

REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-04-

0486

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE Final		3. DATES COVERED (From - To) 1 January 2002 - 31 December 2004	
4. TITLE AND SUBTITLE Command and Control: Planning The Organization and Flow of Information and Decisions in Complex Operations				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER F49620-02-1-0135	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Warren B. Powell				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Princeton University Department of Operations Research and Financial Engineering Princeton, NJ 08544				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research 4015 Wilson Blvd Mail Room 713 Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S) AFOSR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Distribution Statement A. Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT We have focused on the development of a class of technologies which we group under the name "optimizing simulator." This strategy simulates decisions with up to four classes of information. Two of the classes require the use of iterative learning strategies. One captures the value of resources in the future, while the second tries to match decisions made by the model with exogenously specified patterns of behavior.					
<div style="font-size: 2em; font-weight: bold; margin: 0;">20040930 025</div>					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Warren B. Powell
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code)



Castle Laboratory

Department of Operations Research and Financial Engineering
Princeton University

Final Project Report:

Command and control: Planning the organization and flow of information and decisions in complex operations

Principal Investigator:

Warren B. Powell
Princeton University
Department of Operations Research and Financial Engineering
Princeton, NJ 08544
powell@princeton.edu

Prepared for:

Air Force Office of Scientific Research
Grant: F49620-02-1-0135

September, 2004

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

Table of Contents

1. Summary of activities	1
2. The "Optimizing Simulator"	2
2.1. Modeling the organization and flow of information and decisions	2
2.2. Results of calibration exercises	3
2.3. Selected papers	3
3. Sensitivity results from simulations	4
4. Approximate dynamic programming for high-dimensional problems	5
4.1. The three curses of dimensionality	5
4.2. A solution strategy	5
4.3. Performance	6
4.4. Publications	6
5. Two-stage resource allocation problems	7
5.1. A separable approximation algorithm	7
5.2. Convergence theory	8
5.3. Publications	8
6. Multistage management of simple assets	9
7. Multistage management of complex assets	10
8. Adaptive learning for batch processes	12
9. Multi-agent decision making	14
10. Research reports sponsored by AFOSR	16
10.1. Multicommodity and heterogeneous resource allocation problems	16
10.2. Multiagent modeling and algorithms	16
10.3. Discrete routing and scheduling	16
10.4. Stochastic optimization	16
10.5. Optimizing to match low-dimensional patterns	17
10.6. Dynamic batch service problems	17
10.7. Implementation research	17
10.8. Book chapters	17
10.9. Books	18
10.10. Doctoral dissertations	18
11. Personnel supported	19
12. Interactions/transitions	20
12.1. Participation/presentations at meetings, conferences, etc	20
12.1.1. Invited talks:	20
12.1.2. Conference presentations:	21
12.2. Consultative and advisory functions	23
12.3. Transitions	23

1. Summary of activities

We have focused on the development of a class of technologies which we group under the name “optimizing simulator.” This strategy simulates decisions with up to four classes of information. Two of the classes require the use of iterative learning strategies. One captures the value of resources in the future, while the second tries to match decisions made by the model with exogenously specified patterns of behavior.

Our research has followed several themes. These include:

- The creation of a formal paradigm for resource allocation, using the framework of the “dynamic resource transformation problem” (or DRTP) framework. This framework produced two new problem classes, and creating a modeling framework built around the modeling of information and decisions.
- Approximate dynamic programming for asset management – We have combined the power of dynamic programming with the tools of math programming, allowing us to solve large scale, complex dynamic resource allocation problems. This strategy involves formulating Bellman’s optimality equations in a nonstandard way, and introduces a computationally tractable class of value function approximation strategies.
- Using the value function approximations, we show how to find derivatives of simulations, avoiding the use of noisy “what if” analyses for certain types of questions.
- We have studied in depth the theoretical properties of separable approximations, which are particularly useful for discrete resource allocation problems. These techniques provide near optimal solutions with faster rates of convergence than previously available techniques.
- We have shown that simple linear approximations work quite well for batch processes, scaling easily to large numbers of customer classes. This result opens the door to providing high quality solutions for large-scale networks that previously have been completely intractable.
- We have shown how these methods generalize to multiagent systems, and have introduced a strategy that dramatically improves the rate of convergence.



These techniques are capable of modeling very complex problems, as shown in our work with freight transportation companies. The explicit modeling of information allows us to answer questions about the value of information and technologies that improve the quality of information. We can also study questions addressing the organization of information and decisions.

2. The “Optimizing Simulator”

There are two classes of modeling technologies that have been applied to the military airlift problem: simulation models, which can capture a high degree of realism, and optimization models (linear programming) which sacrifice detail for problem structure which enables algorithms that provide optimal solutions and important sensitivity results.

