A Primal-Dual Iterative Scheme for Solving Capacity Planning Problems under Uncertainty

Abhilasha Aswal[a] and G. N. Srinivasa Prasanna[b]

[a]*Infosys Technologies Limited*
44 Electronics City,
Hosur Road, Bangalore – 560100, India.
Email: abhilasha_aswal@infosys.com

[b]*International Institute of Information Technology*
26/C Electronics City,
Hosur Road, Bangalore – 560100, India.
Email: gnsprasanna@iiitb.ac.in

Abstract. In this paper we present an application of robust optimization to capacity planning problems under uncertainty. We present the framework to handle uncertainty and discuss the computational complexity of capacity planning problems under this framework. We show that the formulation is not only intuitive but the computational complexity of a large variety of problems is the same as linear (in general convex) programming.

Keywords: Capacity Planning, Optimization under uncertainty, Robust programming
PACS: 45.70.Vn; 91.65.My

INTRODUCTION

A supply chain is a network of suppliers, production facilities, warehouses and end markets. Capacity planning decisions involve decisions concerning the design and the configuration of this network. The decisions are made on two levels: strategic and tactical. Strategic decisions include decisions such as where and how many facilities should be built and what their capacity should be. Tactical decisions include where to procure the raw-materials from and in what quantity and how to distribute finished products. These decisions are long range decisions and a static model for the supply chain that takes into account aggregated demands, supplies, capacities and costs over a long period of time (such as a year) will work. The challenge is to make these decisions under uncertainty.

To deal with uncertainty, extensive research has been carried out in both Probabilistic (Stochastic) Optimization and Robust Optimization (constraints) frameworks. However, these techniques face difficulties in conveniently estimating the data that they require. For new products, such data may not even exist. We have proposed an extension of robust optimization to solve this problem intuitively and meaningfully in our earlier work [5], [6], [7], [8].

Below, in Section II we give an overview of the problem. Section III gives a mathematical formulation of the problem, and characterizes the properties of the solution under linear metrics. Section IV gives a simple example of our ideas. Section V gives results on a wide variety of examples, including those with nonlinear costs with breakpoints. Section VI concludes.

OVERVIEW

Our framework is characterized by intuitive specification of uncertainty and its quantification. Classical robust optimization approaches are seen to be too conservative. To make the solutions more attractive, Bertsimas, Sim, Theile [2], [3], [4] have proposed an approach where the level of conservatism or the budget of uncertainty, for each constraint can be controlled based on probabilistic bounds of constraint violation. An advantage of their approach is that the robust counterpart of a linear programming problem is also a linear, thus maintaining computational tractability. But their model specifies uncertainty as a symmetric variation around a nominal point, which may not reflect economically meaningful information.

Strengths of our Formulation

In contrast to this, we represent uncertainty in a constraint based framework naturally derived from basic economic principles. Instead of specifying the data directly, we specify bounds on linear (& quadratic) combinations of the data, incorporating correlations amongst the data elements. The uncertainty sets (constraint sets) form a convex polytope, built from simple and intuitive linear constraints (simple sums and differences of supplies, demands etc) those are derivable from historical time series data, which are meaningful in terms of macro-economic behavior. Specifically, substitutive effects bound the sum of different demands, complementary effects bound differences, revenue constraints bound weighted sums, etc. The budget of uncertainty does not adequately reflect these underlying physical realities.

With our specification, many kinds of future uncertainty can be specified. Not only does this specification avoid ad-hoc gravity models and their variants, as well as ad-hoc probability distributions, but it is also simple and intuitive. Answers are globally valid over the entire range of parameter variation.

In addition, we have a unique ability to quantify information content in the polytope using Shannon's information theoretic concepts, based on which we can quantitatively compare different scenario sets for the future [7]. This is done as follows.

Assuming that in the lack of information, the parameters vary with equal probability in a large region R (taken to be of finite volume for simplicity initially), of volume Vmax. Then the constraints specifying the convex polyhedron CP specify a subset of the region R, of volume V_{CP}. The amount of information provided by the constraints specifying the convex polyhedron can be equated to:

$$I = \log 2\left(\frac{V_{\max}}{V_{Cp}}\right) \tag{1}$$

It can be easily seen that (1) reduces to the Shannon surprisal [9] of getting a set of parameters satisfying the constraints, assuming that the parameters are equiprobable in the large region R. While the estimation of polyhedral volume is a difficult problem and initial results with polyhedra obtained by practical constraints are promising.

