAMA-2011 is mostly unproductive, Diverse-LAMA(4) rarely generates worse final solutions, while for those domains in which LAMA-2011 is more productive later on — such as Floortile, Tidybot, Sokoban and Parking — Diverse-LAMA(4) will occasionally find weaker plans.

Experiment 3: Combining Diverse Any-time Search with ARAS

This section tests the DAS meta-algorithm with LAMA and ARAS. The system is named Diverse-LAMA-Aras(N). The four planners **LAMA-2011**, **LAMA-2011-Aras**, **Diverse-LAMA(4)**, and **Diverse-LAMA-Aras(4)** are tested on all

in 5 runs.

Problem	N(LAMA)	UT(LAMA)	N(DL(4))	UT(DL(4))
elev01	1	1799	4	1350
elev02	1	1798	6	330
elev03	1	1797	4	1349
elev04	3	1794	9	835
elev05	1	1799	3	900
elev06	1	1796	4	446
elev07	1	1798	4	448
elev08	1	1791	4	895
elev09	1	1789	4	1342
elev10	1	1783	4	440
elev11	1	1760	4	1762
elev12	1	1704	4	863
elev13	1	1679	4	354
elev14	1	1725	4	793
elev15	1	1659	4	713
elev16	1	1649	4	795
elev17	1	1576	4	1262
elev18	1	1510	4	294
elev19	1	1152	4	1009
elev20	1	1531	4	296

Table 4: Number of solutions (**N**(0)) and *Unproductive Time* (**UT**(0)) of planners LAMA-2011 (**LAMA**) and Diverse-LAMA(4) (**DL**(4)).

domain	UT	LAMA-2011	Diverse-LAMA(4)
2011-parcprinter	99.99%	20	20.08
2011-barman	99.73%	20	21.76
2011-woodworking	98.21%	20	20.48
2011-elevators	94.14%	20	25.20
2011-visitall	86.89%	20	20.10
2011-pegasol	86.41%	20	19.79
2011-scanalyzer	82.90%	20	20.75
2011-nomystery	81.57%	10	10
2011-transport	76.13%	15	15.69
2011-openstacks	69.17%	20	20.22
2011-floortile	65.11%	6	5.05
2011-tidybot	62.84%	16	15.30
2011-sokoban	59.58%	19	18.59
2011-parking	43.05%	18	17.21
total		244	250.10

Table 5: Plan Quality of LAMA-2011, Diverse-LAMA(4) on IPC-2011. Domains are sorted by decreasing fraction of *Unproductive time* (**UT**) shown in Table 1.

550 problems from IPC-2008 and IPC-2011, the two IPC competitions which emphasize plan quality.

Table 7 shows the plan quality comparison of the 4 planners in each domain. Diverse-LAMA-Aras(4) achieves the highest overall score, improving the baseline planner LAMA-2011 by **59** units, and is best in 18 out of 23 domains.

Diverse-LAMA(4) and Diverse-LAMA-Aras(4) mainly restart based on elapsed time. Figure 2 shows the normalized score curve over 30 minutes of the 4 tested planners over all IPC-2008 and IPC-2011 domains. The three vertical lines indicate the restart points for Diverse-LAMA(4) and

domain	UP	better	worse	total
2011-parcprinter	99.99%	2	0	20
2011-barman	99.73%	19	0	20
2011-woodworking	98.21%	8	0	20
2011-elevators	94.14%	19	0	20
2011-visitall	86.89%	6	3	20
2011-pegasol	86.41%	0	1	20
2011-scanalyzer	82.90%	6	2	20
2011-nomystery	81.57%	0	0	10
2011-transport	76.13%	6	0	15
2011-openstacks	69.17%	4	4	20
2011-floortile	65.11%	0	2	6
2011-tidybot	62.84%	2	7	16
2011-sokoban	59.58%	1	4	19
2011-parking	43.05%	7	4	18

Table 6: Plan Comparison between Diverse-LAMA(4) (one run result) and LAMA-2011 on different domains. The columns **better** indicates in how many problems Diverse-LAMA(4) generates better plans than LAMA-2011 (**worse** means how many worse). Domains are ordered according to the percentages of *Unproductive time* (**UP**) shown in Table 1.