2.1. Modeling the organization and flow of information and decisions

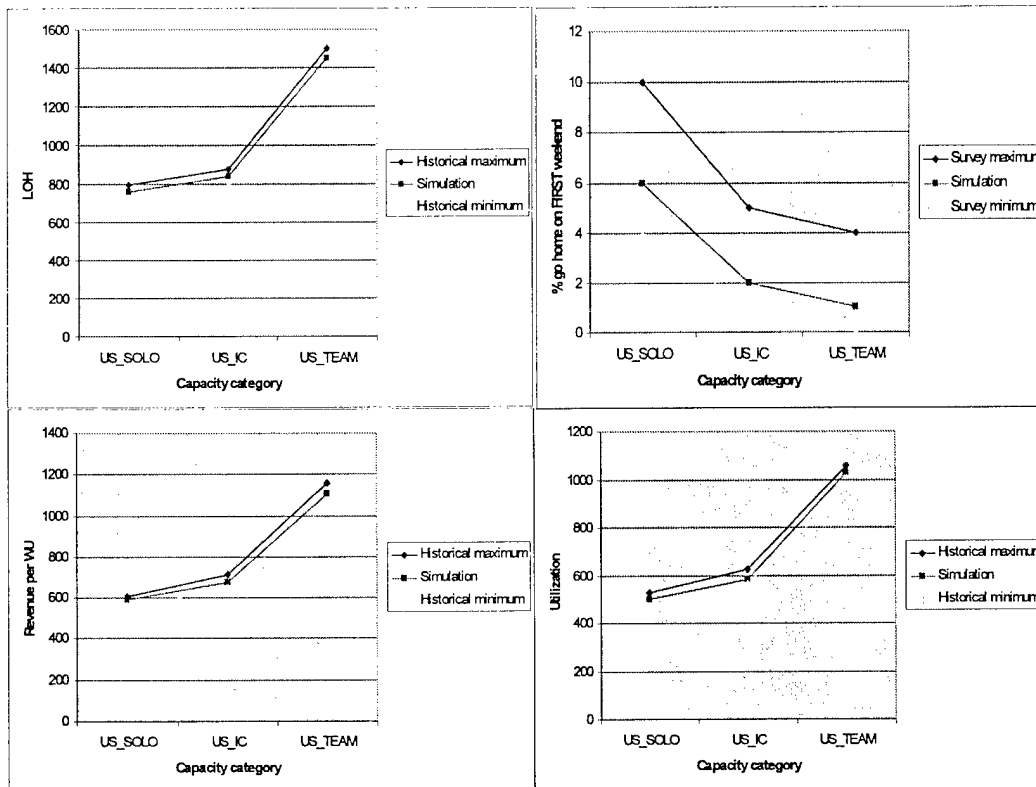
CASTLE Lab has pioneered the use of the “optimizing simulator” as a modeling and algorithmic paradigm for its use in military airlift problems and a variety of civilian freight applications. This technology simulates decision making with up to four classes of information: 1) knowledge (what we know about our system now), 2) forecasts of future *exogenous* information (demands, weather delays, equipment failures), 3) forecasts of the impact of decisions now on future, or the impact of decisions made by one agent on another, and 4) forecasts of decisions themselves. By explicitly modeling the information available to make a decision, we can test the value of information. The first three classes work to improve the quality of a solution measured in terms of a quantifiable objective function (illustrated in the figure below). The fourth class aims to make the simulator more realistic in the eyes of a knowledgeable expert.



The optimizing simulator concept also produces some practical benefits. Traditional simulators can require a fair amount of tuning when first applied to new situations. The optimization logic typically provides reasonable solutions very quickly by taking advantage of a cost function that helps to make complex tradeoffs. The fourth information class, the low dimensional patterns, still allows a knowledgeable user to influence the behavior of the model to produce more realistic simulations.

2.2. Results of calibration exercises

The optimizing simulator has been applied to two railroads for locomotive management, a car distribution problem at a railroad, and most recently a major fleet management problem at Schneider National, the largest truckload motor carrier in the United States. As with the other projects, we undertook an intensive, year-long calibration exercise at Schneider to ensure that the model was accurately reproducing the behavior of the company. The company chose a series of operating statistics and produced, from history, upper and lower limits for these measures. The challenge of the model was to produce numbers that fell within these ranges.



2.3. Selected papers

Wu, T. T., W. B. Powell and A. Whisman, "The Optimizing Simulator: An Intelligent Analysis Tool for the Military Airlift Problem" Working paper.

Powell, W.B., T. Wu, H. P. Simao and A. Whisman, "Using Low Dimensional Patterns in Optimizing Simulators: An Illustration for the Military Airlift Problem," *Mathematical and Computer Modeling*


Shapiro, J. and W.B. Powell, "A Metastrategy for Dynamic Resource Management Problems based on Informational Decomposition," *Inform Journal on Computing* (to appear).

Powell, W.B., "Dynamic Models of Transportation Operations," *Handbooks in Operations Research and Management Science: Supply Chain Management*, (S. Graves and T. A. G. de Tok, eds.), Elsevier, Amsterdam, pp. 677-756, 2003.

3. Sensitivity results from simulations

One appeal of linear programming models is sensitivity analysis – the ability to understand the marginal impact of increasing different types of resources (aircraft, airbase capacity) without having to run difficult “what if” analyses (small changes to inputs to a simulator can produce random responses). The problem is that making small changes to inputs and rerunning any simulator (including an optimizing simulator) can produce fairly noisy outputs. It was very important to run multiple samples and average across the runs, a process that could be extremely time consuming.

The figure below shows a base case (where we plotted costs from each of the last eight iterations as a measure of the variability of the base case), and three scenarios where we added 50, 100 and 200 drivers. For each scenario, we ran eight repetitions with different samples. The purple line shows an average through these points, and the red line shows the estimate of the slope we obtained from gradients produced by the base case. This analysis shows that the gradients accurately match the marginal value of resources produced by the “what if” approach.



These runs show that we can obtain gradient estimates from an optimizing simulator, avoiding the need to perform “what if” analyses for resources (e.g. the value of more cargo aircraft). These gradients can also be used to produce estimates of gradients of other capacity constraints such as airbases, although we have not tested this yet.

Selected papers include:

H. Topaloglu and W. B. Powell, “Sensitivity Analysis of a Dynamic Programming Approximation-Based Dynamic Resource Allocation Policy”, Working paper, 2004.

Topaloglu, H. and W.B. Powell, “Dynamic Programming Approximations for Stochastic, Time-Staged Integer Multicommodity Flow Problems,” *INFORMS Journal on Computing*, (to appear).

4. Approximate dynamic programming for high-dimensional problems

A major contribution of CASTLE Laboratory has been the integration of math programming, with its power for solving high dimensional optimization problems, and dynamic programming, which allows us to model the sequencing of information and decisions over time. The techniques are central to the optimizing simulator, as they combine stepping forward through time (as we would in a simulation) while “learning” the value of being in a state from previous decisions.

4.1. The three curses of dimensionality

Fundamental to dynamic programming is the solution of Bellman’s optimality equation:

$$V_i(R_i) = \max_{x \in X} c_i(R_i, x_i) + E\{V_{i+1}(R_{i+1}) | R_i\}$$

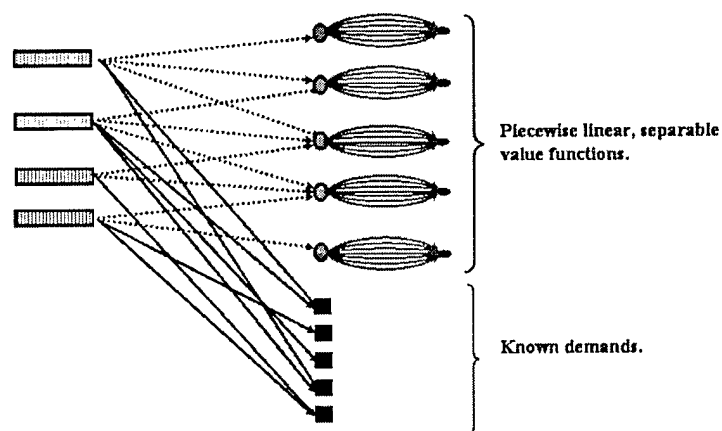
The equation requires that we compute $V_i(R_i)$ for each possible value of the (resource) state variable R_i . If R_i is a vector, we face a computationally intractable problem.

