Proceedings of the First Canadian Workshop on Distributed Artificial Intelligence,
Banff, AB, May 15, 1994.

Supply Chain Coordination via Mediated Constraint Relaxation

J. Christopher Beck[1] & Mark S. Fox
Department of Computer Science
University of Toronto
Toronto, Ontario, M5S 1A4
chris@cs.utoronto.ca msf@ie.utoronto.ca

Abstract

Coordination of the participants in the supply chain of a manufacturing enterprise is a key to agile reaction to unexpected events. As a starting point, we take a mediated approach to coordination: a single agent is responsible for recovery of the supply chain from a disruptive event. This mediator gathers commitment information from other agents and forms a constraint graph. If the event is truly disruptive, this graph will reflect an infeasibility: a subset of agents can no longer meet commitments. Repair of the graph is done via constraint relaxation controlled by the mediating agent. We present a schema for constraint relaxation algorithms and experimental results on Partial Constraint Satisfaction Problems (PCSPs). We sketch the coordination protocol that is being developed.

1.0 Introduction

Dynamic events in a multiagent environment can have significant impact on the ability of agents to meet commitments made to other agents. If the network of commitments is viewed as a constraint graph, an environmental event that prevents the meeting of a commitment creates an infeasible constraint graph. We apply constraint relaxation directed by a mediating agent in an attempt to optimally reconfigure the commitment graph. This model is applicable to a wide range of multiagent domains where response is needed to unexpected events. An example of such a domain is that of supply chain management.

2.0 Supply Chain Management

The ability to quickly respond to environmental changes has been recognized as a key element in the success and survival of corporations in today's market [Nagel 91] . This agility includes an ongoing monitoring of events both inside and outside the corporation, quick recognition of the impact of exogenous events, and rapid re-planning and reconfiguration to allow the enterprise to take advantage of opportunities and minimize incurred costs.

In a manufacturing enterprise, the entire supply chain is subject to unexpected events for which reactions are required. The supply chain flows from the customer order taken by the sales division through planning, production, distribution, field service, and reclamation. Exogenous events are many and varied: change in the customer order, unavailability of a particular resource, price change in a resource, late delivery of a resource, breakdown of a machine, an urgent order from a good customer, and so on. Handling these events requires close coordination and cooperation among sales, marketing, accounting, material planning, production planning, production control, and transportation. The following example illustrates the scope of the problem [Fox 92] .

1. Supported by Natural Sciences and Engineering Research Council Centennial Fellowship.

The Canfurn, Inc. furniture company produces a variety of furniture with options on wood type and upholstery. Leo's, the largest and best-paying customer of Canfurn, places a large order for delivery in six months and Canfurn is able to schedule delivery as requested. Two months before the original delivery date, Leo's requests a significant change in the order but still wants to maintain the delivery date. Canfurn's sales department immediately contacts the manufacturing division. Manufacturing has a number of options:

- Can the new order be manufactured? Are extra shifts needed to meet capacity requirements? What does personnel think of extra shifts? Are the materials for production in stock? If not, can a supplier be found that can make delivery?
- Can another order be delayed (and possibly delivered late) in order to meet Leo's order? What does sales think of this?
- Can the job be subcontracted to another manufacturer? What does marketing and strategic planning think of this? What does accounting say about reducing the margins? Can we afford to take a loss on the order?

Clearly, the manufacturing division cannot make the decisions on its own. It must canvas a number of other divisions within the company and some external bodies (suppliers and subcontractors) in order to choose an alternative that is optimal.

The supply chain extends over the breadth of the enterprise and, as in actual corporations, the inter-agent coordination is hierarchical. For example, with multiple production centers, coordination among agents within one center is at a level of abstraction below the coordination among the centers. The latter abstraction is enterprise-wide logistics. It has a global view of the enterprise and is be concerned with sales, customer delivery, and all aspects of inter-production center coordination. We will focus on this level.

Each production center is viewed as a single resource with the ability to perform multiple activities resulting in the production of a quantity of some resource. The activities that each factory can perform and the capacity of each factory is known.[2] Scheduling at this level involves assigning factories to supply specific quantities of resources at particular times.[3] Figure 1 shows a schematic of the logistics level assignments when an order is received via the Order Acquisition agent.

