Proceedings of the 2006 Winter Simulation Conference
L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, eds.

A SIMULATION-BASED APPROACH TO TRADE-OFF ANALYSIS OF PORT SECURITY

Junko Sekine
Enrique Campos-Náñez
John R. Harrald
Hernán Abeledo

School of Engineering and Applied Science
The George Washington University
Washington, DC, U.S.A.

ABSTRACT

Motivated by the September 11 attacks, we are addressing the problem of policy analysis of supply-chain security. Considering the potential economic and operational impacts of inspection together with the inherent difficulty of assigning a reasonable cost to an inspection failure call for a policy analysis methodology in which stakeholders can understand the trade-offs between the diverse and potentially conflicting objectives. To obtain this information, we used a simulation-based methodology to characterize the set of Pareto optimal solutions with respect to the multiple objectives represented in the decision problem. Our methodology relies on simulation and the response surface method (RSM) to model the relationships between inspection policies and relevant stakeholder objectives in order to construct a set of Pareto optimal solutions. The approach is illustrated with an application to a real-world supply chain.

1 INTRODUCTION

The terrorist attacks of September 11, 2001, have dramatically changed public awareness of national security, in particular, the vulnerability of ports and waterways in the United States. The U.S. Coast Guard, Maritime Administration, Transportation Security Agency (TSA), Department of Homeland Security (DHS), and U.S. Customs and Border Protection have thus established programs designed to reduce the vulnerability of ports and waterways (see U.S. Custom and Border Protection 2004, Secretariat-UNCTAD 2004, U.S. Customs and Border Protection 2006), as 95 % of U.S. international trade moves in via water (U.S. Department of Transportation 1999). These new procedures, which combine law, regulations, government intelligence, and public-private partnership programs, can potentially alter the landscape of the supply-chain security. The programs' implications on the operation of global supply chains and ultimately the nation's economy are still largely unknown. Clearly, excessive security costs could threaten the economic viability of ports and of the maritime industry (Harrald et al. 2004, Looney 2002).

While both public and private sectors acknowledge the need for a secure transit of goods, the economic consequences of increased levels of security have become an increasing concern (Stana 2004, Harrald 2005). As a result, businesses involved in global supply chains around the world now face the task of enforcing security regulations set by the U.S. federal authorities while streamlining the operations of the supply-chain due to the drastic increase in the volume of international commerce worldwide (World Trade Organization 2004). Supply chains are complex economic and organizational systems, and achieving a reasonable level of security will require a systemic policy in order to guarantee their economic feasibility (Harrald et al. 2004).

In this context, trade-offs naturally arise as the organizations involved in the supply chain pursue different and often conflicting goals (Reese 2003). The methods of multicriteria decision making (Soland 1979, Hwang and Yoon 1981) form a natural basis for the examination of this problem. One potential approach is to build a multiattribute utility function (Lindley 1994) or take a cost-minimization approach (Tulsiani, Haimes, and Li 1990). While this approach has been successful in countless applications, there are substantial hurdles for its use in supply-chain security. One of the most salient is the fact that assessing the uncertainty of such rare events based on little or no empirical observations is extremely difficult (Bier et al. 1999, Lambert et al. 1994), and expert input (when available) can be unreliable and subjective (Bier et al. 1999, Rosoff and von Winterfeldt 2005). Another substantial hurdle is assigning a monetary value to the consequences of an inspection or security failure or other "surprise" events (Bier et al. 1999). As a consequence, most risk analysis models are scenario based (Rosoff and von Winterfeldt 2005).

An alternate approach is to characterize the trade-off between identified objectives in supply-chain security. In this approach, the concept of Pareto optimality (or Pareto efficiency) arises naturally and captures those operational

1-4244-0501-7/06/$20.00 ©2006 IEEE

modes in which no objective can be improved without sacrificing performance of the others (Soland 1979). The characterization of the set of Pareto optimal solutions can be very valuable in understanding how a private sector's objective, such as cycle time, is affected while trying to comply with one of the public sector's objectives such as attaining a certain inspection failure rate. In this approach, we must characterize the objective functions in order to understand the true relationship between input variables and relevant objectives. In some limited cases, closed-form results can be obtained for the objectives of interest. In all other situations, simulation can be used as a source of estimates as long as the objectives are quantifiable. Furthermore, since the scale of such systems makes simulations computationally expensive, one can resort to so-called metamodels that can be obtained through statistical sampling methods such as response surface methodology (RSM) to construct approximations of the system's behavior (Giddings et al. 2001, Srivastava et al. 1999, Shang et al. 2004).