For resource allocation problems, there are actually three curses of dimensionality: the state space (when there are different types of aircraft, in different locations, with different attributes), the outcome space (different types of exogenous information such as demands, weather and breakdowns) and the action space (the vector of all the things you can do with an aircraft). In other words, even for a single instance of R_i , we face the daunting challenge of computing the expectation and finding the best decision.

4.2. A solution strategy

We solved the “three curses” problem using several tricks. First, we formulate the well-known Bellman optimality equations around a novel version of the state variable. Normally, Bellman’s equations are written using the state of the system right before we make a decision. Instead, we capture the state of the system right after we make a decision, which makes our “post-decision” state variable a deterministic function of our decision.


Next, we replace the value function with specially chosen approximations. It is important to find approximations that do not destroy problem structure. We have found approximations that require that we solve sequences of fairly small linear programs, where finding



integer solutions is especially easy. These linear programs sometimes look like the network shown above.

4.3. Performance

We have run this logic on deterministic problems where we can obtain optimal solutions. Solutions within one percent of optimal are common. In addition, we always return integer solutions, something that can be quite difficult with standard solvers.



On problems where we do not know the future, we can outperform the deterministic solution by five or ten percent or more. These results have been found both on laboratory datasets, and problems based on actual fleet management problems in rail and trucking.

4.4. Publications

J. Si, A. Barto, W.B. Powell and D. Wunsch (eds.), *Learning and Approximate Dynamic Programming: Scaling up to the Real World*, John-Wiley and Sons, New York, 2004.

Powell, W.B. and B. van Roy, "Approximate Dynamic Programming for High Dimensional Resource Allocation Problems, (in *Learning and Approximate Dynamic Programming: Scaling up to the Real World*, J. Si, A. Barto, W.B. Powell and D. Wunsch, eds.), John-Wiley and Sons, New York, 2004.

Godfrey, G. and W.B. Powell, "An Adaptive Dynamic Programming Algorithm for Single-Period Fleet Management Problems I: Single Period Travel Times," *Transportation Science*, Vol. 36, No. 1, pp. 21-39 (2002).

Godfrey, G. and W.B. Powell, "An Adaptive Dynamic Programming Algorithm for Single-Period Fleet Management Problems II: Multiperiod Travel Times," *Transportation Science*, Vol. 36, No. 1, pp. 40-54 (2002).

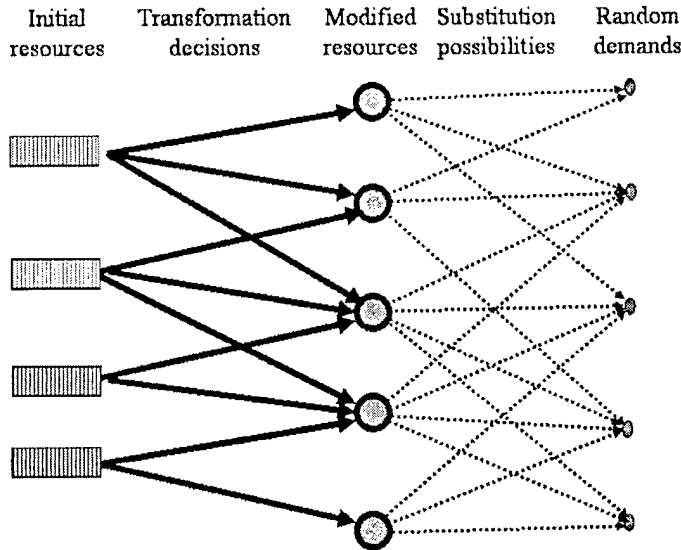
Topaloglu, H. and W.B. Powell, "Dynamic Programming Approximations for Stochastic, Time-Staged Integer Multicommodity Flow Problems," *Inform Journal on Computing*, (to appear).

Topaloglu, H. and W.B. Powell, "An Algorithm for Approximating Piecewise Linear Concave Functions from Sample Gradients," *Operations Research Letters*, Vol. 31, No. 1, pp. 66-76 (2003).

5. Two-stage resource allocation problems

Fundamental to our work on approximate dynamic programming has been our study of two-stage resource allocation problems.

This basic problem arises in a vast array of settings. The basic problem requires that we first make a decision to design/purchase/allocate a resource that may be an aircraft, fuel depots, complex equipment such as components for the electric power grid, as well as sensors and detectors. After we make the decision, we learn the demand for different



types of resources or assets, including the ability to make substitutions (at a cost). The problem is depicted below.

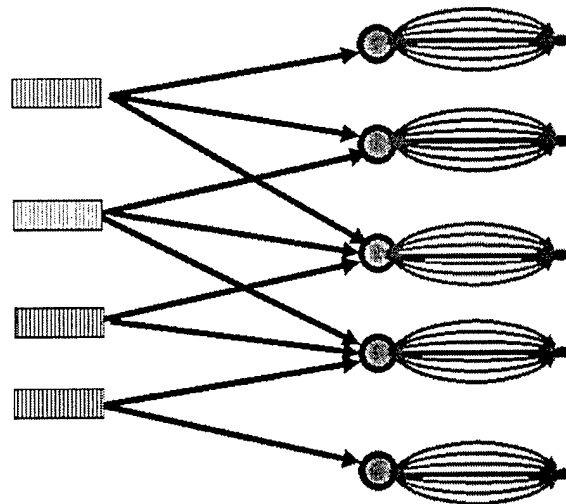
These problems have long been recognized by the stochastic programming community under the name of "two stage stochastic programs with network recourse." The most popular solution strategy has been to approximate the second stage using a series of "cuts" – inequalities that

successfully improve the approximation of the second stage.

5.1. A separable approximation algorithm


We have introduced and studied in depth the use of piecewise linear, separable approximations. The result is sequences of relatively simple linear programs. A major property is the ease with which we can obtain integer solutions. By constructing piecewise linear functions, we produce approximations that make it extremely easy to obtain integer solutions, something that can be quite difficult with other stochastic programming methods.

