Strategic Design of Robust Global Supply Chains under Uncertainty

Marc Goetschalckx, Shabbir Ahmed, Alexander Shapiro, Tjendera Santoso, and Gonzalo Cordova

School of Industrial and Systems Engineering,
Georgia Institute of Technology

Abstract

We will present a model for the strategic design of the global supply chain of an individual company. The model presented is a multi-period, two-stage stochastic, multi-country, multi-product, multi-echelon formulation based on forecasted exchange rates and with bill of materials (BOM) flow conservation. The objective is maximization of the expected value of the time-discounted world-wide after-tax net cash flows *NPV(NCF)*. The uncertainty of the data is explicitly considered through the inclusion of scenarios in the deterministic equivalent problem (DEP). However, the problem size of the DEP is so large that convergence of standard Benders decomposition requires an excessive amount of computer time. We have implemented several acceleration techniques for Benders decomposition. The following observations are based on the statistical analysis of the required computing times for an industrial case study. 1) Acceleration techniques significantly decrease the computation times. 2) Strengthening of the dual variables contributes the most, while a primal heuristic contributes the least. 3) The largest reductions in computer time are achieved when all acceleration techniques are used in concert. We are continuing the numerical experiment to further differentiate the effectiveness of the acceleration techniques.

1 Introduction

In today's global economy corporations constantly have to evaluate and redesign their supply chains to respond to the rapidly changing conditions. Capital investment decisions have to be made long before all the relevant data are known with certainty. Evaluation of the capital investments and supply chain configurations has to be consistent with generally accepted accounting principles and financial reporting, incorporate global trade and taxation laws and regulations, and be compatible with the risk preferences of the corporation.

We will present a model for the strategic design of the global supply chain of an individual company. The model presented is a multi-period, two-stage stochastic, multi-country, multi-product, multi-echelon formulation based on forecasted exchange rates and with bill of materials (BOM) flow conservation. The objective is the maximization of the expected value of the time-discounted world-wide after-tax net cash flows, or in other words, the net present value of the net cash flows *NPV(NCF)*.

The uncertainty of the data is incorporated explicitly in the model through the inclusion of scenarios. The accuracy of the expected value of the objective function, which in our case is the *NPV(NCF)*, in two-stage stochastic programming problems increases with the number of scenarios included in the deterministic equivalent formulation or DEP. The DEP problem is a very large mixed integer programming problem, where the flow variables have five indices: scenario, time period, product, origin, and destination. The DEP is solved for in the framework of the sample average approximation (SAA) method, which itself requires tens of master iterations and each master iterations requires hundreds of solutions of the second-stage recourse problem, which is a capacitated multi-period multi-commodity network flow problem with BOM constraints.

2 Model and Solution Algorithms for Strategic Supply Chain Design under Uncertainty

2.1 Scenario-Based Model

The accurate assessment of the risk associated with a strategic supply chain configuration requires the use of a large number (hundreds or thousands) of scenarios. At the same time, increasing the number of scenarios in the DEP yields unacceptable computation times. We will describe an algorithm that integrates a recently proposed sampling strategy, the Sample Average Approximation (SAA) scheme, with an accelerated Benders decomposition algorithm to solve the strategic supply chain design problem with continuous distributions for the uncertain parameters, and hence an infinite number of scenarios.

We will approximate the two-stage stochastic optimization formulation for the strategic design of supply chains (1.1) with its deterministic equivalent problem (1.3) based on randomly sampled scenarios.

$$\begin{aligned} Max \quad & cy + \mathrm{E}\left[Q(y,\xi)\right] \\ s.t. \quad & Hy \le g \\ & y \in \{0,1\} \end{aligned} \tag{1.1}$$

Where the second stage or inner optimization problem is defined as

$$\begin{aligned} Q(y,\xi) = \quad Max \quad & d(\xi)\,x \\ s.t. \quad & F(\xi)\,x \le h(\xi) + E(\xi)\,y \\ & x \ge 0 \end{aligned} \tag{1.2}$$

With:

y	binary facility, size, and technology status variables
x	continuous material flow and storage variables
c	fixed facility cost vector (costs are negative)
$d(\xi)$	operational sales and cost vector consisting of variable (positive) sales revenues and (negative) transportation, purchasing, production, and inventory costs
$h(\xi)$	right-hand side of the technological constraints
$E(\xi), F(\xi)$	technology and conservation of product flow and storage matrix
H	relationships between facility status variables
g	right-hand side of the facility relationship constraints
$\mathrm{E}[v]$	expected value of random variable v
ξ	random vector

The above stochastic optimization problem can in theory be transformed into an equivalent deterministic optimization problem (DEP) by including all possible scenarios weighted by their associated probabilities.