In this paper, we consider the problem of designing inspection policies in a supply chain and present a simulation-based approach that can be used to characterize the set of Pareto optimal solutions or the corresponding *Pareto efficient set*. We subsequently apply this technique to the analysis of a real-world process, in order to understand the trade-off involved among system-wide Type I & Type II errors, cycle times, and cycle time variance. Here system-wide Type II errors represent public safety, while Type I errors, cycle time, and cycle time standard deviation are determinant to the economic viability of the supply chain. While qualitatively our results are intuitive, we argue that this additional knowledge—the best way to operate under certain security levels—is vital to the decision making and negotiation processes.

The paper is organized as follows: In Section 2 we define the problem of trade-off analysis as a multiobjective decision problem, and introduce our approach to construct a Pareto efficient set and Pareto efficient frontier. In Section 3, we describe the application of our approach to a real-world supply chain involving a Japanese firm. Finally, in Section 4, we present our conclusions and suggest future research.

2 SIMULATION-BASED TRADE-OFF ANALYSIS

The diverse objectives involved in implementing a security policy in supply chains clearly motivate the formulation of a multicriteria decision problem (Lahdelma et al. 2002). As argued before, the difficulties involved in estimating the probabilities of an extremely costly event make it practically impossible to construct an adequate utility function. Our approach calls instead for characterizing the trade-offs in the different objectives in order to support decision-making and negotiation among the parties involved.

2.1 Mathematical Framework

We consider a set A of possible security policies in the supply chain. For example, the most common security policy seen worldwide is the use of scanning machines on luggage (or items). In this paper, we will focus on a model of the supply chain in which multiple two-step inspections are placed in order to reduce the possibility of an unchecked container reaching the end of the supply chain.

Associated to each possible security policy $\alpha \in A$, there is a set of supply chain performance measures

$$f_1(\alpha), f_2(\alpha), \ldots, f_m(\alpha),$$

which represent the objectives of the stakeholders involved, e.g., TYPE II errors or the time that a product is to be delivered by the supply chain (cycle time). We assume that the stakeholders' objectives are advanced when functions $f_1, \ldots, f_m$ are minimized. In this context, the following definitions are useful. A security policy $\alpha \in A$ is called *Pareto optimal* or Pareto efficient if for every $\alpha' \in A$ such that $f_k(\alpha') < f_k(\alpha)$ for some objective k, then there will exist another objective $i \in \{1, 2, \ldots, m\}, i \neq k$, such that $f_i(\alpha) < f_i(\alpha')$. Alternately, a policy α' is said to be *Pareto dominated* if there exists another policy which is better in every performance measure, i.e., if $f_k(\alpha') < f_k(\alpha), \quad k = 1, 2, \ldots, m$.

The set of all Pareto efficient solutions, also called the *Pareto efficient set*, will be denoted $\mathcal{P}(A) \subset A$. Associated to these solutions, we define the Pareto efficient frontier as the set $F(A) = \{(f_1(\alpha), \ldots, f_m(\alpha)) | \alpha \in \mathcal{P}(A)\}$. This set captures naturally the trade-offs involved in the system, because it consists of the set of security policies that are *better than every other in at least one performance measure*. In a scenario in which a given level of security is to be enforced—as a consequence of a threat or available intelligence—the decision maker can understand how to enforce current regulations while advancing his or her own objectives. On the other hand—perhaps more importantly—it provides the basis for negotiation. Therefore, our approach, outlined in the following section, will be to characterize the set $\mathcal{P}(A)$ through a combined analysis/simulation-based approach.

2.2 Characterization of the Pareto Efficient Frontier

In general, it is difficult to characterize all the objective functions in closed-form and sometimes it is necessary to estimate some of these through simulation. In addition, given the complexity of the systems being modeled, simulations to obtain estimates of system-wide performance measures (e.g., cycle time) can be expensive.