A central feature of this strategy is that the recourse functions are concave. We estimate these functions by first solving the allocation problem using the current approximations of the second stage. We then draw a Monte Carlo sample of the second stage and solve a deterministic problem. This problem yields estimates of the value of an additional unit of each



type of resource, which we use to help build up an approximation of the second stage.

The algorithmic strategy reduces a very complex problem into a sequence of very simple allocation problems that can be solved using standard linear programming packages. The method is particular well suited to allocating discrete assets: where to locate a sensor, where to locate a facility, where to send an aircraft, how many aircraft to purchase. Obtaining integer solutions with this class of approximations is quite easy and does not require specialized algorithms. The algorithm also seems to exhibit much faster rates of convergence than competing techniques such as Benders (see graph below).



5.2. Convergence theory

For this fundamental problem class, we have been able to obtain several important convergence results. We have proven that several methods for maintaining concavity still allow us to converge to the current concave function. We

have proven convergence of a separable approximation for continuous resources (e.g. fuel). We have shown that for discrete problems, sequences of piecewise linear approximations will produce optimal solutions when the underlying problem is separable, even when we start with arbitrarily bad initial approximations. The algorithm is not provably convergent when the second stage is nonseparable, but extensive experimental work has shown that the solutions are extremely close to optimality, and provide very fast convergence.

5.3. Publications

Powell, W.B., A. Ruszczyński and H. Topaloglu, "Learning Algorithms for Separable Approximations of Stochastic Optimization Problems," *Mathematics of Operations Research* (to appear).

Topaloglu, H. and W.B. Powell, "An Algorithm for Approximating Piecewise Linear Concave Functions from Sample Gradients," *Operations Research Letters*, Vol. 31, No. 1, pp. 66-76 (2003).

Godfrey, G. and W.B. Powell, "An Adaptive, Distribution-Free Algorithm for the Newsvendor Problem with Censored Demands, with Application to Inventory and Distribution Problems," *Management Science*, Vol. 47, No. 8, pp. 1101-1112, (2001).

Powell, W. B. and H. Topaloglu, "Stochastic Programming in Transportation and Logistics," *Handbooks in Operations Research and Management Science: Stochastic Programming* (A. Shapiro and A. Ruszczyński, eds.), Elsevier, Amsterdam, pp. 555-635, 2003.

6. Multistage management of simple assets

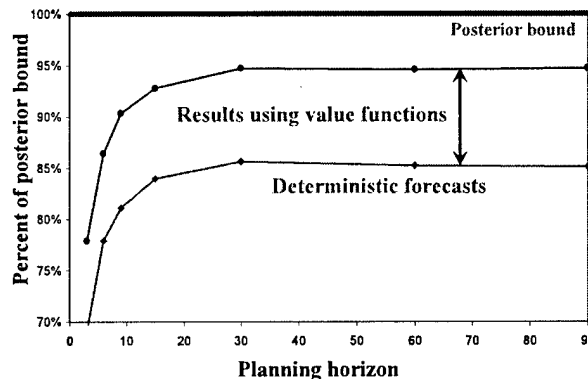
Simple assets are those where the attributes of the asset are not too complicated, which means there are not very many different types of assets. If we are modeling box cars for a railroad, the attribute of a box car can be represented by the type of box car and its location. A real problem might involve allocating 10 types of box cars between several hundred locations, producing several thousand box car attributes. Similar problems arise in the management of intermodal containers, truck trailers, or the management of people or equipment who can be characterized by relatively simple vector of attributes. The key characteristic is that we can handle thousands of attributes, but not millions.

These problems may be characterized by a variety of uncertainties:

- What are the future demands for these resources?
- Weather delays or other forms of variability in transportation movements.
- Equipment failures.
- New equipment entering the system, or equipment being taken from the system.

These problems are typically solved using models which either ignore the future, or which make best estimates of the future (“point forecasts”) and then solve the problem as if these estimates are going to come true.

These multistage problems can be solved as sequences of two-stage problems using the methods described above. Instead of solving one large problem, we step through time, solving sequences of small problems. The result is much higher solution quality (see figure), but the technique has other useful features:



- Explicit modeling of uncertainty – we can run simulations testing the effect of reducing the level of variability (for example, more reliable aircraft, improved forecasting of future requirements).
- The model produces an estimate of the value of each type of resource without having to run “what if” scenarios.
- It provides a very small and compact model for solving real-time allocation problems.

Sample papers include:

Godfrey, G. and W.B. Powell, “An Adaptive Dynamic Programming Algorithm for Single-Period Fleet Management Problems I: Single Period Travel Times,” *Transportation Science*, Vol. 36, No. 1, pp. 21-39 (2002).


Godfrey, G. and W.B. Powell, “An Adaptive Dynamic Programming Algorithm for Single-Period Fleet Management Problems II: Multiperiod Travel Times,” *Transportation Science*, Vol. 36, No. 1, pp. 40-54 (2002).

7. Multistage management of complex assets

In real applications, assets often take on a variety of attributes. We may start by assuming that we are managing eight or ten aircraft types, but over time we learn that there are different configurations, they may be loaded or empty, we may have to capture maintenance issues as well as fuel status. People and complex equipment such as aircraft always qualify as “complex assets”, but our experience is that even relatively simple pieces of equipment can sometimes take on a richer set of attributes as a project matures.




We can represent a complex piece of equipment using an attribute vector $a = (a_1, a_2, \dots, a_N)$ where the elements a_n may describe the location, type, status, direction, readiness, or any of a host of items. Representing resources with a general vector of attributes introduces tremendous flexibility in the modeling of complex assets, especially when we do not know what attributes are important at the beginning of a study. The problem is that we no longer assume that we can enumerate all possible values of a . This introduces a significant complication for any modeling methodology that wants to make decisions now which account for future, downstream events.



We typically face the challenge of estimating the value v_a of an asset with attribute a . When the number of different types of assets is large, we encounter a statistical problem of estimating these values with limited information. We have solved this problem using a hierarchical estimating strategy that estimates a value at different levels of aggregation. Whenever we need an estimate of the value of an asset, we use a weighted combination of values at different levels of aggregation. These weights handle the tradeoff between statistical noise and structural error. The weights are estimated adaptively (see figure), and show that initially we put the highest weight on the most aggregate estimates, with more weight on more disaggregate estimates as the algorithm progresses. The weighted

estimates produce faster, more stable convergence than estimates at a particular level of aggregation.