Diverse-LAMA-Aras(4) at 450, 900 and 1350 seconds.

Before the first restart, DAS and non-DAS versions of the same planner are nearly the same³. Immediately after the first restart, the DAS planners show a quick jump in solution quality. For many problems, restarting quickly finds new, better solutions. The second restart provides a less pronounced but still visible gain.

By producing more plans, Diverse-LAMA-Aras(4) also provides more input plans for ARAS. Compared to Diverse-LAMA-Aras(4), the improvement from the first restart is more pronounced in Diverse-LAMA(4). Using ARAS smoothes out some of the variance in solution quality between different runs.

As shown by the previous two experiments, using either ARAS or DAS improves plan quality. The next experiment shows that these improvements are independent, and even synergetic, in the case of LAMA. Table 8 contains details of experiments to support this claim. The score of each system is split into two components: the raw scores of the best plans produced by the baseline planner as part of the combined system, and the improvement achieved through running Aras. For comparison, the baseline planner’s scores are also shown. The raw scores are a little worse than the baseline scores, since the planner has less time when used in conjunction with Aras. The improvement of Aras over Diverse-LAMA(4) is even slightly larger than the improvement of Aras over LAMA-2011. This demonstrates that the improvements due to the DAS enhancement, from Diverse-LAMA(4) finding substantially different overall plans seems to be largely independent from the local plan improvements

³To fully utilize our computational resources, we run several processes simultaneously on a multi-core machine. While all versions use the same memory limit, the restarts cause DAS to use less memory. The resulting decrease in memory contention accounts for the small differences in planner performance.

found by Aras. The following two examples illustrate this point:

- ARAS can help DAS in cases where restarting prevents the search from running long enough to find a good plan. In 2011-floor tile 05, the best raw plan generated by Diverse-LAMA(4) has cost 132, while LAMA-2011 can find a cost 63 plan. ARAS can improve the cost 132 plan to a cost of 63 as well.
- Diverse Search provides more input plans for ARAS. In 2011-woodworking 01, plans of cost 1600, 1630 and 1620 are produced in time segments 1, 3 and 4, while LAMA-2011 only finds one solution of cost 1600. ARAS improves the three solutions as follows: 1600 \rightarrow 1460, 1630 \rightarrow 1290 and 1620 \rightarrow 1380. The worst input plan is easiest to improve, while the best input plan is poorest after post-processing. Out of the 244 problem instances solved by LAMA-2011, in 44 cases the best final plan produced from Anytime Aras does not come from Lama's best plan. In Diverse-LAMA-Aras(4), this ratio increases dramatically, to 86/244 **problems**.

domain	LAMA	DL(4)	LAMA-Aras	DL-Aras(4)
08-cybersec	30	30.00	30.00	30.00
08-elevators	30	35.82	38.50	43.34
08-openstacks	30	30.25	29.97	30.35
08-parcprinter	30	30.00	30.09	30.10
08-pegssol	30	29.78	30.00	30.00
08-scanalyzer	30	31.35	34.00	34.18
08-sokoban	28	27.15	27.66	27.48
08-transport	29	31.51	35.14	36.73
08-woodworking	30	30.92	33.10	34.28
11-barman	20	21.76	23.99	24.20
11-elevators	20	25.20	26.01	31.16
11-floor tile	6	5.01	6.77	6.77
11-nomystery	10	10.00	10.00	9.89
11-openstacks	20	20.22	19.98	20.11
11-parcprinter	20	20.08	20.10	20.05
11-parking	18	17.21	18.93	19.54
11-pegssol	20	19.79	20.00	20.01
11-scanalyzer	20	20.75	23.35	23.46
11-sokoban	19	18.59	20.23	20.54
11-tidybot	16	15.21	16.77	16.17
11-transport	15	15.69	17.70	19.68
11-visitall	20	20.10	20.45	20.53
11-woodworking	20	20.48	20.96	21.78
total	511	526.89	553.69	570.35