In our approach, we look at a continuous set of security policies, namely the percentage of items that will be inspected at a set of predefined points in the supply chain. Under this working assumption, we use simulations to sample performance of the system for a subset of the set of policies A, and rely on the response surface method (RSM) (Khuri and Cornell 1996) to construct a function that imitates the behavior of the supply chain under the prescribed security policy (which is also called a metamodel or approximation architecture). The methodology of RSM involves the follows the steps:

1. Select appropriate inputs for an objective function and design the experiment to extract the relationship between inputs and their corresponding objective function values (e.g., the average cycle time in a supply chain would be a function of inspection rates and the data is obtained via simulation). In this case, we already know both the input variables, which are the levels of inspection, and their range—the interval $[0, 1]$.
2. Assume a form of the mathematical model with the inputs $\alpha_1, \alpha_2, ..., \alpha_n$, typically a polynomial. For instance, if a linear model is presumed for objective i, we would propose the form

$$\hat{f}_i(\alpha, \beta) = \beta_0 + \beta_1\alpha_1 + \cdots + \beta_n\alpha_n.$$

   In later steps, we will try to estimate the parameters $\beta = (\beta_0, \beta_1, \ldots, \beta_n)$ in such a way that

$$\hat{f}_i(\alpha, \beta) = f_i(\alpha) + \varepsilon_i,$$

   where ε_i is white noise, for every $i = 1, 2, \ldots, m$.
3. Select some data points $\{\alpha^1, \alpha^2, \ldots, \alpha^K\} \subset A$ for sampling.
4. Obtain the observations from the designed experiment $f_i'(\alpha^1), \ldots, f_i'(\alpha^K)$ for the pre-selected data points.
5. Estimate the parameters of the mathematical model (metamodel) assumed using Mean Squared Error (MSE) method. If the model fit (adjusted R^2) is low, go back to #1 to redesign the experiment for the inputs.
6. Test the performance of the metamodel ($\hat{f}_i(\alpha)$) with estimated coefficients (e.g., $\hat{\beta}_0, \hat{\beta}_1, \ldots, \hat{\beta}_n$). In our research, we compared the Mean Squared Error (MSE) in the observations used for constructing the metamodel (MSE_1) with the MSE obtained with additional sampling points(MSE_2). If $\mathrm{MSE}_1 \geq \mathrm{MSE}_2$ then accept the metamodel to predict other objective function values. Otherwise, go back to #1 to redesign the experiment. There MSE is obtained through the formula

$$\sqrt{\frac{\sum_{j=1}^{K}[f_i'(\alpha_j) - \hat{f}_i(\alpha_j, \hat{\beta})]^2}{K}}$$

When all objective function have been obtained, it is sometimes possible to fully characterize the Pareto efficient frontier. For instance, in (Soland 1979, Wilson 1967, Nakayama 1980) it is shown that if the objective functions $f_1, \ldots, f_m$ are continuous and convex in the set of possible security policies A, then there is a one to one correspondence between the Pareto efficient set and the set of solutions to optimization problems where the objective is a convex linear combination of the objectives. In our approach, we use the actual values obtained for $\alpha \in A = \prod_{i=1}^{n} [0, 1/K, 2/K, \ldots, 1]$ and identify Pareto optimal solutions with respect to this set of policies by inspection.

In summary, the Pareto efficient set and Pareto efficient frontier can be constructed as follows:

1. Find the set of security policies and objectives.
2. Characterize the objective function values. This can be achieved through analysis (closed-form), or through simulation. In the latter case, use response surface methods to create an approximation architecture of the objective function in question, also called metamodel.
3. Use the objective functions (or, in its case, the corresponding approximations) to characterize the Pareto efficient set.
4. Construct the Pareto efficient frontier that is simply the set of objective function values associated to the Pareto efficient set for trade-off analysis.

A similar approach has been applied successfully in the past in different contexts (Srivastava, Hacker, and Lewis 1999, Shang, Li, and Tadikamalla 2004). The application in the next section illustrates the approach.

3 APPLICATION TO THE CASE STUDY SUPPLY-CHAIN MODEL

As a part of the public sector program, TSA has funded Boeing and General Electric (GE) to develop a highly sophisticated explosive detection system (EDS) since 2002 and has tried to test it in the second quarter of calendar year 2005 (McCarter 2005). This equipment will be used as an initial inspection in a 2-step inspection procedure.