The chart below shows that using a weighted combination outperforms using any single level of aggregation.



Sample papers include:


Topaloglu, H. and W.B. Powell, "Dynamic Programming Approximations for Stochastic, Time-Staged Integer Multicommodity Flow Problems," *Informs Journal on Computing*, (to appear).

Powell, W.B., J. Shapiro and H. P. Simao, "An Adaptive, Dynamic Programming Algorithm for the Heterogeneous Resource Allocation Problem," *Transportation Science*, Vol. 36, No. 2, pp. 231-249 (2002).

George, A., W.B. Powell and S. Kulkarni, "Value Function Approximation Using Hierarchical Aggregation for Multiattribute Resource Management," working paper.


8. Adaptive learning for batch processes

Most of our work has focused on “coupling problems”: coupling one pilot with an aircraft, couple a load of freight with a truck or aircraft, couple a locomotive with a train. A different problem class arises with quantities of freight and people have to be loaded into a large container. These problems have proven to be intractably difficult for classical optimization methods, and tend to be solved using myopic policies in a simulation framework (dispatch the vehicle when it is at least X percent full if it has been waiting longer than Y).



Using a single-link example, we found that the value function has a monotone, K -convex shape as shown to the left. This suggested two approximation strategies: linear, and concave (capturing the behavior when the number of customers is less than the capacity of the vehicle, which is the most common states to visit). Compared to the optimal solution, which we can find if there is only one type of customer, we found that both methods work well. The simpler linear approximation works better initially (there is only one parameter to estimate), while the nonlinear worked better in the limit (see below). The linear approximation worked relatively better on stochastic problems. When there is uncertainty, the expected value function becomes smoother and closer to a linear approximation.

The real value of these approximations is that they scale to problems with many customer classes, a problem that cannot be solved exactly. For a problem with multiple customer types we proved optimality of a particular class of greedy policies, and compared the approximate dynamic programming strategy against the best myopic policy we could design. The strategy was tested on a problem with 100 customer classes (although there was no real limit to this) and showed that it consistently outperformed the myopic heuristic which never outperformed the approximate DP strategy, and only performed comparably when the cost of holding a customer was very low.



Publications include:

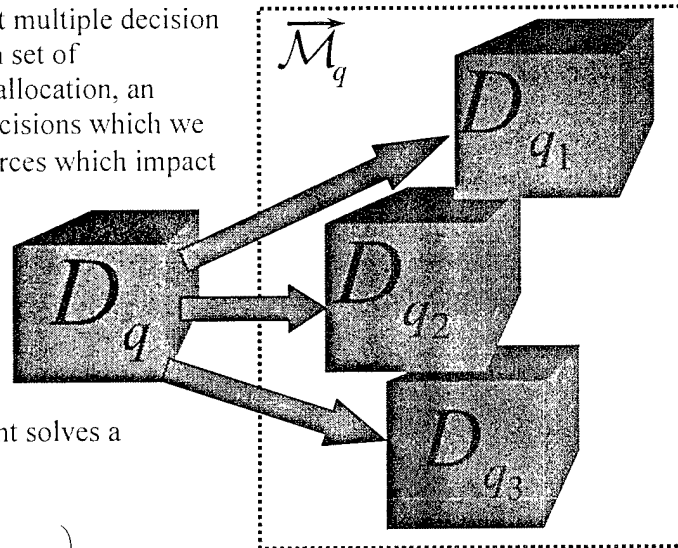
Dall'Orto, L. C., T. G. Crainic, J. E. Leal and W. B. Powell, "The Single-Node Dynamic Service Scheduling and Dispatching Problem," *European Journal of Operational Research* (to appear)

Papadaki, K. and W.B. Powell, "An Adaptive Dynamic Programming Algorithm for a Stochastic Multiproduct Batch Dispatch Problem," *Naval Research Logistics*, Vol. 50, No. 7, pp. 742-769, 2003.

Papadaki, K. and W.B. Powell, "Exploiting structure in adaptive dynamic programming algorithms for a stochastic batch service problem," *European Journal of Operational Research*, Vol. 142, No. 1, pp. 108-127, 2002.

9. Multi-agent decision making

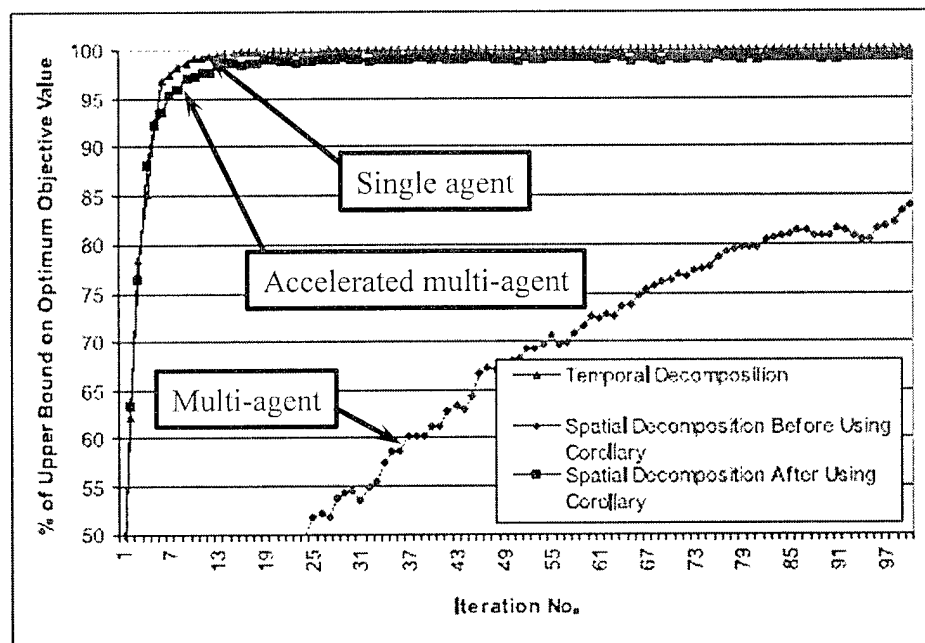
Virtually all complex problems exhibit multiple decision makers, each characterized by the own set of information. In the arena of resource allocation, an agent, call it q , may control a set of decisions which we denote by D_q , which may move resources which impact another agent, which we might call q' . We call the set of agents q' that we may impact the *forward reachable set*, denoted by $\vec{\mathcal{M}}_q$. We can capture these interactions using an approximate dynamic programming framework by assuming that each agent solves a problem of the form:



$$X_q(I_q) = \max_{x_q \in \mathcal{X}_q} \left(C_q(x_q) + \theta \sum_{q' \in \vec{\mathcal{M}}_q} V_{qq'}(x_{qq'}) \right)$$

The function $V_{qq'}(x_{qq'})$ captures the effect of decisions by q on q' .