The identification of polyhedral volume with information content yields several powerful methods of data analysis. For example, we can construct new constraint sets equivalent in information content to an original set, by using orthogonal transformations, on the original set.

The formulation of uncertainty is clearly very powerful. Simultaneously, however, it does not substantially increase the computational complexity of the optimization problems, and we shall show a few important cases in this paper.

In general, with linear constraints, it is easy to model most optimization problems as linear programs. However, in practice a number of non-convex constraints like cost/price breakpoints and binary 0/1 facility location decisions change the problem from a standard LP to an non-convex ILP problem, and heuristics are necessary for obtaining the solution even with state-of-the-art programs like CPLEX. While such optimizations are very difficult, it is possible to bound the performance of the optimal solution reasonably simply. We show below that Linear Programming allows us to determine bounds on performance of any metric given one or more solutions. These solutions may have been obtained by ad-hoc or other heuristics. An ensemble of such solutions enables us to find tight bounds for the metrics encountered in various classes of optimization problems. Many classical problems can be generalized and solved using such a representation of uncertainty.

Although getting the optimal answers is difficult but note that even bounds for the optimal solution are very useful information in the complex supply chain framework. Given this, the advantages of our approach are that bounds can be quickly given on any candidate solution using LP/ILP, since the equations are then linear/quasi-linear in the demands/supplies/other parameters, which are linearly constrained (or using Quadratic programming with quadratic constraints). The best case, best decision and worst case, worst decision are clearly global bounds, solved directly by LP/ILP. Details are skipped for brevity (see [9]).

THE CAPACITY PLANNING PROBLEM

From a theoretical viewpoint, the classical multi-commodity flow model [1] is the natural formulation for capacity planning. In the static multi-commodity flow model (at a single time instant) for each commodity, inflows Φ_{ij} are equal to outflows Φ_{ik} at a node j. In addition, the flows Φ_{ij} are bounded. If we need to minimize cost we get the optimization (written for a single commodity):

$$\begin{aligned}
&\min \sum_{ij} c_{ij}\phi_{ij} \\
&\sum_{i \in \mathrm{Pred}\ j} \phi_{ij} - \sum_{k \in Succ\ j} \phi_{jk} = d_j \\
&0 \le \phi_{ij} \le C_{\max}
\end{aligned} \tag{2}$$

Here i varies over all predecessors of j, and k varies over all successors of j, and d_j is the amount of sourced/sinked flow at node j. The optimization task is to decide the flows in order to minimize the cost. A variety of optimizations result, based on auxiliary assumptions, and these are detailed below. We show that while these optimizations are NP-hard/non-convex, heuristic techniques exist which give solutions within 10-20% of optimal.

A general supply chain is given in FIGURE 1, connecting suppliers to factories and markets. Capacity planning problem on such a chain would mean finding the optimum set of factory locations (suppliers and demands are taken as given), capacities and distribution policy for a given demand, including in general some uncertainty.

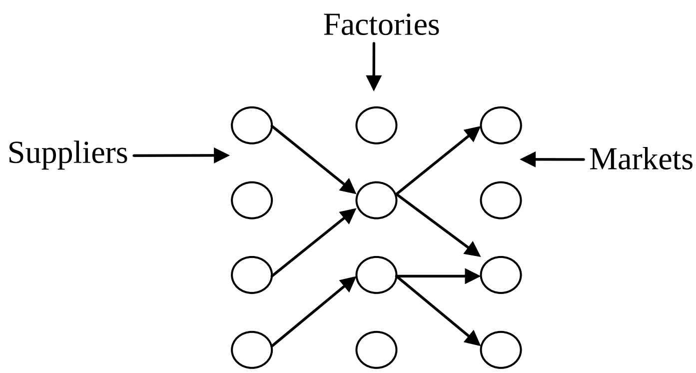

FIGURE 1. A general supply chain

This problem can be formulated in various flavors, with a variety of assumptions on demands being fixed or variable, locations of factories being fixed / variable, costs being linear/piecewise linear, variables having integrality constraints etc. The computational complexity of these problems varies – we illustrate this by looking at problems at two extremes of the spectrum (next section has more details).

At one extreme is a problem with fixed demands, locations, and linear costs. This is directly solvable using linear programming, as shown below. At the other extreme are problems with variable demands, locations, and nonlinear costs. This problem is computationally difficult, but we present a heuristic which comes to within 10-20% of optimal below.