---

2. These capacities represent aggregate information based on previous performance. Environmental events (e.g. machine breakdown) can dynamically impact these capacities.
3. For now we ignore transportation between factories and the delivery to the customer.

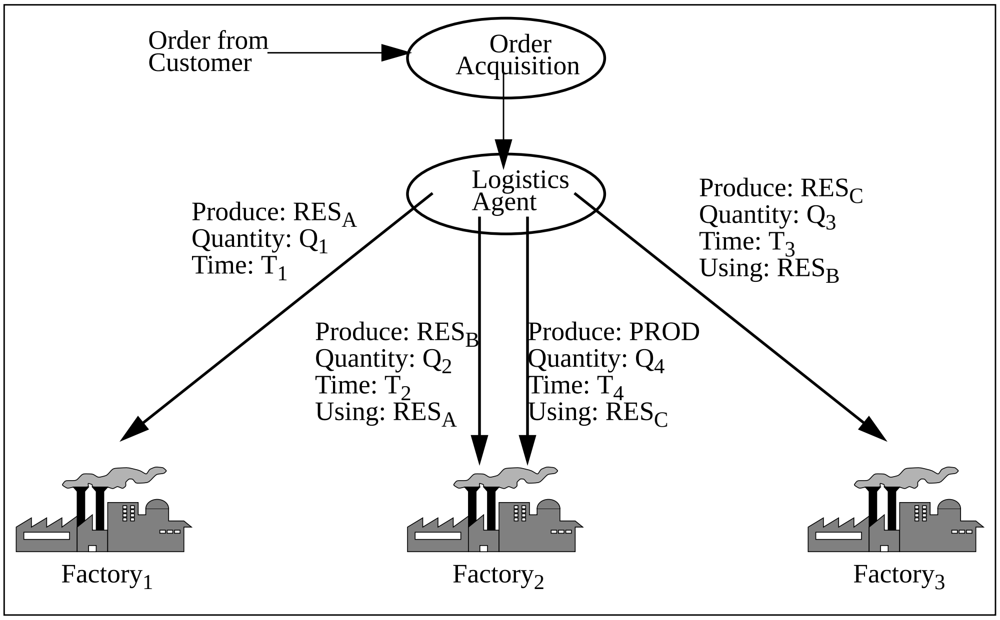

FIGURE 1. Logistics Level Scheduling of a Customer Order

When a customer order is scheduled each factory commits to the execution of a number of activities. Each activity has a number of characteristics such as quantity of the resource to be produced and the time of completion. On the basis of these commitments activities at other factories are scheduled, producing a constraint graph of inter-dependent activities such as shown in Figure 2.

This is a simple graph, dealing with a single order. Given multiple orders and activities at each factory, a full constraint graph will certainly grow to a non-trivial size. A search for a near-optimal reconfiguration has to handle the combinatorial explosion of interdependent alternatives in resource choice, transportation method, and execution times.

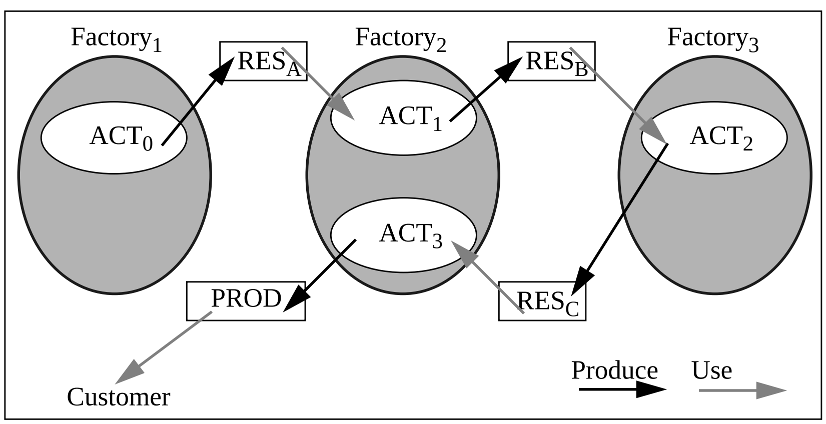

FIGURE 2. A Simple Activity/Resource Constraint Graph

The Enterprise Integration Laboratory at the University of Toronto is pursuing the development of an Integrated Supply Chain Management System (ISCM) addressing these and other problems. The project is based on a distributed simulation of an enterprise, with departments encapsulated as software agents. Given this distribution, the inter-departmental coordination in real corporations is manifest in inter-agent coordination in the simulation. Rationale for the distributed model is manifold such as the exploitation of concurrent processing and the reduction in the complexity of the knowledge base of each agent. Part of the project is the examination of enterprise knowledge and the forming of a reasonable partitioning of it into functional agents. We will use the following agents in our logistics-level coordination example below.

