Integration of a Refinement Solver and a Local-Search Solver

Alexander Nareyek
Cork Constraint Computation Centre
University College Cork
Cork, Ireland
e-mail: alex@ai-center.com

Stephen F. Smith and Christian M. Ohler
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3891, USA
e-mail: sfs@cs.cmu.edu, ohler@cs.cmu.edu

Abstract

We describe an integration of a refinement solver and a local-search solver for constraint satisfaction, with the goal to use information from the local-search solution process as a basis for directing a backtracking-based refinement search. In this approach, the decision about the next refinement step is based on an interposed phase of constructing/revising a complete (but not necessarily feasible) variable assignment by local search, i.e., as a kind of variable- and value-ordering strategy. The assignment generated by local search is called a "probe."

We investigate the efficiency of this hybrid search approach in the combinatorial domain of job-shop scheduling. First, we evaluate methods for improving probe-based guidance, by basing refinement decisions not only on the final assignment of the probe-construction phase but also on information gathered during the probe-construction process. We show that such techniques can result in a significant performance boost.

Second, we consider the relative strengths of probe-based search control and search control that is directly based on refinement search and biased by more classically motivated variable- and value-ordering heuristics (incorporating domain-specific knowledge). Our results indicate that — while probe-based search performs better than an uninformed search — use of domain-specific knowledge proves to be a much more effective basis for search control than information about constraint interactions that is gained by local-search probes, and leads to substantially better performance.

1 Introduction

A broad range of combinatorial problems is naturally formulated as constraint satisfaction problems (CSPs), and as such, the design of efficient and general search techniques for solving CSPs has attracted much attention in recent years.

Generally, these search methods are based on the idea of *refinement*, i.e., stepwise reduction in the value domains of the decision variables until each variable's domain has exactly one admissible value. Refinement-based search methods are classically formulated within a backtracking search framework. In this framework, promising refinements are recommended by search-control heuristics, which can vary in their problem-specific or more domain-independent nature. Constraint propagation techniques are applied at each refinement step, and provide a basis for detection and early pruning of infeasible states. For example, suppose we have two finite variables A and B with domains of $\{1 \ldots 10\}$ and a constraint $B > A$. In the case of a chosen refinement of $A \in \{5 \ldots 10\}$, the propagation will entail a domain reduction of B to $\{6 \ldots 10\}$. Refinement decisions and propagation are iteratively repeated until a solution is found. If the refinement that was made turns out to be inconsistent later on, which means that no more possible values are left in a variable's domain, backtracking is applied to choose another refinement option.

Search-control heuristics for refinement search typically exploit measures of flexibility/inflexibility in the evolving partial solution — the sizes of domains of various variables in the most basic case. The effectiveness of these heuristics rests on the ability to properly distinguish more-constrained decisions from less-constrained decisions, which can be more or less evident from propagation results at different stages of the search and at different levels of solution difficulty. Since, in the case of domain-independent heuristics, guidance is usually based directly on (local) constraint propagation results, they can suffer from the lack of a global perspective on decision constrainedness. Domain-specific heuristics (e.g., [2, 3]) offer one approach to providing more globally-motivated search guidance.

Search methods that are based on local search, on the other hand, generate a solution by repeatedly revising a concrete value assignment of the variables. For example, the variables A and B could initially be assigned with values of $A = 1$ and $B = 1$. As long as constraints are inconsistent with this assignment, such as $B > A$, the assignment will iteratively be revised. Changes are typically guided by some inconsistency measurement. For example, if the repair methods allow search to increase or decrease a variable, B might be changed to $B = 2$ in order to lower the inconsistency of the constraint $B > A$. Such repairs are iteratively applied until a solution is found or the search is stopped.

Local-search methods tend to build up a global picture of the "hot spots" currently causing inconsistency, but they lack a more systematic way of resolving these constraint interactions.

Given the complementary strengths of refinement and local search, it makes sense to combine them. Use of local search can inject a global perspective to refinement search's control decisions; refinement search can provide a systematic basis for exploring tightly constrained regions of the underlying search space. We focus specifically on use of the inconsistency measure of local search as guidance for determining which

refinements to apply within a refinement search. Some theoretical justification for this model is given in [9], where it is proved that a backtracking-based refinement search guided by information of a complete assignment has a much better average run-time behavior than previously analyzed algorithms that were not.