Fixed Demands and Fixed Locations with Linear Costs

If the demands and the locations are known exactly, then a variety of optimizations can be formulated, and we mention two examples below. The problem of finding the minimum cost flow can be formulated as the following linear program.

$$\begin{aligned} \text{Minimize} \quad & C^T \Phi \\ \text{Subject to} \quad & A\Phi = d \\ & \Phi \geq 0 \end{aligned} \tag{3}$$

Here, C is the vector of linear costs, Φ is the vector of flows, d is the demand vector and A is the incidence matrix defining the supply chain network. The solution to this LP is easily obtained. The maximum cost flow problem can be similarly solved.

In order to find the maximum flow on any link satisfying the demand, the problem can be specified as follows.

$$\begin{aligned} &\text{Minimize} \quad z \\ &\text{Subject to} \quad z \geq C_{ij}\Phi_{ij} \; \forall i, j \\ &\qquad A\Phi = d \\ &\qquad \Phi \geq 0 \end{aligned} \tag{4}$$

Variable Polyhedral Demand and Fixed Locations with Linear Costs

If the locations are fixed but the demands are variable but constrained within a convex polyhedral set ($(CP)d \leq e$), the problem becomes more complex. The uncertainty manifested in the demand is classically tackled using *recourse*- adjustment of the solution *after* the demand has materialized. We note that another, static solution where all decisions are made *before* the demand has materialized can also be considered, but will be skipped for brevity.

The optimization using recourse, takes an optimal decision, after the demand has materialized. To evaluate its performance, we have to determine the demand for which the optimal routing is most expensive. This is formulated as follows:

$$\begin{aligned} &Max_{d_{ij}} \left(Min_{\Phi_{ij}} C^T \Phi \right) \\ &S.t. \\ &A\Phi = d \\ &\Phi \geq 0 \\ &(CP)d \leq e \end{aligned} \tag{5}$$

The maximization is over the demand variables dij’s, and the flow Φ for each set of demands is optimally chosen. As written, this problem is not an LP, since it is a max-min. However, utilizing the theory of duality, we can transform this to a tractable optimization.

Assuming strong duality with respect to the flow variables Φ, with dual variables ν, the above problem can be formulated as follows

$$Max_{d} \begin{pmatrix} \text{Maximize}_{\nu} & -\mathrm{d}^T \nu \\ \text{Subject to} & A^T \nu \geq -c \\ & (CP)d \leq e \\ & \nu \geq 0 \end{pmatrix} \tag{6}$$

$$\Rightarrow \quad Max_{d} \begin{pmatrix} \text{Maximize}_{\ v} & -\mathrm{d}^{T} v \\ \text{Subject to} & \begin{pmatrix} -A^{T} & 0 \\ 0 & CP \end{pmatrix} \begin{pmatrix} v \\ d \end{pmatrix} \leq \begin{pmatrix} c \\ e \end{pmatrix} \\ & d, v \geq 0 \end{pmatrix}$$

$$\Rightarrow \quad \begin{array}{ll} \text{Maximize}_{\ d,v} & -d^{T} v \\ \text{Subject to} & \begin{pmatrix} -A^{T} & 0 \\ 0 & CP \end{pmatrix} \begin{pmatrix} v \\ d \end{pmatrix} \leq \begin{pmatrix} c \\ e \end{pmatrix} \\ & d, v \geq 0 \end{array} \tag{7}$$

This is a quadratic program and generally difficult to solve optimally, with the solution being possibly at an interior point. Also, the metric is not positive semi-definite (it has a saddle point at $(d, v) = (0, 0)$). However, we can considerably simplify the solution using the following lemma.

Lemma: The optimal value of d and v are at the vertices of the polytopes defined by $c + A^{T} v \geq 0$ *and* $(CP)d \leq e$.

Proof: Note that

- For a fixed *d* the metric $d^{T} v$ is linear in *v* and vice versa.
- The constraints on *d* and *v* are decoupled.

The metric $d^{T} v$ is linear in *d* for a fixed *v* and vice versa. If the optimal point is denoted as (d^{*}, v^{*}) d^{*} has to be at a vertex of the d polytope ($(CP)d \leq e$). Similarly the v^{*} is at a vertex of the polytope $c + A^{T} v \geq 0$. Hence the optimal point (d^{*}, v^{*}) is found amongst the vertices of the polyhedron specified by ($(CP)d \leq e, C + A^{T} v \geq 0$). Since, the number of vertices is exponential in the number of constraints and variables, an exhaustive search is computationally infeasible. At this point of time we do not have a fast polynomial method for this optimization, and we propose a heuristic based on Lemma 1