$$\begin{aligned} Max \quad & cy + \sum_{s=1}^{N} p_s d_s x_s \\ s.t. \quad & Hy \le g \\ & y \in \{0,1\} \\ & -E_s y + F_s x_s \le h_s \\ & x_s \ge 0 \end{aligned} \tag{1.3}$$

With

S	scenario index, $s = 1, 2 \ldots N$
p_s	probability of scenarios s, in the SAA method $p_s = 1/N$

If the probability distribution functions of the parameters are discrete, the resulting number of scenarios is finite but extremely large. For example, a problem with 1000 parameters, each having a probability distribution with five discrete values, yields 5^{1000} scenarios. If the probability functions of the parameters are continuous the number of scenarios is infinite. Solving such large DEP instances is not feasible in practice because of the excessive computation times.

In the SAA method, a random sample of N realizations (scenarios) of the random vector ξ is generated, and the expectation $E\left[Q(y,\xi)\right]$ is approximated by the sample average function $\frac{1}{N}\sum_{n=1}^{N} Q(y,\xi_n)$. Consequently, the original stochastic problem is approximated by the following deterministic equivalent problem containing N scenarios

$$DEP_N : \max_y \left\{ \hat{f}_N(y) \equiv cy + \tfrac{1}{N}\sum_{n=1}^{N} Q(y, \xi_n) \right\}. \tag{1.4}$$

Kleywegt et al. [8] have shown that the SAA method converges to the optimal solution and solution value. Based on the solutions y_i a lower bound to the original minimization problem and an incumbent feasible solution to the original problem can be computed. From these the optimality gap can then be derived. Further details can be found in Santoso et al. [14] and Kleywegt et al. [8]. Finally, the SAA algorithm uses an exterior sampling method, since the samples can be generated independently of the optimization method. This independence allows a modular structure of the overall algorithm, so that different optimal and heuristic methods can be used to solve (1.4). The evaluation of a given configuration with its corresponding recourse actions requires the solution of a large number (*N'*) of recourse problems that are themselves large multicommodity network flow problems with identical network structure. Typically, *N'* is chosen much larger than *N*. However, the most time consuming step in the overall algorithm remains the solution of the DEP problem with a large number of scenarios.

The DEP exhibits a block-diagonal structure which makes it suitable for primal (Benders) decomposition developed by Benders [1]. Van Slyke and Wets [17] developed the same primal decomposition algorithm in the framework of scenario-based stochastic programming. This algorithm is known as L-Shape decomposition. L-shaped decomposition is sometimes also referred to as stochastic Benders decomposition.

Benders decomposition separates the variables of the problem to be solved into easy and hard variables. The hard variables correspond to the configuration of the supply chain, which are the decision variables in the first stage of the stochastic program. Constraints on the supply chain configuration together with constraints or cuts based on the solution of the sub-problems constitute the master problem. The sub-problems correspond to the recourse decisions for the tactical production and distribution planning after the strategic variables have been determined. The sub-problems contain the easy variables, which are continuous variables in multicommodity network flow problems, with the temporarily fixed values of the hard variables used as parameters.

The computation times by a standard MIP procedure of a commercial solver (CPLEX MIP) for the monolithic model, by the standard Benders decomposition algorithm, and by the accelerated Benders decomposition algorithm for industrial sized supply chain design project in South America in function of the number of scenarios in the DEP were compared in Santoso [13]. Standard Benders decomposition becomes more efficient than the monolithic MIP algorithm if the number of scenarios exceeds 15. Both standard algorithms are strongly dominated by accelerated Benders, especially for large number of scenarios in the DEP. The time ratio is already 50 for a DEP with 20 scenarios. The accelerated Benders decomposition is the only algorithm that can solve the DEP in a reasonable amount of time for more than 20 scenarios.