Table 7: Plan Quality of LAMA-2011 (LAMA), LAMA-2011-Aras (**LAMA-Aras**), Diverse-LAMA(4) (**DL(4)**) and Diverse-LAMA-Aras(4) (**DL-Aras(4)**) on all 550 problems from IPC-2008 and IPC-2011.

Experiment 4: Parameter Test for Number of DAS Phases

Figure 3 shows the performance of DAS when varying the parameter N , the number of phases. $N = 1$ corresponds to LAMA-ARAS, and $N = 4$ to the algorithm used in the




Figure 2: Normalized Score Curve of the 4 tested planner

previous experiments. The graph shows results of Diverse-LAMA-ARAS(N) on the IPC-2011 domains, varying N in the range from 1 to 120. Overall, the score differences are small, with the best results for N from 3 to 6. For $N \leq 3$, plan quality increases with N , taking advantage of the diversity and number of plans generated. For $N > 6$, the solution quality slowly decreases as the runtime for both Lama and ARAS becomes ever shorter.⁴ **It also raises the trade-off we discussed before in selecting a good N : 1), a too big N causes that planners don't have enough search time to find high quality input plans; 2), a too small N causes that planners are not able to generate more diverse plans.**

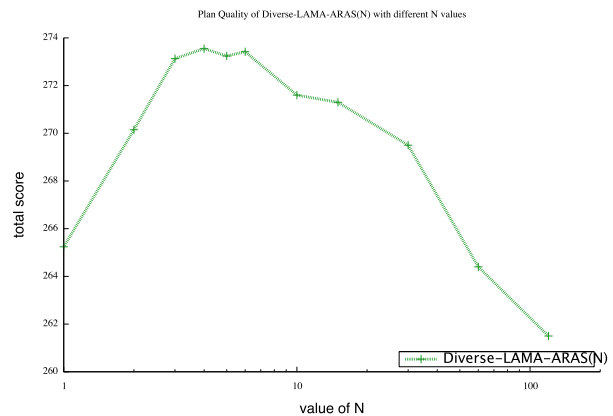


Figure 3: Plan Quality of Diverse-LAMA-Aras(N) with different N values.

Experiment 5: Applying DAS to Anytime Explicit Estimation Search

Anytime explicit estimation search (AEES) is an anytime search algorithm introduced by (Thayer, Benton, and Helmert 2012). AEES uses explicit estimation search (EES)

⁴The time of each time segment equals to T/N , where T is the total time (30 mins).

domain	LAMA	DL(4)	LAMA-Aras			DL-ARAS(4)		
			raw _{LAMA}	Δ_{ARAS}	Final	raw _{DL}	Δ_{ARAS}	Final
2011-barman	20.00	21.76	20.00	3.99	23.99	21.70	2.51	24.20
2011-elevators	20.00	25.20	20.21	5.80	26.01	24.82	6.35	31.16
2011-floortile	6.00	5.01	5.21	1.55	6.77	4.67	2.10	6.77
2011-nomystery	10.00	10.00	10.00	0.00	10.00	9.86	0.03	9.89
2011-openstacks	20.00	20.22	19.98	0.00	19.98	19.94	0.17	20.11
2011-parcprinter	20.00	20.08	20.00	0.10	20.10	20.00	0.05	20.05
2011-parking	18.00	17.21	16.84	2.09	18.93	16.33	3.21	19.54
2011-pegasol	20.00	19.79	19.57	0.43	20.00	17.65	2.36	20.01
2011-scanalyzer	20.00	20.75	18.89	4.46	23.35	20.39	3.07	23.46
2011-sokoban	19.00	18.59	16.58	3.65	20.23	16.10	4.44	20.54
2011-tidybot	16.00	15.21	15.15	1.62	16.77	14.75	1.42	16.17
2011-transport	15.00	15.69	14.00	3.70	17.70	16.32	3.37	19.68
2011-visitall	20.00	20.10	20.00	0.45	20.45	20.06	0.47	20.53
2011-woodworking	20.00	20.48	20.00	0.96	20.96	20.48	1.30	21.78
Total	244.00	250.10	236.44	28.80	265.24	243.06	30.84	273.90