Heuristic-DV: Heuristic-DV alternates between the *d* and *v* spaces. An initial candidate solution satisfying the uncertainty constraints (CP)d <= e, is found in the *d*-space, and the vertex in the *v*-space which optimizes the metric is determined using an LP. Then this vertex in the *v*-space is kept fixed, and the best vertex in the d-space which optimizes the metric is determined using an LP. The process is continued till convergence is attained. Unfortunately, this search can get stuck in local optima, as the following example with 2 demands shows. Consider the vertices in the $d = [d_1, d_2]$ space as [1, 1] and [2, 1/2]; and in the $v = [v_1, v_2]$ space as [1, 1] and [0, 4.5], then looking at FIGURE 4 it can easily be verified that if we start at $d = [2, 0.5]$, then we get stuck at the local optimum -2.25 when clearly the global optimum is -2.

Heuristic-DV is then enhanced using simulated annealing, where we perturb the solution obtained, and repeat the process. FIGURE 3 demonstrates the heuristic.

$d \downarrow$ $\vec{v}$	$(1,1)$	$(0,4.5)$
$(1,1)$	-2	-4.5
$(2,0.5)$	-2.5	-2.25

FIGURE 2: Heuristic-DV example

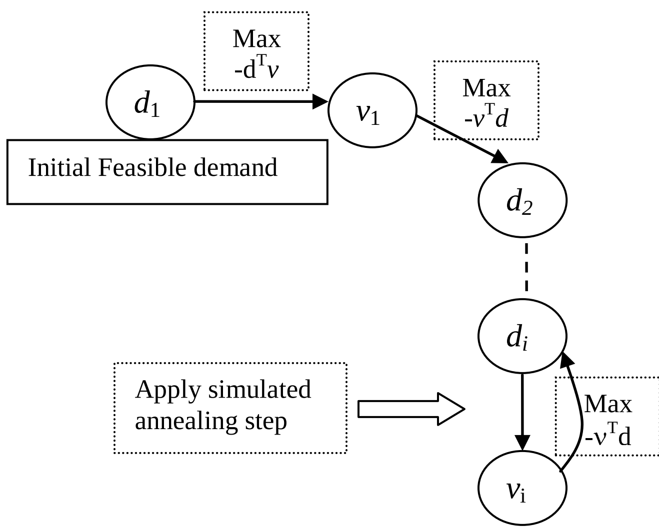

FIGURE 3: Heuristic-DV

Variable Polyhedral Demand and Variable Locations with Linear Costs

In this case, the problem becomes an ILP as the locations are now variable so we need a location variable x_i for each potential location i ($x_i \in \{0,1\}$).

$$
\begin{aligned}
\text{Minimize}_{\Phi_{ij},\, x_i} \quad & \left(C^T \Phi \right) \\
\text{subject to} \quad & A\Phi = d \\
& \Phi \geq 0 \\
& (CP)\mathrm{d} \leq \mathrm{e} \\
& x \in \{0,1\}
\end{aligned}
\tag{8}
$$

The problem with recourse becomes:

$$
\begin{aligned}
& Max_{d_{ij}} \left(Min_{\Phi_{ij}, x_i} C^T \Phi \right) \\
& \text{subject to} \quad A\Phi = d \\
& \qquad \Phi \geq 0 \\
& \qquad (CP)d \leq e \\
& \qquad x_i \in \{0,1\}
\end{aligned}
\tag{9}
$$

Taking the dual, the resultant problem becomes a MIQP problem (details omitted for brevity). The problem is NP-hard in this case and cutting plane, branch and bound and similar techniques are required.

Variable Polyhedral Demand and Variable Locations with Breakpoints and Multiple Fixed and Variable Costs

In practice, there are seldom any costs that are without breakpoints. The problem is highly non-convex in this case, and heuristics are in general required. However, this computational difficulty is *intrinsic* to the cost function, and not a result of our uncertainty formulation.

SIMPLE EXAMPLE

The following simple illustrative example demonstrates all the aspects of the proposed formulation. It shows how uncertainty can be specified and how capacity bounds are derived for varying assumptions.

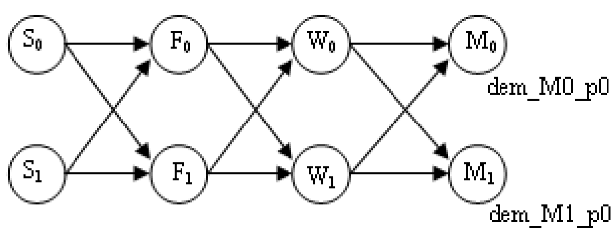

FIGURE 4: Simple supply chain

The single-commodity supply chain in FIGURE 4 consists of 2 suppliers, 2 plants, 2 warehouses and 2 market locations. We want to minimize/maximize the total cost of the supply chain while satisfying the demand for the product at the markets.