A number of researchers have proposed approaches along these lines. [5] and [14] applied the GSAT local-search method [11] to generate complete assignments within a backtracking refinement search in the context of solving propositional satisfiability (SAT) problems. Similarly, [10] studied the use of a linear programming solver as a complete-assignment generator to guide a backtracking refinement search, obtaining results that outperformed all alternative algorithms applied to the problem of minimizing perturbations in dynamic scheduling. A variation of this approach was presented in [7], using a local-search approach instead of a linear programming solver.

In this paper, we extend and further evaluate the idea of integrating local search as a basis for more globally informed control of refinement-based CSP search. First, we consider the added benefit of factoring information related to the local search's assignment history into search-control decisions. The work mentioned above has relied strictly on analysis of the final assignment generated. Second, we consider the relative strengths of probe-based heuristics in relation to uninformed refinement search and a search control bias that is provided by more classical CSP variable- and value-ordering heuristics that incorporate knowledge of the problem domain at hand.

We investigate these issues in the application domain of job-shop scheduling. We concentrate specifically on the problem of generating a feasible job-shop schedule. In brief, a problem consists of n jobs, each of which requires execution of a sequence of m tasks. Each task i of job j has a fixed duration $p_{j,i}$ and requires a specific resource $r \in R$. $|R| = m$, and each job visits all resources in some order (determined by its task sequence). All jobs must be completed by a specified global due date h, and a resource can only perform one task at a time. The goal is to find starting times for the tasks that satisfy all time and resource constraints. Contention for resources by competing tasks is what makes the problem difficult (NP hard). Resource contention is also reflective of global constraint interactions, and hence representative of the sort of domain where local search might be expected to boost refinement-search performance.

We start by describing the assumptions and components of our integrated search framework.

2 Solver Integration

To combine local and refinement search, [14] favor a system that uses attributed variables, i.e., which allows the local-search solution's values to be stored directly within the refinement search's solution structures, using the same problem modeling for both search techniques. While this certainly eases problem modeling, it has the major drawback that models cannot be created in a way to enable structure exploitation for both search approaches. Search performance is highly dependent on appropriate models, and we thus instead use an integration of two solvers, both based on constraint programming but specifically designed for local and refinement search respectively. Both of our solvers have different internal representations of the problem, and were not

initially designed to work in an integrated way. The solver and their integration are described in the following.

2.1 The Refinement Solver

The base refinement solver is built using the Comirem planning and scheduling framework [13], which promotes a general view of problem solving as an incremental constraint posting process. At its core is a Simple Temporal Problem (STP) constraint-network solver [6]. All input tasks, durations, sequencing constraints, and deadline constraints are translated into a graph of time points and distance constraints. Constraint propagation is performed to establish initial time bounds on the start and end of each task. The solver is then invoked to update these start-time and end-time bounds and to detect constraint conflicts (infeasible states) as constraints corresponding to each new scheduling decision are added to the network.

The process of scheduling in this model is not concerned with assigning task start times, but instead aims to feasibly sequence the set of tasks that are competing for each resource.[1] A basic refinement step in the search involves two decisions: (1) selecting an as yet unsequenced task and (2) selecting where to insert this task into the partial sequence that has been established thus far on the required resource's timeline. We refer generally to a feasible position in a resource timeline as a "slot." In the underlying constraint network, the selection of a slot may result in the posting of one or two precedence constraints between competing tasks, depending on whether the task is inserted at the beginning, at the end, or somewhere in the middle of the resource's timeline. Thus, a final solution will typically designate a set of possible start times for any given task (all feasible) rather than committing to a single point.

At each refinement step, the base refinement procedure recomputes the feasible options (slots) for each uninserted task and applies search control heuristics to determine the next decision. In the event that an infeasible state is detected, the search backtracks chronologically and considers alternative slots for previously inserted tasks. Thus, the basic scheduling procedure is complete if given enough time to execute, and search control heuristics are used to improve average case performance.

We make use of two different heuristics for task selection, aimed at selecting the most temporally constrained task for insertion next:

- *Least Options First* (`LOF`):
  $Min_{i \in UninsertedTasks}(|Slots_i|)$. In case of ties, then
  $Min_{i \in UninsertedTasks}(Min(lft_i - est_i - p_i, \sum_{s \in Slots_i}(lft_s - est_s - p_i))$,
  where lft_x is latest-finish-time of x, est_x is the earliest start time, and p_i is the processing time of i.