Table 8: Combined effect of DAS and post-processing.

(Thayer and Ruml 2011) as its main search component. EES is a bounded sub-optimal best first search algorithm. EES only expands nodes whose cost is at most the sub-optimality bound. It expands these nodes in best-first order according to first the inadmissible distance-to-go and then the inadmissible cost-to-go heuristic. Therefore, EES can take advantage of both cost-to-go and distance-to-go heuristics. AEES is an anytime version of EES, which lowers the sub-optimality bound whenever a new best solution is found, and keeps searching with this bound.

The AEES algorithm’s goal is to minimize the time between solutions, and generate more solutions. This makes it a good test case for DAS.

We repeat the same set of IPC experiments, running DAS with AEES configured as follows: it uses the two planning-specific enhancements of deferred evaluation and preferred operators, and the three heuristics Landmark-cut (admissible cost-to-go heuristic) (Helmert and Domshlak 2009), FF-cost (inadmissible cost-to-go heuristic) and FF-distance (distance-to-go heuristic) (Hoffmann and Nebel 2001). The scores shown use the IPC metric with AEES as a baseline. If L is the plan computed by AEES, then the score of a given plan P is calculated by $cost(L)/cost(P)$. The experimental results are shown in Table 9. Similar to the LAMA-2011 experiments, Diverse-AEES-ARAS(4) gets the highest score, improving the baseline planner AEES by 73.7 units from 440 to 513.7, and achieving the best score in 14/23 domains.

Conclusions and Future Work

The search performance of the current state of the art planner LAMA-2011 suffers from a large amount of *unproductive time*. These time can be safely used for improving plan quality, the new meta-algorithm of DAS (Diverse Any-time Search) tries to utilize this unproductive time with randomized restarts. This generates more, and more diverse, plans that post-processing systems such as ARAS can improve upon. Experimental results show significant improvements over the state of the art on IPC-2008 and IPC-2011 domains for both LAMA-2011 and the AEES algorithm.

domain	AEES	DE(4)	AEES-Aras	DE-Aras(4)
08-cybersec	29	31.20	29.00	32.36
08-elevators	30	33.80	41.67	45.20
08-openstacks	30	31.35	30.00	31.15
08-parcprinter	25	25.16	25.70	25.82
08-pegasol	30	29.96	30.11	30.05
08-scanalyzer	30	30.53	34.35	34.16
08-sokoban	27	26.47	26.71	26.73
08-transport	28	31.09	38.43	40.86
08-woodworking	20	20.25	21.14	21.32
11-barman	20	20.88	22.74	22.85
11-elevators	19	23.03	25.03	28.06
11-floortile	6	5.50	6.00	6.00
11-nomystery	10	9.94	10.00	9.91
11-openstacks	20	20.92	20.00	21.00
11-parcprinter	11	11.14	11.88	11.92
11-parking	15	15.15	16.61	18.40
11-pegasol	20	19.92	20.11	20.04
11-scanalyzer	20	21.00	24.73	24.43
11-sokoban	17	16.73	16.60	16.96
11-tidybot	13	13.56	16.21	15.41
11-transport	13	14.60	18.22	19.51
11-visitall	3	3.32	7.30	7.39
11-woodworking	4	3.98	4.12	4.13
total	440	459.50	496.67	513.67

Table 9: Plan Quality of AEES (**AEES**), Diverse-AEES(4) (**DE(4)**), AEES-Aras (**AEES-Aras**) and Diverse-AEES-Aras(4) (**DE-Aras(4)**) on all 550 problems from IPC-2008 and IPC-2011.