To evaluate the performance of our solution, we calculate the absolute bounds on cost – the minimum cost under the best decision for the best demand (min-min), and the maximum cost (max-max) assuming the worst possible decision for the worst case demand. Both are directly obtained using LPs, and serve as bounds for the performance of heuristic DV. If the costs are linear and the locations are fixed, we can use heuristic-DV to find the solution.

In this example, the direct edges are low cost, 10 units/unit of flow, and the cross edges are 15 units/unit of flow. An example of a constraint set ((CP)d $\leq$ e) for the uncertain demand derived from historical data is below:

```
dem_M0_p0 + dem_M1_p0 >= 200
dem_M0_p0 + dem_M1_p0 <= 400
dem_M0_p0 - dem_M1_p0 <= 20
dem_M0_p0 - dem_M1_p0 >= -20
```

These have an economic interpretation - the first two bound the sum of the first and second demand, a substitutive effect. The next two put bounds on differences, a complimentarity effect.

Assuming linear costs and fixed locations, the min-min solution to the problem is given by FIGURE 5 and the max-max solution is given in FIGURE 6, where dark edges carry flow and light edges do not carry any flow. The min-min solution has low demand (as low as permissible under the constraints). It carries flow only on the least cost edges (each costing 10 units per unit of flow), and the max-max pushes all the flow to the cross edges, each costing 15 units per unit of flow

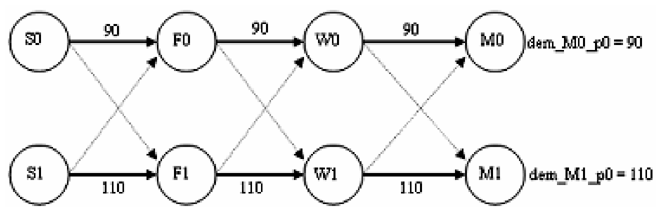

FIGURE 5: Optimal routing for the min demand with fixed locations and linear costs (lower bound on cost)

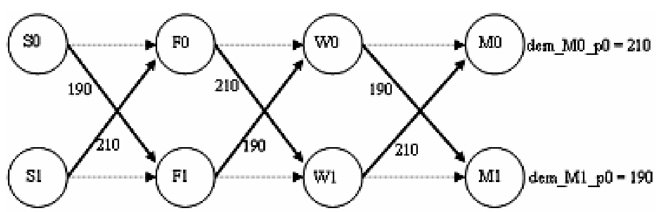

FIGURE 6: Worst routing for the max demand with fixed locations and linear costs (upper bound on cost)

The min-min cost for operation was 6400.00 units. The max-max cost for this problem was 18000.00 units. We note that these are absolute bounds – with maximum optimism and pessimism about the demand and decision both. In practice, these will not be used for operational decision, but furnish performance limits.

Given these bounds, we are interested in finding the cost of optimal routing with worst case demands (a max-min) and for that we can use Heuristic-DV.

We first take an initial random vertex in the demand space (satisfying constraints 1.1) as follows:

$$\vec{D} = [d_1 \quad d_2] = [110 \quad 90]$$

We maximize with respect to the dual variables (flow variables) and get an objective function value of 7350. We now find the values of dual variables $\vec{\nu}$ and fixing those, we maximize with respect to the demand variables and get an objective function value of 9450. Then we repeat the process. Table 1 summarizes the result of this exercise.

The cost for the random demand vertex was 7350. In the next step, maximizing with respect to demand variables keeping dual variables fixed, the cost was 9450. The new demand vertex is now [210 190]. We maximize with respect to dual variables and get a solution of 14850, fix the dual variables and maximize with respect to the demand variables. At this point we come back to the previous demand vertex [210 190]. It could be a local optimum and so we perturb the demand vector and choose another random point [150 130]. The steps of the heuristic are summarized in Table 2. The heuristic again converges to the same solution.

TABLE 1. Heuristic-*DV* For Small Example With Initial Random Demand Vector = [110 90]

Demand vertex	Objective function value with D fixed	Objective function value with V fixed
[110 90]	6400	12800
[210 190]	12800	12800
[210 190]	12800	12800

TABLE 2. Heuristic- *DV* For Small Example With Initial Random Demand Vector = [150 130]

Demand vertex	Objective function value with D fixed	Objective function value with V fixed
[150 130]	8960	12800
[210 190]	12800	12800
[210 190]	12800	12800

The maximum value for the objective function that we obtained was 12800.00. The min-min solution for this example was 6400.00 and the absolute maximum was 18000.00. Thus the answer obtained by Heuristic-DV is within 29% of the upper bound on the cost, and reflects the improvement obtained by making optimal decisions, once the demand has materialized.