- *Minimum Slack First* (`MinSlack`):
  $Min_{i \in UninsertedTasks}(Min(lft_i - est_i - p_i, \sum_{s \in Slots_i}(lft_s - est_s - p_i))$,
  where lft_x, est_x, and p_i are defined as above.

[1] Hence, constraint propagation through the STP network in this context implements enforcement of disjunctive resource constraints as described in [1].

Intuitively, the number of feasible options (slots) remaining provides one basic estimate of temporal constrainedness. However, given the nature of the search space, the insertion of a task into a given resource's schedule (sequence) can sometimes increase the number of options for other pending tasks that require this same resource, and hence LOF decisions may not always be directly correlated to temporal constrainedness. MinSlack provides a simpler but potentially more direct measure of a task's current degrees of freedom [12].

For option selection, we adopt the following heuristic:

- *Maximum Slack First* (`MaxSlack`):
  $Max_{s \in Slots_i}(lft_s - est_s - p_i)$.

Intuitively, the option that retains the most temporal flexibility is selected by `MaxSlack`.

To provide an alternative basis for search control, we also define non-deterministic versions of these task selection and option selection heuristics. Following the principle of calibrating the level of non-determinism to the discriminatory power of the search heuristic in a given decision context [4], we randomize these heuristics in a value-biased manner. Specifically, we modify the heuristic values assigned to each choice by a bias function of the form $value^b$, and use the resulting numbers to define choice probability. We refer to the randomized counterparts of the above heuristics as `LOF(b)`, `MinSlack(b)`, and `MaxSlack(b)` respectively.

2.2 The Local-Search Solver

The local-search solver is a modification of the DragonBreath engine[2] (see [8] for details). The solver is a general constraint-programming system and not specialized to scheduling problems. However, it is easy to express the job-shop scheduling problem by way of the given constraint types.

Two high-level constraint types are used to model the problem: a non-overlap constraint for a machine's tasks and a linear-inequality constraint to describe a temporal relation between tasks. In the search approach of the DragonBreath engine, every constraint calculates an inconsistency value using a constraint-specific measure indicating how far off the involved variables' assignment is from a consistent solution. For example, the non-overlap constraint returns costs related to the time spans during which involved tasks are overlapping.

The sum of the constraints' inconsistencies represents the total inconsistency of the solver's current assignment. In each improvement iteration then, an inconsistent constraint is selected, and the constraint selects one of its constraint-specific heuristics to change its variables' values in order to reduce the constraint's inconsistency. For example, a non-overlap constraint might select a heuristic to shift a task to a time at which it causes less overlaps.

The local search can always escape local minima in the search space because the applicability of the constraints' heuristics is not restricted to greedy improvements of the overall inconsistency. Randomization and learning techniques are applied to promote exploration.

[2]The engine is freely available at:
`http://www.ai-center.com/projects/dragonbreath/`

In contrast to the refinement solver, the local-search solver has concrete values assigned for all variables — the tasks' starting times in this case — at any time. The local-search solver therefore also has a complete picture about the currently involved inconsistencies, which makes it easier to identify potential "trouble spots", i.e., regions where it is not easy to establish consistency.

2.3 The Solvers' Interaction

The interaction of the solvers is shown in Figure 1. Refinement search is the master process, using the local-search solver for heuristic guidance. Both solvers have different internal representations of the job-shop scheduling problem and communicate in reference to decision variables, which are the tasks' starting times. For every refinement decision of the refinement solver, the local-search solver is notified to internally add a corresponding constraint <4> so that both solvers keep working on the same problem. Thus, the search of the local-search solver is not restarted every time a refinement decision is to be made but the solver keeps working on the problem, only adding the constraints that are implied by the refinement decisions. In case of backtracking, these constraints are removed again <5>.