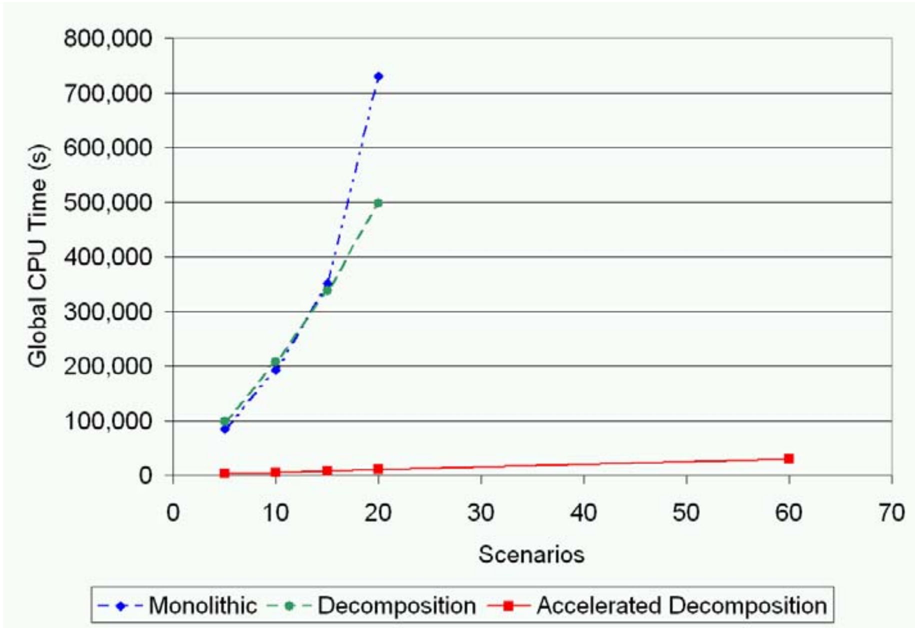

Figure 1: Computation Times in function of Number of Scenarios

2.2 Acceleration Techniques

For the complex strategic supply chain design problem Benders decomposition has poor convergence properties and a large relative difference between upper and lower bounds in the initial iterations as researchers like Magnanti and Wong [10], Santoso [13], Santoso et al. [14], and Santoso et al. [15] have reported. Many different acceleration techniques for Benders decomposition have been proposed. However, most of these acceleration techniques make a successful implementation even more difficult than the standard implementation. Geoffrion and Powers [6] report on the declining use of Benders decomposition in commercial software. Acceleration techniques for Benders decomposition and their descriptions can be found in the work of Magnanti and Wong [10], Magnanti and Wong [11], Dogan and Goetschalckx [5], Cordeau et al. [4], Santoso [13], Santoso et al. [14], Santoso et al. [15], and Goetschalckx [7]. The most powerful acceleration techniques developed up to date are the logistics constraints, trust region methods, primal heuristic methods, warm starting techniques, knapsack constraints, disaggregation of the Benders cuts, cut strengthening methods, and linear relaxation (Magnanti and Wong [10], Ruszczynski [12]; Magnanti and Wong [11]; Linderoth [9]; Santoso [13]; Santoso et al. [14]; Santoso et al. [15]; Goetschalckx [7]). These acceleration techniques are briefly explained in the following paragraphs considering a profit maximization problem for the master problem.

2.2.1 Logistics Constraints

Logistics constraints are also known as pre-processing constraints. They have been in existing for many years and have been proposed by a multitude of authors. The reason for the lack of unique authorship is that the logistics constraints uniquely depend on the problem structure and problem data. Hence, every problem demands a different set of logistics constraints. An example of such a constraint is a condition that specifies that at least one facility has to be established in each echelon of the supply chain so that a material flow path from external suppliers through the supply chain to the final customers can exist. Logistics constraints force the values of some of the master configuration variables to be either zero or one at the start of the solution algorithm. This forces the sub problems to yield much more feasible solutions and results in larger upper bounds during the first iterations of the decomposition algorithm, as reported by Santoso et al. [14], [15]. Finally, implementing and using the logistics constraints is not computationally expensive since the number of such constraints is considerably less than the total number of constraints in the final master problem after applying Benders decomposition.

Following is a list of such valid constraints. Not all constraints may apply to a particular supply chain formulation.

- The existence of paths or segment of paths with sufficient capacity to satisfy the demand at the termination nodes. The capacity must be sufficient for each commodity separately and for all commodities combined.
- The facilities of each echelon must also have sufficient capacity to satisfy the total customer demand on a per commodity basis and for all commodities combined.
- If only certain facilities can supply certain commodities to customers, one or more of these facilities must be enabled, which generates additional valid constraints.
- The disaggregated linkage constraints between machines or manufacturing lines inside a facility and the facility status also generate additional valid constraints.
- The linkage constraints between customer single sourcing variables and the status variable of the sourcing facility provide further valid constraints.
- Status variables for transportation channels are linked to the status of origin and destination facilities.
- Each enabled facility must have at least one incoming and outgoing channel enabled.

The existence of such valid constraints is checked in a backwards fashion, starting at the customers and working up the supply chain to the raw material suppliers. All of these constraints only have to be generated once for the root formulation. The disadvantage of adding all constraints at the root node is that a very large number of constraints that may be generated. The advantage is that the constraints remain valid for the whole branch-and-bound tree and constraint management is not required.

2.2.2 Trust Region Methods

The notion of trust region or regularized decomposition methods was introduced by Rusczynski [12]. The trust region method prevents the master problem solution from moving considerably from one region of the feasible domain to another during the first iterations by considering a regularized term into de objective function of the master problem. Linderoth and Wright [9] developed the idea of box-shaped trust regions, which have the advantage of allowing direct control over the size of the trust region. It is important to observe that in order to apply the trust region acceleration technique while running Benders decomposition, an initial master solution different from zero is required. Implementing the trust region method is not computationally demanding as reported by Santoso et al. [14], [15].