If we assume that agents all make their decisions in a particular sequential order, we can apply the same approximate dynamic programming methods we have been developing, using the agent index q as we did earlier with time. It is important to emphasize the



double indexing of the value function. We found, however, that a straightforward application of this method was exceptionally slow. In the graph, we show the convergence if we use a single agent (the red line that is highest) versus a straightforward implementation of approximate DP (the lowest line). We found an accelerated strategy (labeled "accelerated multi-agent") which was almost as fast as the original single agent formulation.

Unresolved is the effect of asynchronous solution. Initial experiments indicated that asynchronous solution produced much poorer results. Additional research is needed to study the effects of asynchronous solution and communication.

Our papers on this topic include:

Shapiro, J. and W.B. Powell, "A Metastrategy for Dynamic Resource Management Problems based on Informational Decomposition," *Inform Journal on Computing* (to appear).

Topaloglu, H. and W.B. Powell, "A Multi-Agent Decision Making Structure for Dynamic Resource Allocation with Nonlinear Functional Approximations," *Operations Research* (to appear).

10. Research reports sponsored by AFOSR

10.1. Multicommodity and heterogeneous resource allocation problems

- Topaloglu, H. and W.B. Powell, "Dynamic Programming Approximations for Stochastic, Time-Staged Integer Multicommodity Flow Problems," *Inform Journal on Computing*, (to appear).
- Powell, W.B., J. Shapiro and H. P. Simao, "An Adaptive, Dynamic Programming Algorithm for the Heterogeneous Resource Allocation Problem," *Transportation Science*, Vol. 36, No. 2, pp. 231-249 (2002).
- Godfrey, G. and W.B. Powell, "An Adaptive Dynamic Programming Algorithm for Single-Period Fleet Management Problems I: Single Period Travel Times," *Transportation Science*, Vol. 36, No. 1, pp. 21-39 (2002).
- Godfrey, G. and W.B. Powell, "An Adaptive Dynamic Programming Algorithm for Single-Period Fleet Management Problems II: Multiperiod Travel Times," *Transportation Science*, Vol. 36, No. 1, pp. 40-54 (2002).

10.2. Multiagent modeling and algorithms

- Shapiro, J. and W.B. Powell, "A Metastrategy for Dynamic Resource Management Problems based on Informational Decomposition," *Inform Journal on Computing* (to appear).
- Topaloglu, H. and W.B. Powell, "A Multi-Agent Decision Making Structure for Dynamic Resource Allocation with Nonlinear Functional Approximations," *Operations Research* (to appear).

10.3. Discrete routing and scheduling

- Spivey, M. and W.B. Powell, "The Dynamic Assignment Problem," *Transportation Science* (to appear).
- Chen, Z.L. and W.B. Powell, "Exact Algorithms for Scheduling Multiple Families of Jobs on Parallel Machines," *Naval Research Logistics*, Vol. 50, No. 7, pp. 823-840, 2003.
- Spivey, M. and W.B. Powell, "Some Fixed Point Results for the Dynamic Assignment Problem," *Annals of Operations Research*, Vol. 124, (G. Mitra, ed.), Kluwer Academic Publishers, pp. 15-33 (2003).

10.4. Stochastic optimization

- Powell, W.B., A. Ruszczyński and H. Topaloglu, "Learning Algorithms for Separable Approximations of Stochastic Optimization Problems," *Mathematics of Operations Research* (to appear).
- Topaloglu, H. and W.B. Powell, "An Algorithm for Approximating Piecewise Linear Concave Functions from Sample Gradients," *Operations Research Letters*, Vol. 31, No. 1, pp. 66-76 (2003).

10.5. Optimizing to match low-dimensional patterns

Powell, W.B., T. Wu, H. P. Simao and A. Whisman, "Using Low Dimensional Patterns in Optimizing Simulators: An Illustration for the Military Airlift Problem," *Mathematical and Computer Modeling* 29, pp. 657-675 (2004).

10.6. Dynamic batch service problems

Dall'Orto, L. C., T. G. Crainic, J. E. Leal and W. B. Powell, "The Single-Node Dynamic Service Scheduling and Dispatching Problem," *European Journal of Operational Research* (to appear)

Papadaki, K. and W.B. Powell, "An Adaptive Dynamic Programming Algorithm for a Stochastic Multiproduct Batch Dispatch Problem," *Naval Research Logistics*, Vol. 50, No. 7, pp. 742-769, 2003.

Papadaki, K. and W.B. Powell, "Exploiting structure in adaptive dynamic programming algorithms for a stochastic batch service problem," *European Journal of Operational Research*, Vol. 142, No. 1, pp. 108-127, 2002.

10.7. Implementation research

Powell, W.B., A. Marar, J. Gelfand, and S. Bowers, "Implementing Operational Planning Models: A Case Application from the Motor Carrier Industry," *Operations Research*, Vol. 50, No. 4, pp. 571-581 (2002).

10.8. Book chapters

Powell, W.B. and B. van Roy, "Approximate Dynamic Programming for High Dimensional Resource Allocation Problems, (in *Learning and Approximate Dynamic Programming: Scaling up to the Real World*, J. Si, A. Barto, W.B. Powell and D. Wunsch, eds.), John-Wiley and Sons, New York, 2004.

Powell, W.B. and H. Topaloglu, "Fleet Management," in *Applications of Stochastic Programming*, (S. Wallace and W. Ziemba, eds.), Math Programming Society - SIAM Series in Optimization. (to appear).

Powell, W.B., "Dynamic Models of Transportation Operations," *Handbooks in Operations Research and Management Science: Supply Chain Management*, (S. Graves and T. A. G. de Tok, eds.), Elsevier, Amsterdam, pp. 677-756, 2003.