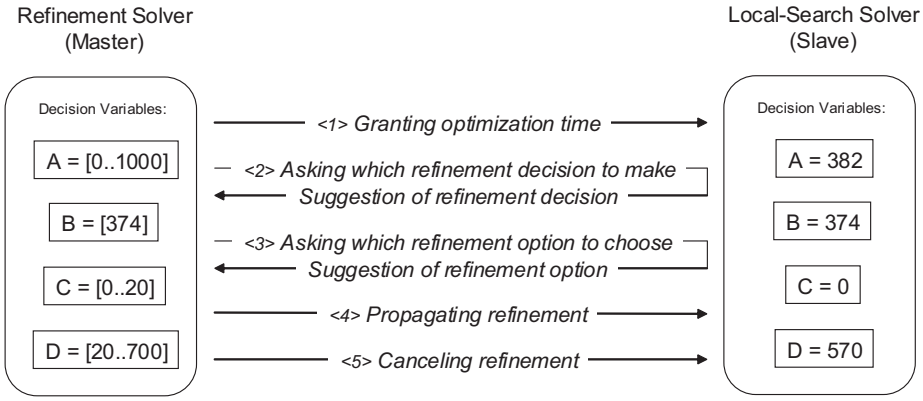

Figure 1: Solver integration.

The local-search solver can suggest, at each refinement step, which refinement decision to make and which refinement option to choose. For the job-shop scheduling problem, the local-search solver suggests for which task an ordering decision should be made next <2> and which ordering constraints (before or after) relative to other tasks on the same machine should be chosen <3>. Since we are also interested in exploring the interplay of local search and the base refinement heuristics, we also admit configurations that restrict the use of local-search guidance to either task selection <2> or ordering decisions <3>.

Refinement recommendations of the local-search solver are generally made in a way to focus on critical regions, i.e., recommending variables for refinement that are related to the highest amount of inconsistency, and to steer the local-search solver's

assignment away from inconsistencies, i.e., recommending values that minimize these inconsistencies (see following sections for details).

At each refinement step, the local-search solver is given some time for optimization <1> to adapt to the current situation (if constraints were added or removed) before a recommendation is queried. In our current implementation, this time includes a time for general optimization (500 iterations for the test runs given below), and a time in which the local-search solver concentrates on satisfying only the refinement decisions already taken (maximally 200 iterations for the test runs given below). There is also some optimization time granted to the local-search solver before a refinement (1000 iterations for the test runs given below). If at any point during the overall search, the local-search solver finds a feasible solution to the problem at hand, then this solution is returned and the search terminates in success.

2.4 The Translation Layer

The fact that the solvers use different search paradigms makes it necessary to add functionality to translate search-state information. In the model of the Comirem framework, search decisions are related to refining variables' domains, while the DragonBreath engine does not usually need to maintain information regarding such decisions. As described above, we have decided to use a communication with respect to the refinement ontology between the solvers, which leaves it to the DragonBreath engine to maintain additional structures in order to generate meaningful output in terms of refinement search.

The additional translation structures (which are explained in the following sections) are very expensive in terms of required memory and CPU because every potential refinement decision/option that can be suggested needs to be tracked. This results in a large overhead to be maintained. Changing the value of only a single task's starting variable requires re-analyzing all suggestions that are somehow affected by this variable. Especially if the full history, i.e., the state after each single improvement iteration of local search, is to be analyzed, this becomes quite costly. It can surely pay off, but when this kind of architecture is considered, respective costs have to be kept in mind.

3 Test Problems

Our test problems are taken from the OR-library[3]. They include 14 problems with 10 jobs and 10 resources, one with 15 jobs and 15 resources, and five with 20 jobs and 10 resources. Note that the commonly used makespan-optimization benchmarks for these problems are not applicable here because we tackle the *satisfaction* problem – searching for a solution within a fixed horizon.

For each test run, there is a cut-off at 300 nodes (600 for the larger instances), one "node" corresponding to one task and slot selection. Most variants discussed below involve randomization components. 50 test runs are thus run per configuration.

Three difficulties are defined per problem to obtain problems of different hardness. Problem difficulty is a function of the tightness of the global due date. To determine

[3] `http://graph.ms.ic.ac.uk/info.html`

problem difficulties, the temporal distance between an infinite capacity solution (lower bound) and a due date equal to $\sum_{i \in Tasks} p_i$ (upper bound) was divided into 100 increments. The "hard" difficulty represents the loosest global due date (moving from upper to lower) for which the refinement solver alone (i.e., using only the base heuristics `LOF` and `MaxSlack`) cannot produce a solution with a cut-off of 1,000 nodes. The "medium" difficulty is obtained by setting the global due date one step toward the upper bound and "easy" is obtained by moving three steps further in that direction.