2.2.3 Primal Heuristic Methods

Several primal heuristic methods have been developed to improve the performance of Benders decomposition. They focus on particular characteristics of the problem under consideration. Santoso et al. [14], [15] propose a method in order to improve the upper bound of the master problem. They divide the master variables in major and minor variables, fix the major variables and then solve optimally for the minor variables. The major variables correspond to facilities and the minor variables correspond to different manufacturing capacity in the facility. Not only do they generate a new feasible solution of better quality, they also contribute to the knapsack constraints and to the reduction of the optimality gap.

2.2.4 Warm Start

Since the structure of the sub problems remains unchanged between iterations in the Benders decomposition algorithm, one way to improve the performance is take advantage of the optimal basis of the different sub problems in the previous iteration. This avoids having to solve the sub problems from scratch at each iteration. Details about taking advantage of warm starting are presented by Bertsimas and Tsitsiklis [2]. The computational burden of applying warm starts assumes that sufficient memory or fast disk cache is available since it involves keeping track of the sub problem and master problem bases in an efficient way.

2.2.5 Knapsack Constraints

The primal decomposition does not ensure that the lower bound or primal feasible solution is non-decreasing during the iterations. The knapsack constraints ensure that this event never happens, Santoso et al. [14], [15]; Goetschalckx [7]. In terms of computational complexity, this acceleration technique is not hard to implement since it involves adding a single constraint to the master problem.

2.2.6 Disaggregation of the Benders Cuts

The Benders decomposition algorithm adds one aggregated cut at every iteration. Cut disaggregation is an acceleration technique that disaggregates the cut into several cuts whose weighted sum is equivalent to the aggregated cut. The weight of each disaggregated scenario cut is given by the probability of each scenario. This technique was introduced by Birge and Louveaux [3] as the multi-cut. Obviously, the size of the master problem becomes greater in the case of disaggregated cuts. However, they allow a faster convergence of the master problem (Dogan and Goetschalckx [5]; Santoso et al.

[15]; Goetschalckx [7]. The computational effort required to apply this acceleration technique is considerable since the size of the master problem grows very quickly.

2.2.7 Strengthening Dual Variables

Magnanti and Wong [10] developed a method to strengthen the constraint or cut that is added to the master problem at every iteration. This is done by maximizing the function value of the cut at a point that is close to the true optimal solution or contained in the set of the feasible solutions of the master problem. Santoso [13] reports that using the strengthened cut is not always more efficient than the original cut, because the additional computational effort to find the stronger cut is considerable. A secondary linear program has to be solved at each iteration. Careful consideration should be given to the computational tradeoffs before applying this acceleration technique. Goetschalckx [7] propose strengthening of Benders cuts through heuristic adjustment of the dual variables. This acceleration technique requires the disaggregation of the problem by scenario. Dogan and Goetschalckx [5] developed a specialized algorithm for the adjustment of the dual variables in a domestic supply chain design formulation.

2.2.8 Relaxations

Another acceleration technique for Benders decomposition consists in using the objective function value of the linear relaxation of the master problem as an upper bound for the original master problem. This eliminates many solutions that are not acceptable in practice during the early iterations of the decomposition algorithm and significantly reduces the burden of solving the primal master problem (in the early iterations). The recourse problems are solved with the primal master variables found with the linear relaxation and may not yield cuts that are as strong as those generated with primal feasible master variables. However, the benefits of the acceleration technique seem to overcome this burden as Goetschalckx [7] and Dogan and Goetschalckx [5] report.

3 Statistical Experiment

Some of the previously mentioned acceleration techniques have been successfully implemented to solve real supply chain design problems in a reasonable amount of computing time. Magnanti and Wong [10] report that using their strengthening of Benders cuts resulted in a problem with at least two or three times fewer cuts than in the original Benders decomposition. Dogan and Goetschalckx [5] argue that disaggregation of the cuts reduced the computation time by a factor of 480 compared to a monolithic model solution. Cordeau et al. [4] applied what they call valid constraints to the master problem, achieving better computational results than in standard Benders decomposition. In summary, all of the research results agree that the application of a single acceleration technique improves the computational time of Benders decomposition.

Clearly, the additional performance of implanting an acceleration technique must be balanced against the additional effort to implement and maintain the programming code. Thus, it is desirable to determine the relative efficiency of the different acceleration techniques for Benders decomposition. Santoso [13], Santoso et al. [14], [15] provide computational results that demonstrate that the combination of several acceleration techniques yield the largest performance improvement. However, these results do not provide any statistical evidence defining the efficiency of the different acceleration techniques for Benders decomposition. Hence, a statistical experiment is designed to provide with the statistical evidence required and its development and future analysis is described in the following paragraphs.