RESULTS

In this section we present some results from our simulations. All of our results were produced on an Intel Celeron 1.60 GHz processor, with a 512 MB RAM.

First, we illustrate Heuristic-DV with a larger example. We consider a supply chain with 5 suppliers, 5 factories, 5 warehouses and 5 markets. There is only 1 product demand for which the supply chain services hence there are 5 demand variables (one for each market). The demands are constrained as follows:

```
dem_1 + dem_2 + dem_3 + dem_4 + dem_5 >= 200
dem_1 + dem_2 + dem_3 + dem_4 + dem_5 <= 700
dem_1 - dem_2 <= 100
dem_1 - dem_2 >= -100
dem_4 - dem_3 <= 80
```

```
dem_4 - dem_3 >= 20
dem_5 - dem_3 - dem_1 >= 10
dem_5 - dem_3 - dem_1 <= 130
```

Given per unit transportation costs through all the links in the supply chain, we want to find the demand for which the optimal routing is most expensive. This is the robust optimal solution. We formulate the problem as given by (5) and taking the dual, we get (7). We now apply Heuristic-DV. We first take a random vertex in the demand space as follows

$$\vec{D} = \begin{bmatrix} d_1 & d_2 & d_3 & d_4 & d_5 \end{bmatrix} = \begin{bmatrix} 0 & 100 & 3.34 & 83.34 & 13.34 \end{bmatrix}$$

We maximize with respect to the dual variables and get an objective function value of 7000.00. We now find the values of dual variables $\vec{V}$ and fixing those, we maximize with respect to the demand variables and repeat the process. Table 3 summarizes the result of this exercise.

TABLE 3. Heuristic-*DV* Example 2

Demand vertex	Objective function value with D fixed	Objective function value with V fixed
[0 100 3.34 83.34 13.34]	7000	24500
[216.67 116.67 0 20 346.67]	24500	24500
[216.67 116.67 0 20 346.67]	24500	24500

The cost for the random demand vertex was 7000.00. In the next step, maximizing with respect to demand variables, the cost was 24500.00. From here we came to the demand vertex [216.67 116.67 0 20 346.67]. We fix the demand vertex, maximize with respect to the dual variables, fix the dual variables and get the same demand vertex back. If the heuristic proceeds further, it will keep cycling through this demand vertex. The maximum value for the objective function that we obtained was 24500.00. The absolute minimum cost for this example was 6000.00 and the absolute maximum was 31500.00. Thus the answer obtained by Heuristic-DV tightens the upper bound on cost by 22%. This could be a local optimum and we can apply a simulated annealing step here to come out of it.

In another example, supply chain given in FIGURE 7 is considered. In this example there are 10 demand variables and the constraints on demand are given as follows:

```
dem_1 + dem_2 + dem_3 + dem_4 + dem_5 + dem_6 + dem_7 + dem_8 + dem_9
+ dem_10 >= 700
dem_1 + dem_2 + dem_3 + dem_4 + dem_5 + dem_6 + dem_7 + dem_8 + dem_9
+ dem_10 <= 2000
dem_1 + dem_2 + dem_3 - dem_4 - dem_5 - dem_6 >= 100
dem_5 + dem_6 + dem_7 - dem_8 - dem_9 - dem_10 >= 150
dem_3 + dem_4 + dem_5 >= 300
dem_9 + dem_10 <= 100
dem_9 - dem_10 >= 20
dem_2 - dem_3 >= 35
dem_8 - dem_9 >= 20
dem_1 - dem_4 >= 45
dem_4 - dem_6 >= 35
dem_6 + dem_7 >= 150
dem_6 + dem_7 <= 250
```

```
dem_6 - dem_10 >= 40
dem_10 >= 30
```