Many performance parameters depend on the actual implementation of the solvers and their integration. We will thus not focus on specific numbers in the results but on general trends instead.

4 Involving Local Search's Search History

In previous approaches, a recommendation of the local-search solver was based on the variable assignment and resulting inconsistency after a specific amount of time was given to the local-search solver for improvement[4]. We will refer to this as strategy `CI` (like "Current Inconsistency") in the following. For this strategy, after the local-search phase, the inconsistencies of every constraint that a task is involved in are added up, and the task with the highest inconsistency sum is recommended for a refinement decision. For the recommendation of the ordering, all current orderings of a task with respect to the selected task are considered, and the reverse of the ordering involving the highest inconsistency is recommended.

However, this approach only exploits the information that is available after the improvement phase of local search has been completed. If local search is considered to be an intelligent sampler of a search space, this means that only one sampling is returned by this approach. Moreover, the assignments at this time may just represent a short-term anomaly caused by a bad last local-search move. We thus vary the recommendation strategy in the following by additionally involving information gathered during the complete local-search phase, incorporating information of multiple assignment samples thereby.

4.1 Iterations of Inconsistency

The strategy analyzed in this section makes use of demons that record specific features about the constraints and their inconsistency during the local-search improvements.

For every pair of tasks t_a and t_b on a resource, there is a demon d_{ab} recording the number of iterations in which t_a was preceding t_b and constraints involving t_a or t_b were inconsistent. There is a complement demon d_{ba} for situations in which t_b was preceding t_a. When a recommendation is to be made, demons propagate their iteration counts to the tasks involved, which sum these counts. The task with the most "iterations of inconsistency" is recommended for a refinement decision (strategy `II`). For the ordering recommendation, the reverse ordering of the demon with the highest

[4]Note that there are also some approaches in which local search is run until a local minimum is reached, or – using local search only for sub-problems – in which local search is run until a fully satisfying solution is found. This, however, provides little temporal control for the search process.

iteration count is recommended. The demons' counts are reset to zero every time a refinement decision is made or canceled/backtracked by the refinement solver.

4.2 Average Inconsistency

For the following strategy `AI`, we additionally involve the constraints' quantitative inconsistency values. In contrast to strategy `II`, not only the number of inconsistent iterations is recorded by the demons but also the quantitative inconsistency value that is involved. The average inconsistency during inconsistent iterations is then propagated to the tasks involved and used as recommendation criterion.

4.3 Bound Propagation

In addition to notifying local search about the ordering refinements adopted, the refinement-search solver can also inform local search about changes to the time bounds of individual start-time variables, which are updated during constraint propagation when a refinement decision is made. Implicitly, local search already has this information by way of the orderings, but additional explicit constraints may help local search to fulfill these bound constraints. Bounds represent hard constraints in the DragonBreath engine, i.e., a variable's value will immediately be shifted back into the bound if a heuristic tries to assign a value outside the bound. Thus, the propagation of such bounds may have a much larger impact than the implicit information. The extension of strategy `AI` by propagating these bounds to local search will be called `AI-B` in the following.

4.4 Results

Because of the large number of test runs and results, we only provide a very compressed representation. For every problem instance and solving strategy, 50 test runs are computed. Strategies are then compared in a pairwise way for each problem, rating a strategy as better if it was able to find a solution (within the cut-off limit) in 10% more cases than the other strategy.

Figure 2 shows the results for all strategy pairs. It can be seen that the strategy `II` outperforms `CI`, and `AI` performs best of all. This confirms our thoughts that the exploitation of inconsistency information that is gathered during the local-search phase can substantially improve local search's recommendations.

The result of propagating bounds — strategy `AI-B` — does nearly always perform better than strategy `AI`. However, keep in mind that we count refinement steps/nodes and not total computation time. Propagating bounds to the local-search solver involves computation costs, and it depends on the applied solvers and communication speed if this additional overhead pays off. In some rare cases, propagating bounds slightly worsens the results. An explanation for this might be that in certain situations, the application of hard constraints for bounds restricts local search's ability to move to other regions of the search space, sampling less of the relevant assignments thereby.