In order to run the statistical experiment, an industrial single-period global supply chain design problem is selected. The model used in the experiment is described in further detail in Santoso [13]. A computer code was developed through modifications of the computer code supporting the results of Santoso [13]. The code generates a number of random scenarios based on mean values provided by the data from the problem. For the purpose of this experiment, 10 random scenarios are generated for every problem ($N=10$), since 10 scenarios allow the capture of the complexity of a real problem without incurring into extremely long computing times. Then, the computer code solves the problem using Benders decomposition with and without certain acceleration techniques.

A Single-Factor ANOVA model can be used to evaluate the relative efficiency of several combinations of acceleration techniques of Benders decomposition. Every combination of the acceleration techniques for Benders decomposition is treated as a level of the single factor. This experiment is used instead of the traditional Full Factorial ANOVA model since both models are mathematically equivalent and the same statistical results can be achieved. Using the Single-Factor ANOVA model reduces the complexity of the experiment when considering more than two acceleration techniques.

For the Single-Factor ANOVA model selected, a reasonable measure of the performance of every combination of acceleration techniques is given by the total computing time required to solve the problem. Furthermore, in this experiment Logistics Constraints, Trust Region, Knapsack Lower Bound, and Warm Starting acceleration techniques are permanently enabled. Disaggregation of the Cuts (D), Strengthening of Dual Variables (S), and Primal Heuristic (H) are acceleration techniques that can be turned on and off. Therefore, there are eight possible combinations of acceleration techniques applied to Benders decomposition in the present experiment. The same 10-scenario problem is tested under the eight possible combinations of acceleration techniques per replication. A new 10-scenario problem is generated for every replication.

The number of replications to be used for the current experiment is defined based on a specified tolerance and a 95% confidence interval for the computing time using all of the optional acceleration techniques. Based on previous experience running similar experiments and individual results of every possible combination of acceleration techniques, the tolerance of the experiment is set to 50 minutes. The software used to get the relevant results is C-PLEX 8.1 and Concert Technology 1.3, which was run on a Pentium III machine at 700 MHz with 512 MB in memory. The statistical analysis was performed using the Minitab 14 software.

With the defined tolerance, the number of replications is defined through the 95% confidence interval for the computing time using all of the acceleration techniques. For our experiment, based on initial results, 12 replications were required to differentiate the computation times at a tolerance level of 50 minutes. The number of replications may have to be adjusted based on the estimate of the variance of the computation times. At the current time 7 replications have been run for four combinations of acceleration techniques: DS, DH, SH, and DSH. If a single acceleration technique was turned on, i.e., either D, S, or H was enabled by themselves, then the computation times exceeded 24 hours. Hence, at the 95 % confidence level, it can be stated that using any single acceleration technique is different from using any combination of techniques. But because the computer runs were terminated after 24 hours, no statistical information exists to differentiate between individual acceleration techniques. Further computations are being executed to increase the number of replications. Based on the data collected so far, 7 replications yield a resolution of 47 minutes at the 95 % confidence level.

Based on 7 replications the following analysis steps were executed which yielded the following conclusions. The residuals were plotted in a normal probability plot and they tended to be normal. The null hypothesis that the variances are equal for the different levels was tested with Levene's Test, which yielded a p-value equal to 0.076. So at the 95 % confidence level there was not enough statistical evidence to reject the null hypothesis. The ANOVA analysis yielded a p-value of 0.000 indicating there was no statistical evidence that the level means were the same at the 95 % confidence level.

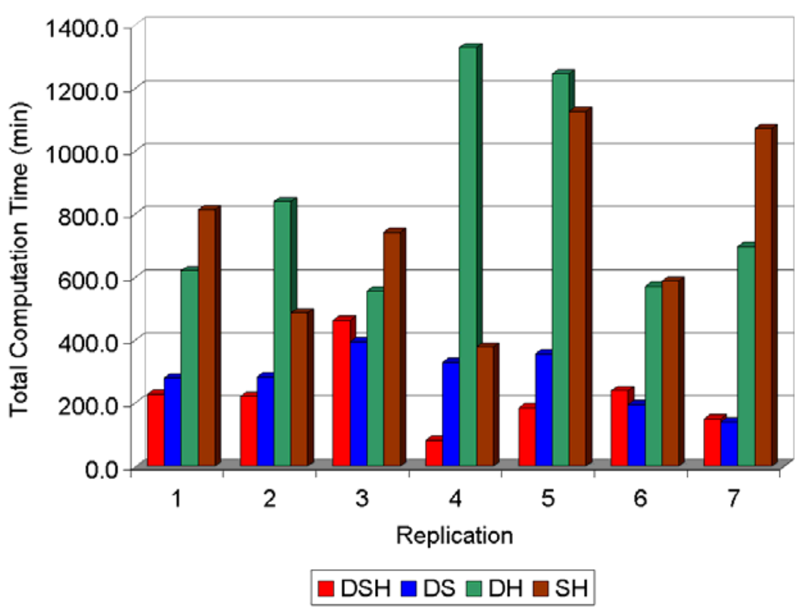

(D: Cut Disaggregation S: Dual Variable Adjustment, H: Primal Heuristic)