- **Logistics**: responsible for the logistics-level scheduling.
- **Order Acquisition** (OA): responsible for all order-related contact with customers including order entry, cancellation, and modification.
- **Resource Manager** (RM): responsible for managing both the consumable resources (e.g. timely ordering of raw materials) and the usable resources (e.g. scheduling of regular maintenance on the machines).
- **Factory Agents**: each factory has an agent that is the factory-level scheduler.

3.0 A Mediated Approach

Much coordination work looks at negotiation among agents as the coordination mechanism [Lesser 81] [Durfee 87a] [Durfee 87b] [Durfee 91] . In contrast we adopt a mediated approach.

[Sathi 89] shows that in some cases a mediated solution can be significantly better than a negotiated solution. The problem investigated is resource reallocation where each agent has some resources and needs other resources. Agents sell the resources they have for those they need. Experimental results show poor performance for a negotiation-based algorithm, while a mediated algorithm using texture measurements[4] on the aggregate constraint graph out-performed the human expert.

We take this work to indicate that coordination via a completely distributed algorithm is not necessarily the best choice. The problems studied demonstrated “keystone” components: critical subproblems for which few solutions exist, but such that once a solution is found the rest of the problem is significantly easier. For problems such as these, a mediated or partially mediated approach may result in superior solutions.

Based on this work, we have a number of reasons for a mediated approach:

- The supply chain is a highly-structured domain that will only require explicit coordination techniques when unexpected events occur.
- A mediator minimizes the coordination knowledge overhead at each agent.
- The constraint graph upon which relaxation is performed is partially represented in the existing schedules. Logistics is a natural choice as a mediator at the logistics level as is the Factory Agent within each factory.
- Scheduling difficulties can often be traced to scarce or “bottleneck” resources [Smith 89] [Sadeh 91] [Fox 90] . Coordination in the supply chain is needed when the schedule is found to be infeasible due to environmental events, therefore we expect difficulties will arise from similar resource properties. The work of [Sathi 89] indicates that problems with such keystone elements can be solved much better by mediated protocols.

4. Texture measurements [Fox 89] [Sycara 91] assess structural properties of the constraint graph representation of the problem. Based on these measurements, heuristic search decisions are made.

4.0 Constraint Relaxation

We view a multiagent plan as a commitment/constraint graph amongst the participating agents. Each agent is assigned one or more tasks toward the global goal. The tasks that are assigned are subject to constraints: a maximum amount of a shared resource that can be used, latest acceptable completion time for the task, precedence constraints with tasks to be completed by other agents, and the availability of a resource as a precondition on the task, among others. When an agent accepts a task, it is committing to the satisfaction of all constraints on that task. Commitments create a series of inter-dependencies among the agents. Finding an original, feasible constraint graph is the problem of multiagent planning. We assume that such a plan (and the corresponding constraint graph) is in place and focus on recovering from stochastic events that make the graph infeasible.

With the occurrence of an event, an agent may no longer be able meet all the constraints on a task. For example, in a manufacturing domain, a machine breakdown can prevent a task at a particular factory from being completed on time. If this task produces a resource needed for a subsequent task at a different factory, the former factory can no longer meet its commitments. It is necessary to assess the alternatives and adopt the one that has the least negative global impact.

A constraint relaxation algorithm takes an overconstrained constraint graph and attempts to find the minimum cost modification that can be made to a subset of the constraints to produce a feasible graph. The cost associated with modification of the constraints are the key optimization criteria. It is important to note that the cost of a relaxation depends on *how* the constraint is relaxed. If the constraint on the end time of an activity is relaxed by three hours, the cost that will be incurred (due to late or rush delivery of the final product to the customer) will be significantly less than if the relaxation allows three more days.