Having a closer look at problem instances that did not behave according to our expectations, it quickly becomes apparent that `AI` and `AI-B` actually perform very poorly (i.e., similar to `CI`) for the five larger problem instances. Our suspicion is that

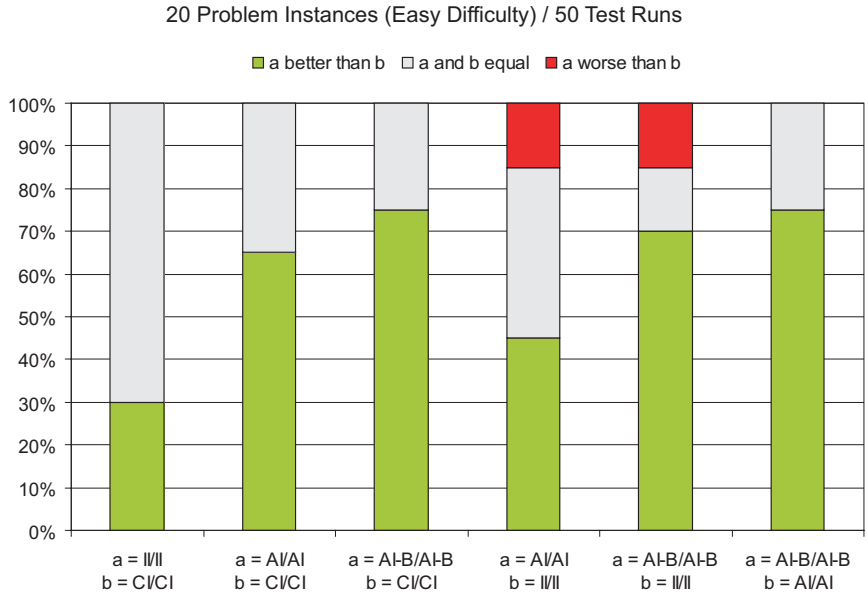

Figure 2: Comparing variants of local-search recommendations.

this is a cost-balancing problem at the local-search side, because, with the growth of the problem, the number of inequality constraints is growing much faster than the number of non-overlap constraints. We need to investigate this effect further, and can only state its existence for now.

In conclusion, we have shown that the advice from local search can be significantly enhanced by exploiting the information that is available during the local-search phase. It is a "safe" extension to involve the number of iterations of inconsistency for this, while one should be careful with quantitative cost information. Moreover, the performance can be much improved by making refinement search's bound-propagation information available to local search.

5 Comparison to Uninformed Recommendations

We have gained an improved local-search advice, but have not yet shown that this advice actually guides refinement search in a better way than a simple uninformed/randomized choice for task selection and ordering at each refinement step (`RND`).

The results for a number of configurations are shown in Figure 3, and imply that `AI-B/MaxSlack` is better than `RND/MaxSlack` is better than `AI-B/AI-B` is better than `AI-B/RND`. This shows that the `AI-B` task ordering is better than a random choice, while using the deterministic heuristic for the task ordering seems to be even better (this topic is detailed in the following section). For the task selection, `AI-B` turns out to be better than a random choice as well, showing that `AI-B` is a well-informed

heuristic in general.