Figure 2: Total Computation Times for Different Acceleration Techniques

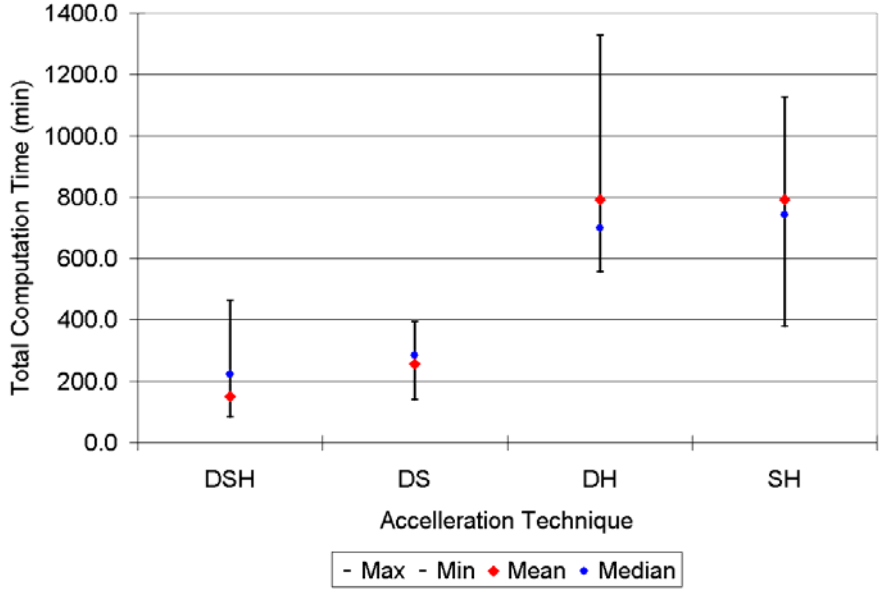

(D: Cut Disaggregation, S: Dual Variable Adjustment, H: Primal Heuristic)

Figure 3: Box Plots of Computing Times for Different Acceleration Techniques

The above box plot demonstrates that the combination of dual variable adjustment (S) and cut disaggregation (D) yields shorter computation times. The contrast of both techniques is different from zero at the 95 % confidence level. Using the primal heuristic (H) acceleration method with one other acceleration technique yields the longest computation times for any pair of acceleration techniques. Using the primal heuristic with the two other acceleration techniques does not yield a statistical different in average computation time. The contrast of using this acceleration technique is not different from zero at the 95 % confidence level. Based on these preliminary results for this specific case study, the implementation of this particular primal heuristic appears of very limited practical benefit.

3.1 Discussion of the Solution Robustness

In this research, the supply chain is optimized (designed) with respect to the expected value only, but after the optimization the supply chain is evaluated with random sampling to determine the expected value and variability measure. It would be more realistic to define the objective as to maximize the difference of the expected value of *NPV(NCF)* of the corporation over the planning horizon minus a compatible measure of the variability of the same *NPV(NCF)* weighted by a risk-preference parameter α. That is, the weighted multi-objective optimization problem is

$$\max\left\{\mathbf{E}\left[NPV\left(NCF\right)\right]-\alpha\cdot\mathbf{VM}\left[NPV\left(NCF\right)\right]\right\} \tag{1.5}$$

where $\mathbf{E}[\cdot]$ and $\mathbf{VM}[\cdot]$ denote the expectation and variability measure operators, respectively. One can employ different choices of the variability operator. Obvious choices are the mean absolute deviation (**MAD**), the standard deviation (**SD**) or the range (**R**) as the risk measure. However, for most of the risk measures the objective function becomes non-linear, non-concave, or both, which makes the optimization for large scale instances computationally very difficult. A particular strategic configuration of the supply chain will have a certain expected value and variability measure. Each configuration can be evaluated based on the solution of the recourse sub problems and then plotted in a classical risk analysis graph, with the expected value on one axis and the risk measure on the other axis. If a corporation knows the value of the α parameter that corresponds to its risk preferences, the preferred configuration can be determined immediately. More often, a corporation is interested in identifying several alternative high-quality supply chain configurations for various values of α. The efficiency frontier is the collection of supply chain configurations that are not Pareto-dominated by any other configuration, i.e. for any efficient or non Pareto-dominated configuration no configuration exists that has at the same time a larger expected value and a smaller risk operator. For a given set or sample of supply chain configurations that are located in the risk analysis graph, the upper envelope of those configurations can be determined. It is denoted by the sample efficiency envelope or SEE. This SEE is an approximation of the efficiency frontier based on the optimal solutions found for a number of random replications. The risk analysis graph for the industrial case is shown in the next figure including the SEE.