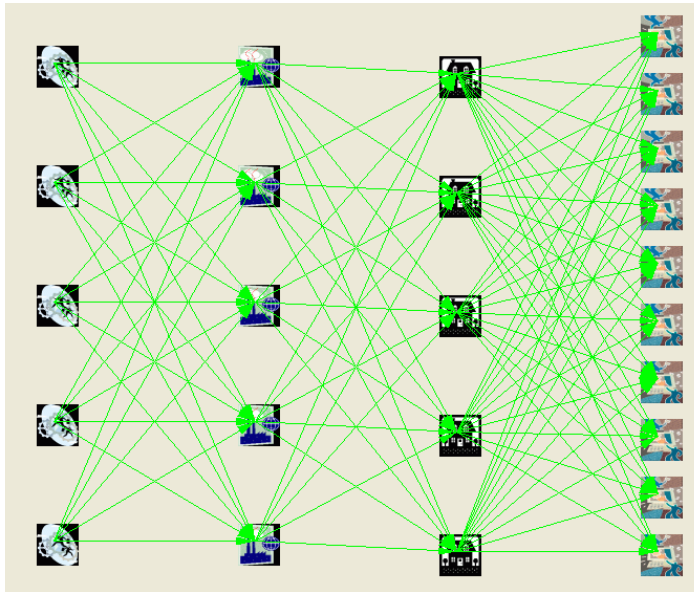

FIGURE 7: A medium supply chain

A random demand vertex is chosen as follows for the first iteration:

$$\vec{D} = [d_1 \quad d_2 \quad d_3 \quad d_4 \quad d_5 \quad d_6 \quad d_7 \quad d_8 \quad d_9 \quad d_{10}] = [200 \quad 350 \quad 100 \quad 105 \quad 195 \quad 70 \quad 180 \quad 70 \quad 50 \quad 30]$$

Table 4 summarizes the results of further iterative procedure.

TABLE 4. Heuristic-*DV* Example 3

Demand vertex	Objective function value with D fixed	Objective function value with V fixed
[200, 350, 100, 105, 195, 70, 180, 70, 50, 30]	42500	63866.67
[663.33, 35, 0, 105, 423.33, 70, 180, 443.33, 50, 30]	63866.67	63866.67
[663.33, 35, 0, 105, 423.33, 70, 180, 443.33, 50, 30]	63866.67	63866.67

The maximum value for the objective function that we obtained was 63866.67. The absolute minimum cost for this example was 2600.00 and the absolute maximum was 90000.00. Thus the answer obtained by Heuristic-DV tightens the upper bound on cost by 29%.

TABLE 5 compares the time taken to find the absolute bounds on cost to the time taken to run Heuristic-DV for the present example as well as for the example in the previous section.

These are the preliminary results from our investigation. We will apply this technique to solve large scale supply chain optimizations with multiple products in the future.

TABLE 5. Time taken to solve heuristic-dv examples of sections IV and V

Nodes	Problem	Variables	Time taken in seconds	Time for LP relaxation in seconds
8	Min-Min	16	0.14	0.06
8	Max-Max	16	0.16	0.05
8	Heuristic-DV	8	0.4	0.25
20	Min-Min	80	0.13	0.05
20	Min-Max	80	0.17	0.05
20	Heuristic-DV	20	1.44	0.81
25	Min-Min	100	0.14	0.06
25	Min-Max	100	0.06	0.03
25	Heuristic-DV	30	0.33	0.15