4.1 A Schema for Constraint Relaxation

4.1.1 Constraint Model

We extend the common variable/constraint model used in CSPs by adding to the constraint representation. Each constraint:

- is defined over a subset of variables, $\{x_i, \ldots, x_j\}$.
- contains a predicate, **Satisfied**($x_i, \ldots, x_j$), which returns TRUE to indicate that the constraint is satisfied by the current variable instantiation and returns FALSE otherwise. The predicate is defined over the Cartesian product of the domains of the relevant variables.
- contains a function, **GenerateRelaxation**, which returns a set of constraints that are relaxations of the constraint.
- contains a function, **RelaxationCost**(c_k), indicating the local cost incurred if the constraint is relaxed to match c_k.

The two functions are keys to a relaxation algorithm. In the former, we put no limitations on the form of the function; it will typically be a center for heuristic decision making. We want to limit the number of relaxations of each constraint that we investigate, therefore heuristic decisions, likely based on texture measurements of the graph, will be used to prune the possibilities. A constraint can be non-relaxable, in which case the **GenerateRelaxation** function returns an empty set. Similarly, we do not put limitations on the form of the **RelaxationCost** function. The cost will depend on the particular sources of cost in the problem model.

4.1.2 Constraint Propagation

The central mechanism for investigation of relaxation costs is the propagation of information through the constraint graph. Propagation is common in consistency algorithms, however the constraint graph structure allows information other than simply the values to be “transmitted” to other variables.[5] A brief example of the propagation of values, costs, and relaxations on the graph in Figure 3 is presented.

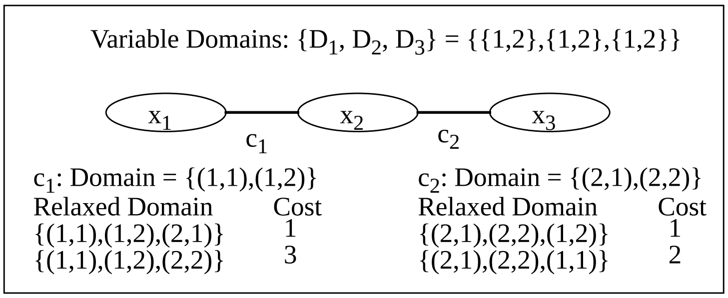

FIGURE 3. A Simple Constraint Graph

Suppose we want to find the minimum cost if $x_1 = 2$. In propagating the value $x_1 = 2$ to x_2, we can not satisfy c_1 as it is. The **GenerateRelaxation** function, in this case, simply returns a singleton set containing the local minimum cost relaxation. We choose that relaxation and propagate the consistent value 1 to x_2. We assign $x_2 = 1$, perform the same greedy procedure on c_2, and propagate 2 to x_3. Once at x_3, the cost of the relaxations is propagated backward. A cost of 1 is propagated from c_2 to x_2. This is summed with the cost at c_1 and propagated to x_1. The cost of the graph, with these relaxations is 2. This is the minimum cost with $x_1 = 2$, but finding the minimum is not guaranteed with the minimum local cost approach. If the cost is acceptable, the relaxation is propagated in the same way as the values. With relaxation propagation we actually replace each constraint with the best relaxation that was tried in the cost/value phase.

More generally, we propose an constraint relaxation schema as follows:

```
select a variable, a set of candidate values, and a set of outgoing constraints
for each value{
   for each outgoing constraint, c{
      find a set of candidate constraints[6], CCc
      propagate value along each constraint in the CCc set
      record the element of CCc that returns the minimum cost
   }
   store the sum of the minimum costs from each outgoing constraint
}
if one of the summed costs is acceptable{
   select corresponding local value
   instantiate the value
   propagate relaxations along the outgoing constraints
}
```

5. This use of propagation builds on work done in propagation of preferences [Sadeh 89] .
6. The set of candidate constraints can contain the constraint itself and any relaxations of it.

Clearly, if there is a possibility of cycles in the graph, this must be dealt with in order to terminate the algorithm.[7] The strength of this schema is the isolation of heuristic decision points in the value selection, the selection of outgoing constraints, and the selection of candidate constraints. If all values and all constraints are selected, we have an exhaustive (and exponential) algorithm. A key to practical algorithms is the exploitation of the problem structure based upon which limits can be placed on the propagation. For example, estimates of costs can be made without propagation or the algorithm can be focussed on a certain type of constraint or a certain subgraph.