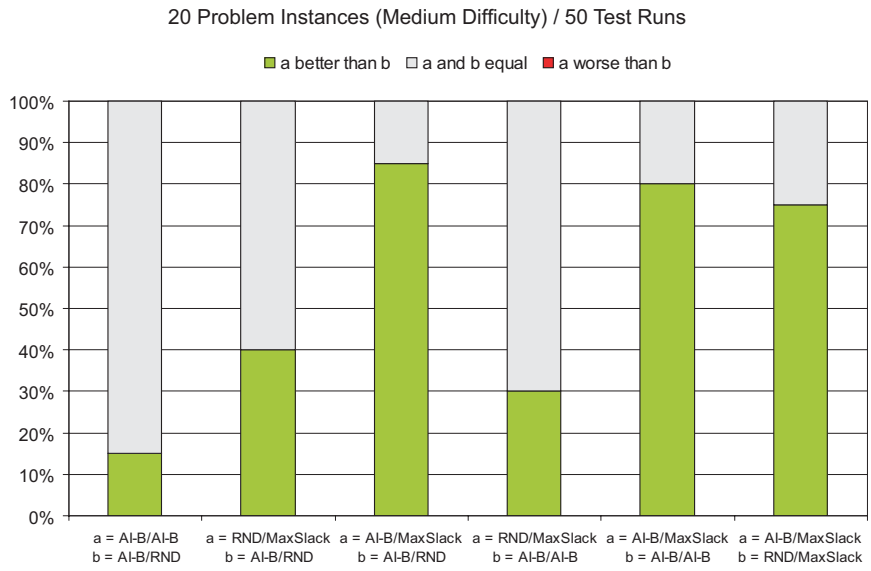

Figure 3: Comparison to uninformed refinements.

If a local-search solver is used, however, the costs of running the local-search optimization and the overhead for solver communication must be considered. In our experimental setting, it turned out to be a substantial overhead that hardly justifies the improvement gained in search guidance. A concrete quantification of this overhead, however, would be misleading because this overhead may generally vary depending of the applied solvers and integration.

6 Comparison to Domain-Specific Heuristics

We have shown that the inconsistency information of local search's sampling is a well-informed heuristic. However, for many problems, it may be easy to come up with domain-specific refinement heuristics. In the previous section, it already became clear that using the `MaxSlack` refinement heuristic for option selection instead of local search's recommendation drastically improves performance.

Figure 4 shows the impact of the refinement heuristics described in Section 2.1 for task selection. Deterministic `LOF/MaxSlack` is, because of our definition of hardness, unable to solve any of the hard instances, but configurations of `LOF(b)/MaxSlack` for various bias values outperform `AI-B/MaxSlack` in most cases. A second configuration `MinSlack(b)/MaxSlack` that utilizes the simpler `MinSlack` heuristic for task selection was found to produce even better results. Moreover, even the deterministic variant `MinSlack/MaxSlack` is able to solve many problems. These

results are all the more significant considering that use of the refinement heuristics is *much* cheaper than calling the local-search solver.

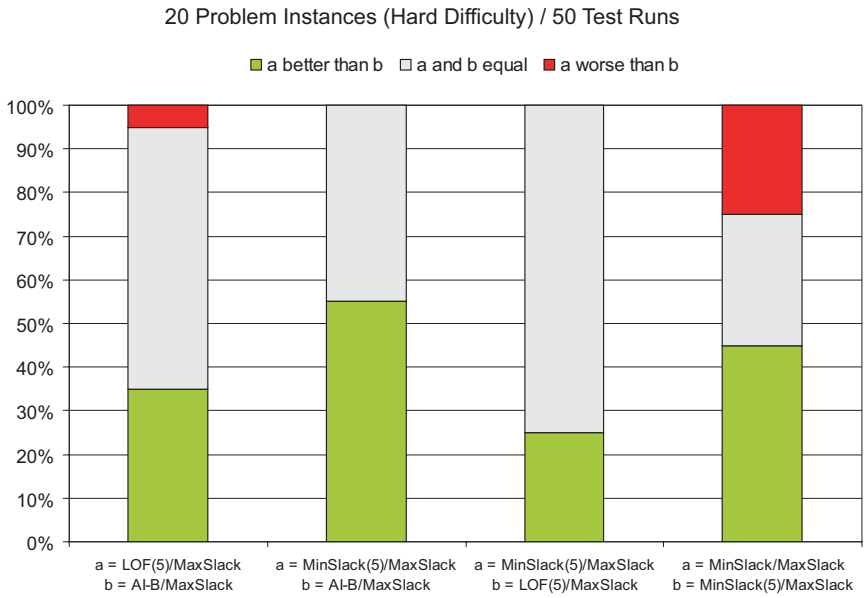

Figure 4: Comparison to domain-specific heuristics.

7 Conclusion

We have studied the integration of local-search advice into a refinement search. The advice was provided by suggesting refinements that guide refinement search away from the inconsistency of local search's assignment samplings.

In the approaches of previous work, local search was allowed to run for a while, and only the final assignment was used for recommendations. The local-search assignment at this time may be of high quality in terms of local search's cost metric, but we have shown that integrating the information of intermediate search steps results in better advice. Intuitively, the advice includes more possible assignments that were sampled by local search this way. Our experiments show that feeding domain-propagation results back into local search can boost performance even further. One should careful, however, with a guidance that is based on the quantitative cost dimension of local search. While this can further improve performance, an unwary cost balance may similarly degrade performance.