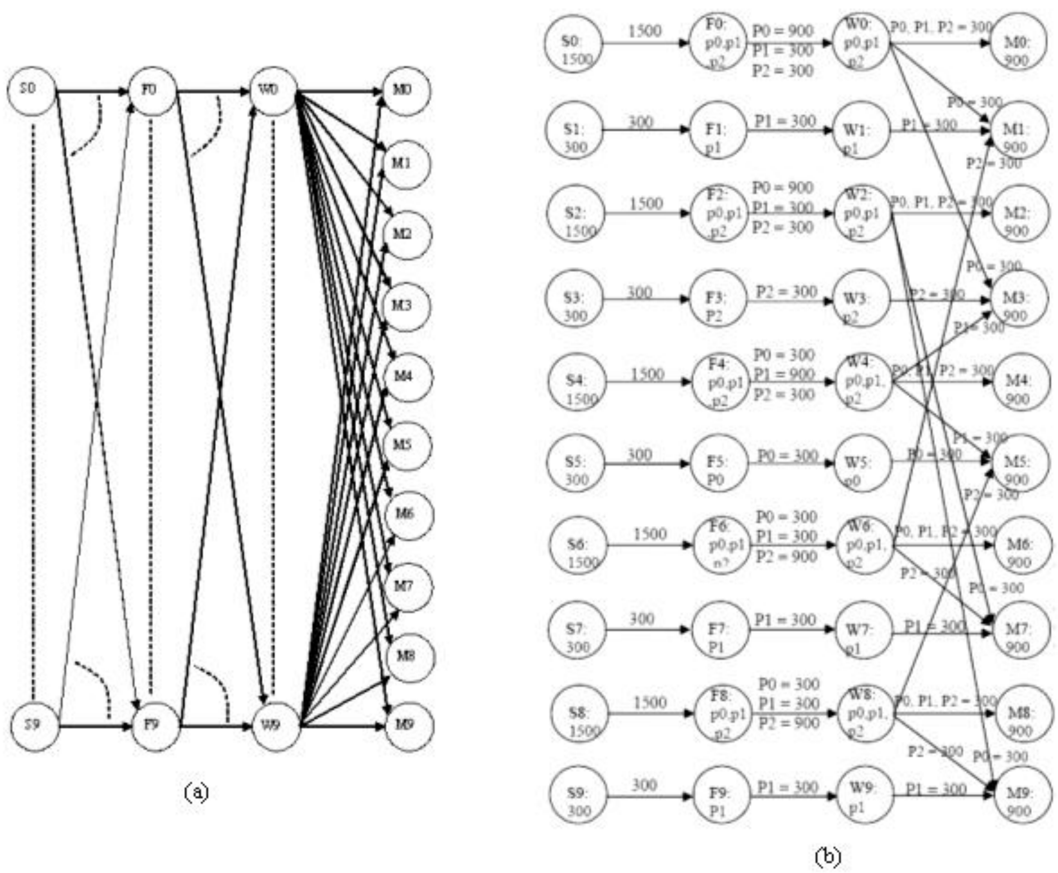

FIGURE 8: (a) A medium sized supply chain example
(b) Solution for the min-min problem

We have solved larger problems of different classes. The example given in FIGURE 8 (a) shows a medium sized supply chain with 40 nodes and 3 products. There is one breakpoint in the costs, and the locations are also variable. If all demands range between a minimum of 100 units and a maximum of 5000 units, then the optimal routing for the minimum demand as found using ILOG CPLEX solver is shown in FIGURE 8 (b).

CONCLUSION

We have shown a technique for solving capacity planning problems under our intuitive uncertainty formulation, which is tractable. Examining a wide variety of capacity planning problems, we focused the problem with recourse in detail, showing that it is a quadratic program, under the assumptions of linear costs and fixed locations. We presented a primal-dual heuristic - *DV* which comes to within 10-20% of optimal. Realistic costs with breakpoints lead to integrality constraints, which increase the computational difficulty, but initial results from our heuristics are promising. This approach presents a promising generalization of the range constraints of Bertsimas, Sim, Theile [2], [3], [4].

REFERENCES

1. Ahuja, Magnanti, Orlin, “Network Flows, Theory, Algorithms and Applications”, Prentice Hall, 1993
2. Bertsimas, D., Sim, M., “The price of robustness”, Operations Research, 52, 1, 2004, pp. 35-53
3. Bertsimas, D., Thiele, A., “A robust optimization approach to supply chain management”, Operations Research, 54, 1, 2006, pp. 150-168
4. Bertsimas, D., Thiele, A., “Robust and Data-Driven Optimization: Modern Decision-Making Under Uncertainty”, 2006
5. Prasanna, G. N. S., “Traffic Constraints instead of Traffic Matrices: A New Approach to Traffic Characterization”, Proceedings ITC, 2003.
6. Prasanna, G. N. S., Aswal, A., Chandrababu, A., Paturu, D., “Capacity Planning Under Uncertainty: A Merger of Robust Optimization and Information Theory applied to Supply Chain Management”, Proceedings of the Annual Convention of the Operational Research Society of India, December 2007, New Delhi
7. Aswal, A., Prasanna G. N. S., “Coupling information theory with optimization: A new approach to robust design of supply chains”, Proceedings of the 23rd Belgian Conference on Operations Research, February 2009, Belgium
8. Aswal, A., Prasanna, G. N. S., “New vistas in inventory optimization under uncertainty”, Proceedings of the International MultiConference of Engineers and Computer Scientists 2009, IMECS 2009, March 18 - 20, 2009, Hong Kong
9. Shannon, C. E., “A Mathematical Theory of Communication”, Bell System Technical Journal, 27, pp. 379–423 & 623–656, July & October, 1948.
10. Aswal, A., “Capacity Planning and Inventory Optimization under Uncertainty”, Master’s thesis, International Institute of Information Technology - Bangalore, June 2008
11. Aswal, A., Prasanna, G. N. S., “Computational Complexity of Capacity Planning Problems under Uncertainty”, Lecture Notes in Engineering and Computer Science: Proceedings of The International MultiConference of Engineers and Computer Scientists 2010, IMECS 2010, 17-19 March, 2010, Hong Kong