We have applied algorithms within this schema to Partial Constraint Satisfaction Problems [Freuder 92] and to a set of scheduling optimization problems [Sadeh 91] . The algorithms applied to the PCSP problems are exponential in complexity and require significant caching of information in order to increase performance and cope with graph cycles. Despite this the algorithms performed well as compared with PEFC3, the best PCSP algorithm investigated by [Freuder 92] . In most cases, the algorithms ran faster than PEFC3 and found solutions with a cost that was equal to or just higher than the cost of solutions found by PEFC3.

Algorithms		Problem Sets					
		1	2	3	4	5	6
PEFC3	Avg. Checks	4664	1460	18077	9700	17221190	592303
	Avg. Cost	3.1	8.9	4.1	12.7	7.6	23.4
MMVLH1	Avg. Checks	1871	1055	5329	2686	-	-
	Avg. Cost	3.3	11.6	4.2	14.2	-	-
MMVLH2	Avg. Checks	2647	1647	8315	5200	-	-
	Avg. Cost	3.1	10.4	4.1	13.9	-	-
SMVLH5	Avg. Checks	1039	1062	2865	2412	28987	25664
	Avg. Cost	3.4	12.0	4.7	14.7	9.0	34.8
SMVLH10	Avg. Checks	1498	1833	5047	4164	50560	43619
	Avg. Cost	3.3	11.3	4.4	13.7	8.6	31.0

TABLE 1. Average number of consistency checks and solution costs for 6 sets of PCSPs.

Table 1 presents results comparing four relaxation algorithms with the PEFC3 algorithm. Each set contains 10 problems. Problem Sets 1, 3 and 5 were used in [Freuder 92] and contain problems with 10, 12, and 16 variables respectively. The cost of not satisfying a constraint is 1. Sets 2, 4, and 6 were created from Sets 1, 3, and 5 respectively by randomly assigning the cost of an unsatisfied constraint on the [1,9] interval. The relaxation algorithms are instantiations of the propagation-based relaxation schema described above, with varying parameters.[8] The results presented are the average number of consistency checks required for a solution to be found for problems in each set and the average cost of the solutions that were found. No results could be found for the 16-variable problems with the MMV relaxation algorithms due to exponential growth in memory use.

The weakness of the both the relaxation algorithms and the PCSP algorithms is their exponential complexity. None of these algorithms will scale-up to problems much larger than 15 variables.

7. Cycle detection is addressed in [Beck 94] .
8. For further information see [Beck 94] .

5.0 Mediated Constraint Relaxation in the Supply Chain

5.1 Constraint Graph Generation

At the logistics level, the normal mode of execution surrounds the scheduler's acceptance a new order and modification of the existing logistics-level schedule to meet the order. Each factory receives its new partial schedule and modifies its factory-level schedule accordingly. As long as the operation of the enterprise does not violate the assumptions and aggregate information upon which the various schedules are based, no coordination is necessary. When an unexpected event occurs that prevents an agent from satisfying a constraint, the agents must coordinate their efforts to respond optimally.

The first step in the coordination is the notification of the mediator that an agent cannot meet its commitments. The current schedule is a partial representation of the global constraint graph, but will typically not represent all relaxation information. The mediator creates the augmented constraint graph by requesting information of other agents and adding it to the schedule constraint graph. The new information is composed of constraints that are not represented at the higher-level of abstraction, changes to the higher-level constraints (e.g. the throughput of a factory is modified when a machine in the factory breaks down), and the cost of relaxing constraints. The mediator must know where it can find information on the relaxations of these constraints. Since the base constraint graph is a schedule there may be relaxations of any of the following relationships:

- the precedence constraints between activities
- the resource requirements of an activity
- the release date of an order
- the due date of an order

The precedence constraints for an order are represented in the process plan defined in the partial enterprise model used by the scheduler. The mediator can access the process plan directly if it is the scheduler or it can ask the scheduling agent for the possible relaxations on the precedence constraints. The resource requirements of an activity are also represented in the process plan, however the resource substitutions may not be. The Resource Manager is responsible for this information. The release date of an order is often constrained by resource availabilities (as well as the current time) so the Resource Manager will model possible relaxations of the release date. Finally, the due date is constrained by the customer's wishes, which are represented inside the supply chain by Order Acquisition. The mediator requests due date information and relaxation cost from OA which may or may not query the customer about possible changes.