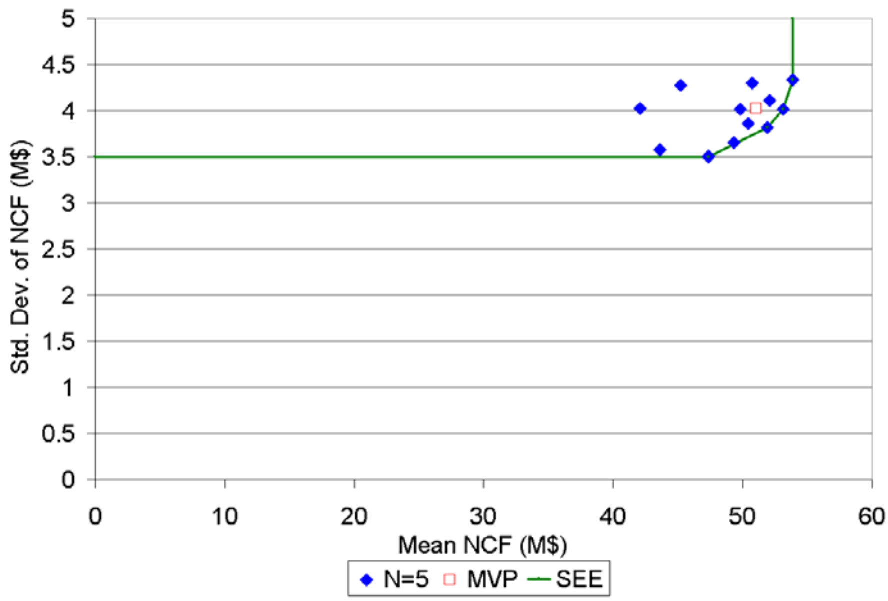

Figure 4: Risk Analysis Graph for Five Scenarios

The preferred supply chain configuration is shown to be very robust with respect to variability of the input data as illustrated in the next figure. Increasing the number of scenarios increases the robustness of the solution. The preferred configuration strongly Pareto-dominates the optimal solution of the mean value problem (MVP) and the difference becomes more pronounced if the data becomes more variable. The solution quality of a stochastic solution with a large number of scenarios in the DEP can only be achieved when using specialized accelerated decomposition algorithms. For the industrial test case of modest size, approximately 72,000,000 continuous flow variables were solved for to optimality to determine (design) and characterize (evaluate) the preferred supply chain configuration.

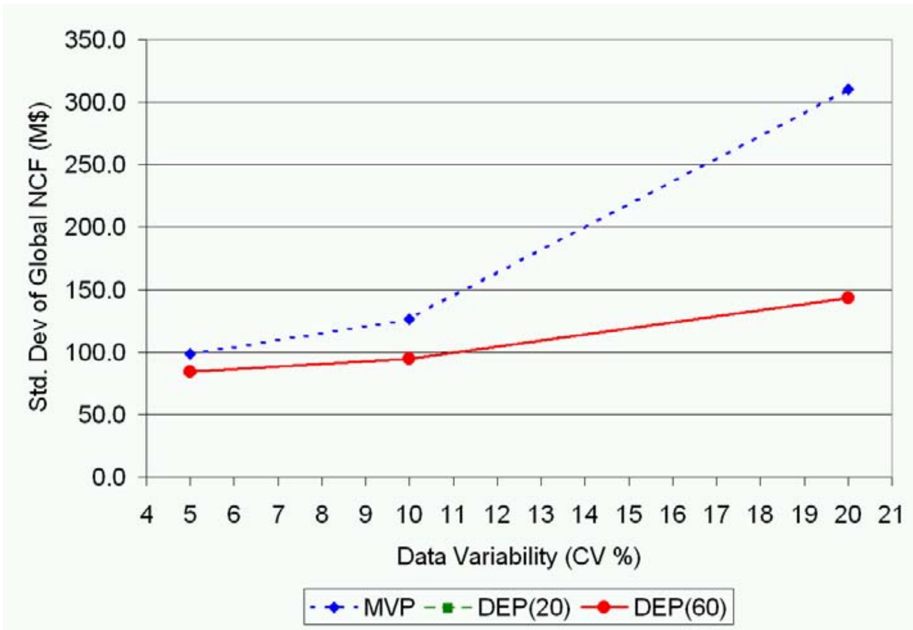

DEP(20) is nearly identical to DEP(60)

Figure 5: Robustness of the Range in function of Data Variability

4 Conclusions

The combination of the modeling and solution methodology provides for the first time a scientific method to evaluate the financial tradeoff between additional investment cost and more stable operational costs when configuring industrial global supply chains for a five year planning horizon. This method allows global corporations to balance high profit strategies against the robustness required to deal with real-word change.