We have shown that the inconsistency of local search's assignment can provide productive heuristic guidance. However, the computational costs for the integration overhead and local-search optimization make this a rather dubious result. Moreover, even simple domain-specific heuristics seem to result in equal or better performance

than the "mushy" information of inconsistency of local-search samplings. This seems to be plausible, but we were surprised how easy it is to outperform the probe-based heuristic.

An interesting question for the future is if we can develop heuristics that can classify their recommendation value, and possibly switch to a probe-based heuristic in contexts where they can provide only little guidance. Preliminary experiments that decrease the probability of using the probe-based heuristic with increasing refinement depth, however, did not look highly promising. The rationale for this approach was that refinement search potentially gets more informed with increasing refinement depth because the variables' domains get smaller, enabling predictions that are more precise.

For other domains than job-shop scheduling, where less domain knowledge can be exploited by refinement heuristics, local-search-based guidance may still be a viable option. In these cases, one should consider the extended integration techniques presented in this paper, i.e., also incorporating information gathered during the local search phase and feeding propagation results back into local search.

8 Acknowledgments

This work was supported in part by DARPA Contract #F30602-00-2-0503, the CMU Robotics Institute, and the Science Foundation Ireland under Grant 00/PI.1/C075.

References

[1] Baptiste, P., LePape, C. and Nuijten, W. *Constraint-Based Scheduling: Applying Constraint Programming to Scheduling Problems*. Kluwer Academic Publishers, Boston, 2001.

[2] Beck, J. C. Texture Measurements as a Basis for Heuristic Commitment Techniques in Constraint-Directed Scheduling. Ph.D. Thesis, Department of Computer Science, University of Toronto, 1999.

[3] Cheng, C., and Smith, S.F. Applying Constraint Satisfaction Techniques to Job shop Scheduling. *Annals of Operations Research* 70: 327–357, 1997.

[4] Cicirello, V., and Smith, S. F. Amplification of search performance through randomization of heuristics. In Proceedings of the Eighth International Conference on Principles and Practice of Constraint Programming (CP 2002), 2002.

[5] Crawford, J. M. Solving Satisfiability Problems Using a Combination of Systematic and Local Search. Second DIMACS Challenge: Cliques, Coloring, and Satisfiability, Rutgers University, NJ, 1993.

[6] Dechter, R.; Meiri, I.; and Pearl, J. Temporal Constraint Networks. *Artificial Intelligence* 49(1): 61–95, 1991.

[7] Kamarainen, O., and El Sakkout, H. Local Probing Applied to Scheduling. In Proceedings of the Eighth International Conference on Principles and Practice of Constraint Programming (CP 2002), 155–171, 2002.

[8] Nareyek, A. Using Global Constraints for Local Search. In Freuder, E. C., and Wallace, R. J. (eds.), *Constraint Programming and Large Scale Discrete Optimization*, American Mathematical Society Publications, DIMACS Volume 57, 9–28, 2001.

[9] Purdom Jr., P. W., and Haven, G. N. Probe Order Backtracking. *SIAM Journal on Computing* 26(2): 456–483, 1997.

[10] El Sakkout, H., and Wallace, M. Probe Backtrack Search for Minimal Pertubation in Dynamic Scheduling. *Constraints* 5(4): 359–388, 2000.

[11] Selman, B.; Levesque, H.; and Mitchell, D. A New Method for Solving Hard Satisfiability Problems. In Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92), 440–446, 1992.

[12] Smith, S.F.; and Cheng, C. Slack-Based Heuristics for Constraint Satisfaction Scheduling. *Proceedings 11th National Conference on Artificial Intelligence* Washington D.C., July 1993.

[13] Smith, S.F.; Hildum, D.; and Crimm, D. Interactive Resource Management in the Comirem Planner. *Proceedings AAAI Workshop on Mixed-Initiative Intelligent Systems*, Acapulco Mexico, August 2003.

[14] Schimpf, J., and Wallace, M. Finding the Right Hybrid Algorithm – A Combinatorial Meta-Problem. *Annals of Mathematics and Artificial Intelligence* 34(4): 259–269, 2002.