In general, the mediator must know where to find relaxation information. The above description is not limiting as additional constraints can be easily handled by assigning an agent to be responsible for the relaxation information and by informing the mediator of the existence of the agent. A case in point is the role of the Transportation Manager (TM) which we will not include in our example. If we are modeling transportation, we need to insert transportation activities between some production activities. The TM has knowledge about possible relaxations to the constraints resulting from transportation alternatives (e.g. the resource shipment can travel by plane instead of truck and arrive much sooner at an increased cost).

5.2 An Example

In our example above, we described the reaction necessary by the supply chain of Canfurn, Inc. to a requested change. Here we will present a particular instantiation of a supply chain simulation and show the actions taken by each agent in reaction to an order change request.

In our supply chain instantiation, we will use the agents described above at the logistics level (Logistics, Resource Manager (RM), and Order Acquisition (OA)) plus three factory agents: Factory$_1$ (with machines: M_{11}, M_{12}), Factory$_2$ (with machines: M_{21}, M_{22}), and Factory$_3$ (with machines: M_{31}, M_{32}).

There are currently three orders each with a release date of t_0 and due date of t_3. Each order consists of three unit-duration activities that each use one unit-capacity machine. The schedule is shown in Table 2.[9]

Start Times	**Machines**					
	M_{11}	**M_{12}**	**M_{21}**	**M_{22}**	**M_{31}**	**M_{32}**
t_0		A_{11}				A_{13}
t_1	A_{31}	A_{12}	A_{22}		A_{21}	
t_2		A_{23}			A_{32}	
t_3			A_{33}			

TABLE 2. Current Schedule in the Supply Chain Simulation

At time t_{-1}, the customer for O_3 requests a change requiring A_{13} to be executed on M_{12}. The OA agent is the first to know about the customer's wishes and notifies Logistics. Logistics investigates the following alternatives:

1. Move activities A_{11} and A_{12} to another machine.
2. Delay delivery of some products.

In order to answer these questions, the mediator requests information from other agents. A representation of the communication for option 1 is as follows:

- Logistics to RM: "What machines can be substituted for M_{12}?"
- RM to Logistics: "M_{22}, but quality will decrease."[10]
- Logistics to OA: "What is the cost of the reduced quality on O_2 and O_1?"
- OA to Logistics: "3 each."

For option 2, Logistics must add the cost of late delivery of each order:

- Logistics to OA: "What is the cost of late delivery on each order?"
- OA to Logistics: "O_1: 5 per time unit, latest acceptable delivery is t_6."
- OA to Logistics: "O_2: 5 per time unit, latest acceptable delivery is t_6."
- OA to Logistics: "O_3: 2 per time unit, latest acceptable delivery is t_6."

9. Note that the due date constraint on O_3 is already relaxed at some cost.
10. Work is progressing on a theory of quality [Kim 94] that will make this answer meaningful.

The information from other agents allows the mediator to form the augmented constraint graph upon which it can execute a relaxation algorithm. Relaxation will not simply choose one of the options, but rather investigate possibilities of a combination of moving some activities and delaying delivery of some products.

6.0 Conclusion

We have presented an mediated approach to multiagent coordination using constraint relaxation. Having modeled the interactions of the agents as a commitment/constraint graph, deviations expected events may produce an overconstrained situation. A mediator gathers information in the form of variables, constraints, and relaxation costs and executes a constraint relaxation algorithm on the augmented graph. The results is a near optimal modification to the constraints to re-establish a feasible graph.

Mediation is one end of a spectrum of coordination techniques. Further work will attempt to move toward a more distributed approach by the combination of negotiation and mediation. The mediation approach will provide empirical data against which future, more distributed approaches can be compared. We are pursuing the development of a theory of coordination where the structure of the problem can be correlated with different coordination algorithms (e.g. at different places along the negotiation/mediation spectrum) resulting in the ability to choose the algorithm that will likely be most efficient for the problem at hand.