However, the timely solution of the large industrial problem instances requires the implementation of acceleration techniques for Benders decomposition. At this time dual variable adjustment and cut disaggregation appear to be the most promising acceleration techniques. The contribution of the implemented primal heuristic appears to be marginal at best.

The supply chains are designed with goal of maximizing the expected value of the profit for multiple scenarios. This optimizing of the expected value has the side benefit of simultaneously reducing the risk of the supply chain configuration. This effect becomes stronger when more scenarios are included in the optimization problem.

The research trends appear to be focused on further integration of the strategic supply chain model and on the incorporation of various measures of robustness. Both trends create ever larger problem instances. Acceleration of Benders decomposition appears to be one of algorithms that are sufficiently fast to solve the problem in a reasonable amount of time. However, the performance of other optimizing or heuristic algorithms needs to be further investigated.

References

[1] Benders, J. (1962), "Partitioning Procedures for Solving Mixed-Variables Programming Problems," *Numerishe Mathematik*, 4, 238-252.

[2] Bertsimas, D. and J. N. Tsitsiklis, (1997), *Introduction to Linear Optimization*, Athena: Scientific Publishing.

[3] Birge, J. R. and F. V. Louveaux, (1988), "A multicut algorithm for two-stage stochastic linear programs," ***European Journal of Operational Research***, Vol. 34, No. 3, pp. 384-392.

[4] Cordeau, J. F, F. Pasin, and M. M. Solomon, (2003), "An Integrated Model for Logistics Network Design," accepted article waiting to be published in *Annals of Operations Research* journal.

[5] Dogan, K. and M. Goetschalckx, (1999), "A primal decomposition method for the integrated design of multi-period production-distribution systems," *IIE Transactions*, 31, 1027-1036.

[6] Geoffrion, A. M., and R. F. Powers, (1995), "Twenty years of strategic distribution system design: an evolutionary perspective", *Interfaces*, 25(5), 105-128.

[7] Goetschalckx, M., (2004), "Acceleration Techniques for Benders Decomposition in Strategic Stochastic Supply Chain Design", TLI-AP Technical Report, Singapore.

[8] Kleywegt, A. J., A. Shapiro, and T. Homem-de-Mello, (2001), "The sample average approximation method for stochastic discrete optimization," *SIAM Journal on Optimization*, Vol. 12, No. 2, pp. 479-502.

[9] Linderoth, J. and S. Wright, (2001), "Decomposition algorithms for stochastic programming on a computational grid," Research report, Mathematics and Computer Science Division, Argonne National Laboratory, 114.

[10] Magnanti, T. L., and R. T. Wong, (1981), "Accelerating Benders Decomposition: Algorithm Enhancement and Model Selection Criteria," *Operations Research*, 29 (3), 464-484.

[11] Magnanti, T. L. and R. T. Wong, (1990), "Decomposition methods for facility location problems," in: Mirchandani, P.B. and Francis, R.L. (Eds.) *Discrete Location Theory*, John Wiley and Sons, Inc., New York, p. 209-263.

[12] Ruszczynski, A., (1986), "A regularized decomposition method for minimizing a sum of polyhedral functions", *Mathematical Programming*, 35, 309-333.

[13] Santoso, T., (2003), "A comprehensive model and efficient solution algorithm for the design of global supply chains under uncertainty," Unpublished Doctoral Dissertation, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia.

[14] Santoso, T., S. Ahmed, M. Goetschalckx, and A. Shapiro, (2003a), "A stochastic programming approach for supply chain network design under uncertainty," Accepted for publication in the ***European Journal of Operational Research***. E-print available at www.optimization-online.org.

[15] Santoso, T., M. Goetschalckx, S. Ahmed, and A. Shapiro, (2003b), "Strategic Design of Robust Global Supply Chains under Uncertainty," submitted to ***IIE Transactions***.

[16] Shapiro, A. and T. Homem-de-Mello, (1998), "A Simulation-based Approach to Two-Stage Stochastic Programming with Recourse, *Mathematical Programming*, Vol. 81, pp. 301-325.

[17] Van Slyke, R. and R. J.-B. Wets, (1969), "L-Shaped linear programs with applications to optimal control and stochastic programming," *SIAM Journal on Applied Mathematics*, Vol. 17, pp. 638–663.