The best parameter N for DAS depends on factors such as the planning domain, randomizing method and search algorithm. However, in the experiments the performance was robust for small values of N between 3 and 6. One interesting future work is to automatically tune N using information on the search so far.

References

- Bonet, B., and Geffner, H. 2001. Heuristic search planner 2.0. *AI Magazine* 22(3):77–80.
- Brafman, R. I.; Geffner, H.; Hoffmann, J.; and Kautz, H. A., eds. 2010. *Proceedings of the 20th International Conference on Automated Planning and Scheduling, ICAPS 2010, Toronto, Ontario, Canada, May 12-16, 2010*. AAAI.
- Chrapa, L.; McCluskey, T.; and Osborne, H. 2012. Optimizing plans through analysis of action dependencies and independencies. In McCluskey et al. (2012), 338–342.
- Gerevini, A.; Howe, A. E.; Cesta, A.; and Refanidis, I., eds. 2009. *Proceedings of the 19th International Conference on Automated Planning and Scheduling, ICAPS 2009, Thessaloniki, Greece, September 19-23, 2009*. AAAI.
- Helmert, M., and Domshlak, C. 2009. Landmarks, critical paths and abstractions: What’s the difference anyway? In Gerevini et al. (2009), 162–169.
- Helmert, M. 2006. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.
- Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.
- McCluskey, L.; Williams, B.; Silva, J. R.; and Bonet, B., eds. 2012. *Proceedings of the Twenty-Second International Conference on Automated Planning and Scheduling, ICAPS 2012, Atibaia, São Paulo, Brazil, June 25-19, 2012*. AAAI.
- Nakhost, H., and Müller, M. 2010. Action elimination and plan neighborhood graph search: Two algorithms for plan improvement. In Brafman et al. (2010), 121–128.
- Nakhost, H.; Müller, M.; Valenzano, R.; and Xie, F. 2011. Arvand: the art of random walks. In García-Olaya, A.; Jiménez, S.; and Linares López, C., eds., *The 2011 International Planning Competition*, 15–16. Universidad Carlos III de Madrid.
- Richter, S., and Helmert, M. 2009. Preferred operators and deferred evaluation in satisficing planning. In Gerevini et al. (2009), 273–280.
- Richter, S., and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research* 39:127–177.
- Richter, S.; Helmert, M.; and Westphal, M. 2008. Landmarks revisited. In Fox, D., and Gomes, C. P., eds., *AAAI*, 975–982. AAAI Press.
- Richter, S.; Thayer, J.; and Ruml, W. 2010. The joy of forgetting: Faster anytime search via restarting. In Brafman et al. (2010), 137–144.
- Thayer, J., and Ruml, W. 2011. Bounded suboptimal search: A direct approach using inadmissible estimates. In Walsh, T., ed., *IJCAI*, 674–679. IJCAI/AAAI.
- Thayer, J.; Benton, J.; and Helmert, M. 2012. Better parameter-free anytime search by minimizing time between solutions. In Borrajo, D.; Felner, A.; Korf, R. E.; Likhachev, M.; López, C. L.; Ruml, W.; and Sturtevant, N. R., eds., *SOCS*, 120–128. AAAI Press.
- Valenzano, R.; Nakhost, H.; Müller, M.; Schaeffer, J.; and Sturtevant, N. 2012. Arvandherd: Parallel planning with a portfolio. In Raedt, L. D.; Bessière, C.; Dubois, D.; Doherty, P.; Frasconi, P.; Heintz, F.; and Lucas, P. J. F., eds., *ECAI*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, 786–791. IOS Press.
- Xie, F.; Nakhost, H.; and Müller, M. 2012. Planning via random walk-driven local search. In McCluskey et al. (2012), 315–322.