The use of constraint relaxation as a coordination operator builds on the power and expressivity of the constraint model. With the appropriate modeling of problems in the constraint formalism, properties of the structure of the problem can be assessed via texture measurements. As mentioned above, the algorithms applied to the PCSP problems are all exponential in complexity and can not be successfully applied to larger problems. We view texture measurements on the graph as an important tool providing information upon which dynamic heuristic decisions can be based. Preliminary work with schedule optimization problems show promising results on graphs of over 150 variables.

7.0 Acknowledgments

This research is supported, in part, by Manufacturing Research Corporation of Ontario, Natural Science and Engineering Research Council, Digital Equipment Corp., Micro Electronics and Computer Research Corp., and Spar Aerospace.

8.0 References

[Beck 94] Beck, J.C. *A Schema for Constraint Relaxation with Instantiations for Partial Constraint Satisfaction and Schedule Optimization.* Master's thesis, University of Toronto, 1994. To Appear.

[Durfee 91] Durfee, E.H. and Montgomery, T.A. Coordination as Distributed Search in a Hierarchical Behavior Space. *IEEE Transactions on Systems, Man, and Cybernetics.* SMC-21(6):1361-1378, 1991.

[Durfee 87a] Durfee, E.H. and Lesser, V.R. Using Partial Global Plans to Coordinate Distributed Problem Solvers. *Proceedings of IJCAI-87,* pages 875-883. 1987.

[Durfee 87b] Durfee, E.H., Lesser, V.R., and Corkill, D.D. Cooperation Through Communication in a Distributed Problem Solving Network. In Michael N. Huhns (Ed.), *Distributed Artificial Intelligence.* Volume 1. Pitman Publishing & Morgan Kaufmann Publishers, 1987, pages 29-58, Chapter 2.

[Fox 92] Fox, M.S. *Integrated Supply Chain Management.* Technical Report, Enterprise Integration Laboratories, Department of Industrial Engineering, University of Toronto, April, 1992.

[Fox 90] Fox, M.S. and Sadeh, N. Why Is Scheduling Difficult? A CSP Perspective. *Proceedings of the Ninth European Conference on Artificial Intelligence*. 1990.

[Fox 89] Fox, M.S., Sadeh, N., and Baykan, C. Constrained Heuristic Search. *Proceedings of IJCAI-89*. 1989.

[Freuder 92] Freuder, E. and Wallace, R. Partial Constraint Satisfaction. *Artificial Intelligence.* Volume 58, pages 21-70, 1992.

[Kim 94] Kim, H.K. and Fox, M.S. Formal Models of Quality and ISO 9000 Compliance: An Information Systems Approach. *48th Annual Quality Congress of the American Society of Quality Control*. 1994.

[Lesser 81] Lesser, V.R and Corkill, D.D. Functionally Accurate, Cooperative Distributed Systems. *IEEE Transactions on Systems, Man, and Cybernetics.* SMC-11(1):81-96, January, 1981.

[Nagel 91] Nagel, R.N. et al. *21st Century Manufacturing Enterprise Strategy: An Industry Led View.* Technical Report, Iacocca Institute, Lehigh University, Bethlehem PA, 1991.

[Sadeh 91] Sadeh, N. *Lookahead Techniques for Micro-Opportunistic Job Shop Scheduling.* PhD thesis, Carnegie Mellon University, 1991. CMU-CS-91-102.

[Sadeh 89] Sadeh, N. and Fox, M.S. *Preference Propagation in Temporal/Capacity Constraint Graphs.* Technical Report CMU-RI-TR-89-2, The Robotics Institute, Carnegie Mellon University, January, 1989.

[Sathi 89] Sathi, A. and Fox, M.S. Constraint-Directed Negotiation of Resource Reallocations. In Michael N. Huhns and Les Gasser (Eds.), *Distributed Artificial Intelligence.* Volume 2. Pitman Publishing & Morgan Kaufmann Publishers, 1989, pages 163-193, Chapter 8.

[Smith 89] Smith, S.F., Ow, P.S., Matthys, D.C., and Potvin, J.Y. OPIS: An Opportunistic Factory Scheduling System. *Proceedings of International Symposium for Computer Scientists*, 1989.

[Sycara 91] Sycara, K., Roth, S., Sadeh, N., and Fox, M. Distributed Constrained Heuristic Search. *IEEE Transactions on Systems, Man, and Cybernetics.* SMC-21(6):1446-1461, 1991.