Powell, W. B. and H. Topaloglu, "Stochastic Programming in Transportation and Logistics," *Handbooks in Operations Research and Management Science: Stochastic Programming* (A. Shapiro and A. Ruszczyński, eds.), Elsevier, Amsterdam, pp. 555-635, 2003.

Powell, W.B. "Transportation and Logistics," *Handbook of Applied Optimization*, (P. M. Pardalos and M. G. C. Resende, eds.), Oxford University Press, New York. Chapter 18 (section 18.1), pp. 679-688. 2002.

Powell, W.B., B. Bouzaiene-Ayari and H.P. Simao, "Dynamic Models for Freight Transportation," *Handbooks in Operations Research and Management Science: Transportation* (G. Laporte and W.B. Powell, eds.), Elsevier, Amsterdam. (Under review).

10.9. Books

J. Si, A. Barto, W.B. Powell and D. Wunsch, (eds.) *Learning and Approximate Dynamic Programming: Scaling up to the Real World*, John-Wiley and Sons, New York, 2004.

Powell, W.B., *Approximate Dynamic Programming for Asset Management* (in preparation – used as the textbook for a combined grad/undergrad course in dynamic programming at Princeton in 2004).

10.10. Doctoral dissertations

The following doctoral dissertations were completed over the last three years.

Tony Wu, 2004, “The Optimizing Simulator for the Military Airlift Problem”

Katerina Papadaki, 2002, “Adaptive Dynamic Programming for Aging and Replenishment Processes”

Arun Marar, 2002, “Information Representation in Large-Scale Resource Allocation Problems: Theory, Algorithms and Applications.”

11. Personnel supported

Faculty:

- Professor Warren B. Powell
- Professor Andrzej Ruszczynski (visiting professor from RUTCOR).

Professional staff:

- Dr. Hugo Simao
- Dr. Belgacem Bouzaiene-Ayari

Graduate students:

- Juliana Nascimento (2nd year)
- Johannes Enders (2nd year)
- Tony Wu (6th year)
- Abraham George (5nd year)
- Katarina Papadaki (graduated 2002)
- Michael Spivey (graduated 2002)
- Arun Marar (graduated 2002)

12. Interactions/transitions

12.1. Participation/presentations at meetings, conferences, etc.

12.1.1. Invited talks:

“CI and OR: The challenge of real-time”, NSF Workshop on Cyberinfrastructure, Washington, D.C., August, 2004.

“Real-Time Optimization for Real-World Problems,” Air Mobility Command, Scott Air Force Base, June, 2004.

“The Optimizing Simulator: Modeling the Organization and Flow of Information and Decisions,” Air Mobility Command, Scott Air Force Base, June, 2004.

“Optimization Technologies for Stochastic Resource Allocation Problems,” Lawrence Livermore National Laboratories, San Francisco, CA, June, 2004.

“The "Optimizing Simulator" for Complex Dynamic Resource Allocation Problems,” Laval University, Quebec, Canada, May, 2004.

“Adaptive Learning Algorithms for Stochastic Resource Allocation,” Faculty Summit, IBM TJ Watson Research Center, May, 2004.

“Modeling Information for Freight Transportation,” Spring School on Transportation and Logistics, University of Montreal, May, 2004.

“Approximate Dynamic Programming for High Dimensional Resource Allocation Problems,” Ohio State Management Sciences Seminar, Ohio State University, April, 2004.

Tutorial: “Approximate dynamic programming for stochastic resource allocation problems,” Twente University, Netherlands, January 2004.

“The Optimizing Simulator,” Eindhoven University, Netherlands, January, 2004.

“Approximate Dynamic Programming for Dynamic Resource Allocation,” Aladdin workshop on dynamic algorithms and applications, New Orleans, January, 2004.

“Cyberinfrastructure challenges: Capturing the organization and flow of information and decisions,” talk invited by Suvrajeet Sen at NSF Grantees conference, Dallas, January, 2004.

“Approximate Dynamic Programming for High Dimensional Resource Allocation,” NSF Workshop, Virginia, November, 2003.

“Optimization Technologies for Freight Transportation,” presented to United Parcel Service Operations Research Group, Maryland, May, 2003.

“The Dynamic Assignment Problem,” Route 2003 Workshop on Vehicle Routing, Denmark, June, 2003.

“Approximate Dynamic Programming for Stochastic Resource Allocation Problems,” Invited tutorial, Canadian Operations Research Meeting, Vancouver, June, 2003.

“Adaptive Learning Strategies for Optimizing Simulators,” AFOSR Grantees meeting, Estes Park, Colorado, May, 2003.

“Real Time Optimization for Real-World Problems,” Invited speaker, Boston Chapter of Informs, March, 2003.

“Real Time Optimization for Real-World Problems,” Hong Kong University of Science and Technology, January, 2003.

“The Optimizing-Simulator,” invited presentation to OOCL, Hong Kong, January, 2003.

“Adaptive Learning Algorithms for Stochastic Resource Allocation,” Hong Kong University of Science and Technology, January, 2003

“Real Time Optimization for Real-World Operations,” Informs Practice Meeting, Montreal, May, 2002.

“Adaptive Dynamic Programming for Large-Scale Resource Allocation: Solving the three curses of dimensionality,” NSF Workshop on Learning and Approximate Dynamic Programming and NSF Workshop on the Electric Power Industry, Mexico, April, 2002.

“The Optimizing Simulator: Raising the ‘IQ’ of Airlift Simulations,” Informs Chapter presentation, Air Mobility Command, Scott Air Force Base, March, 2002.

12.1.2. Conference presentations:

“A Rail Car Distribution Model with Multiple Information Streams”, Informs, Denver, October, 2004 (with B. Bouzaiene-Ayari).

“Using Distributed Computing to Solve Large Transportation Networks”, Informs, Denver, October, 2004 (with J. Day, H. Simao, B. Bouzaiene-Ayari).

“A Distributed Decision-Making Structure for Dynamic Resource Allocation”, Informs, Denver, October, 2004 (with H. Topaloglu)

“The Optimizing Simulator for the Military Airlift Problem with MOG”, Informs, Denver, October, 2004 (with T. Wu and A. Whisman).

“A Dynamic Car Distribution Model with Multiple Lagged Information Processes,” TRISTAN V, Guadaloupe, June, 2004 (with B. Bouzaiene-Ayari and H. Topaloglu).

“Sensitivity Analysis of a Dynamic Vehicle Allocation Policy Using Approximate Dynamic Programming,” TRISTAN V, Guadaloupe, June, 2004 (with H. Topaloglu).

"Hierarchical Aggregation Techniques for Estimating Value Functions for Dynamic Management of Multiattribute Resources," TRISTAN V, Guadaloupe, June, 2004 (with A. George).

"A Scalable Approximate Dynamic Programming Algorithm for the Single-Link Dispatch Problem," Informs National Meeting, Atlanta, October, 2003 (with K. Papadaki).

"The Single-Node Dispatching Problem and Dynamic Service Network Design," Informs National Meeting, Atlanta, October, 2003 (with T. Crainic, L. Dall'Orto, J. E. Leal)

"Adaptive Learning Algorithms for the Dynamic Assignment Problem," Informs National Meeting, Atlanta, October, 2003 (with A. George).

"Adaptive, Hierarchical Algorithms for Short-Term, Tactical Forecasting," Informs National Meeting, Atlanta, October, 2003

"Dynamic Programming Approximation Techniques for Multi-stage Resource Allocation under Uncertainty," Informs National Meeting, Atlanta, October, 2003 (with H. Topaloglu)

"Solving a Large-Scale Driver Management Problem using Informational Decomposition and Data Pattern Matching," Informs National Meeting, Atlanta, October, 2003 (with H. Simao and J. Day)

"Statistical Techniques for Estimating High Dimensional Value Functions in Dynamic Programming," Informs National Meeting, Atlanta, October, 2003 (with A. George)

"The Optimizing Simulator: An Illustration for the Military Airlift Problem," Informs National Meeting, Atlanta, October, 2003 (with T. Wu)

"Separable, Piecewise-Linear Approximations for Two-Stage Stochastic Programs", Informs National Meeting, San Jose, November, 2002 (with H. Topaloglu and A. Ruszczyński).

"Rail Car Distribution Under Uncertainty Using the Separable, Projective Approximation Routine (SPAR)" Informs National Meeting, San Jose, November, 2002 (with H. Topaloglu).

"Approximate Dynamic Programming for Multistage Discrete Resource Allocation Problems," Informs National Meeting, San Jose, November, 2002 (with H. Topaloglu).

"An Information Theoretic Model of Locomotive Operations," IFORS 2002, Scotland (with Belgacem Bouzaiene-Ayari).

"A Multilayered Resource Scheduling Problem," IFORS 2002, Scotland (with Hugo Simao and Raymond Cheung).

“An Information Theoretic Approach to Modeling Car Distribution,” IFORS 2002, Scotland (with Huseyin Topaloglu and S. Melkote).

“Solving a Large Scale Driver Management Problem Using Informational Decomposition,” IFORS 2002, Scotland (with Hugo P. Simao).

12.2. Consultative and advisory functions

I have given several presentations on this research to the analysis group at AMC, as well as at AFIT. This past summer, I gave two talks at AMC:

“Real-Time Optimization for Real-World Problems,” Air Mobility Command, Scott Air Force Base, June, 2004.

“The Optimizing Simulator: Modeling the Organization and Flow of Information and Decisions,” Air Mobility Command, Scott Air Force Base, June, 2004.
and a separate talk at AFIT:

“The Optimizing Simulator for Military Airlift Problems,” Air Force Institute of Technology, July, 2004.

12.3. Transitions

Our transitions have occurred along three lines:

- Communication of ideas to the analysis group at AMC and their subcontractors.
- Sharing our airlift mobility simulator with graduate students doing research on air mobility and related topics (Ken Kopp, McGuire AFB, and Todd Sriver, AFIT).
- Direct implementation of ideas through projects with the corporate partners of CASTLE Lab.
- Making software available for download from the CASTLE Lab web site (the PILOTVIEW diagnostic library is now available).
- Licensing of software through local consulting firms for use in systems for their clients. CASTLE Lab has a relationship with Princeton Consultants, Inc. (www.princeton.com) which implements optimization and simulation models in transportation and logistics.

Specific transitions to the industrial partners of CASTLE Lab over the last three years include:

1. Transition: Optimizing simulator for fleet planning at Schneider National. We have calibrated a system that models the flows of approximately 5,000 drivers of different types. Schneider is interested in knowing what types of drivers are most valuable to the fleet (similar to AMC asking which aircraft types are most valuable). It is almost impossible to answer this question using “what if” analyses. Our logic produces, from one run, estimates of the gradients with respect to each type of driver.

Recipient: Schneider National, the nation's largest truckload motor carrier.

2. Transition: A model for short-term operational forecasting of freight cars using the optimizing simulator. The optimizing simulator framework can handle multiple sources of uncertainty (customer demands, order locations, travel times, equipment failures).

Recipient: Norfolk Southern Railroad

- 3 Transition: Operational, tactical and strategic planning of locomotives. This system uses the optimizing simulator concept, and in particular makes heavy use of techniques for modeling incomplete information through low dimensional patterns. The system was recently approved for production at Norfolk Southern Railroad, making it the first successful production optimization model developed for operational use in North America.

Recipients:

Norfolk Southern Railroad, which uses the system both for strategic planning of the fleet size, and short-term tactical forecasting of surpluses and deficits.

Burlington Northern Sante Fe Railroad, where the primary focus is on routing locomotives to maintenance facilities, where they have to arrive at a particular time.

4. Transition: We have been working for years to developing an operational routing and scheduling system for the multilayered resource scheduling problem. This has been applied to the routing and scheduling of drivers, tractors, trailers, product and customer tanks. This system is now in production and is available for schedulers to use on demand.

Recipient: Air Products and Chemicals

5. Transition: Forecasting of time series activities with multiple calendar effects using hierarchical aggregation.

Recipient: Yellow Freight System, Norfolk Southern Railroad

6. Transition: A driver scheduling system for large-scale less-than-truckload motor carriers using informational decomposition. We combined the adaptive learning techniques for coordinating multiple agents along with low dimensional patterns for capturing incomplete information. The system is now being used to perform short term tactical forecasts of movements of drivers.

Recipient: Yellow/Roadway Transportation

Software was also licensed to Transport Dynamics and Princeton Consultants (which purchased Transport Dynamics in 2003). Through these companies, our research is now running at FedEx Ground, FedEx Custom Critical, Watkins Motor Lines, and de Boer Truckline. These are now successful commercial product lines.