Testing of EDS equipment in real life would require a lot of planning and a massive budget. In order for the inspection operation to be more effective in real life, we can combine the idea of multiple inspections, which have been used as a common approach to increase the inspection

Figure 1: Container Explosive Detection System (EDS)

accuracy, and drop the possible TYPE I and TYPE II error rates (Raouf, Jain, and Sathe 1983, Duffuaa and Al-Najjar 1995, Duffuaa and Al-Najjar 1997, Kaio and Osaki 1989, Qi, Tang, and Sivakumar 2002).

We now apply this methodology to a real-world supply chain. The simulation model for this process was constructed with the input provided by the supply-chain security experts (expert judgment) and publicly available information (e.g., SAIC EDS's product information on its sales web site).

In applying our research methodology to a realistic case, we constructed an "as-is" simulation model and identified a set of possible inspection points. We approached this likely supply-chain process by viewing it as a "system of systems" (Harrald, Hugh, and vanDorp 2004) and separating the supply chain into subsystems each containing a two-step inspection subprocess. The two steps correspond to the use of EDS, followed by a manual inspection to be used only when a container has been deemed as suspicious in the first stage. In order to achieve the objectives of both public and private sectors, the decision maker will have to understand the combined behavior of these subsystems and the trade-offs that arise using different inspection policies, consisting in this case of the inspection rates to be enforced at each point.

The possible trade-offs between the public and private sectors are measured by the following objective functions. First, to represent the public sector's viewpoint, we used the probability of Type II errors—falsely accept a hazardous container—or $\mathrm{PFA}_n(\alpha)$. Second, in the private sector, the probability of Type I errors—falsely rejecting clean containers—or $\mathrm{PFR}_n(\alpha)$, the average cycle time of the containers, denoted $\mathrm{CT}(\alpha)$), and the standard deviation of the cycle time of the containers, denoted $\mathrm{CTSTD}(\alpha)$, will be used to represent the private sector's objectives.

The supply-chain process modeled here consists of three sequential processes occurring in two locations. The first location is a firm's warehouse, where products, which are split into two parts, arrive on palettes (One palette can hold half the product, therefore, two palettes are required per product). They are inspected, packed to a container (three palettes per container), and sealed by a local inspector. The second location is the local port or "port of origin" of containers where container inspection and loading would be conducted.

First, we constructed a base model, which closely replicates the statistics of this particular supply-chain, using parameters provided by experts or estimated by our team based on experts' observations of this particular supply chain in Japan. The statistics and expert observation inputs used are discussed in the following section.

3.1 Case Study Model Definition

The firm has a factory to regularly manufacture a single product. Once a product is manufactured, it is split into two parts and sent to the firm's warehouse. In the warehouse, the split-up product, which is placed on palettes, will be inspected and loaded into a container and sealed by the local custom personnel every Thursday. The local custom inspector can inspect up to three containers per day, and each container can store three palettes. The freight truck will pick up the container and bring it to the local port. At the port's gate, the truck driver has to provide the container's manifest as well as his driver's license. Once the truck passes the gate, the driver will direct the truck to the container waiting area inside of the port, unload the container, and leave. All containers need to be physically stored in the port at least 72 hours prior to loading due to the port regulation in the port of origin. On loading day, containers are moved from the waiting area to the loading area by two forklifts. It will take a day to move all containers to the loading area to the freight ship, which arrives on Wednesdays and leaves on Thursdays. All operations are conducted during regular business hours (Monday through Friday, 8 am to 5 pm). The average cycle time of the product in warehouse is ten days though it varies from one day to twenty-one days. The average cycle time of the container in the port of origin is six days (five days for waiting and one day for loading).

In this supply-chain model, we identify three possible inspection points available. The first inspection point could be in the warehouse before local custom inspection. The second inspection point could be at the gate of the port of origin (see Fig. 2). The last inspection point would be the waiting area of containers in the port of origin. At each inspection point, products or containers are randomly selected with α_i% of the total products/containers arriving to inspection point i, where $i = 1, 2, 3$, then sent to a scanning inspection. If the product/container fails, it would

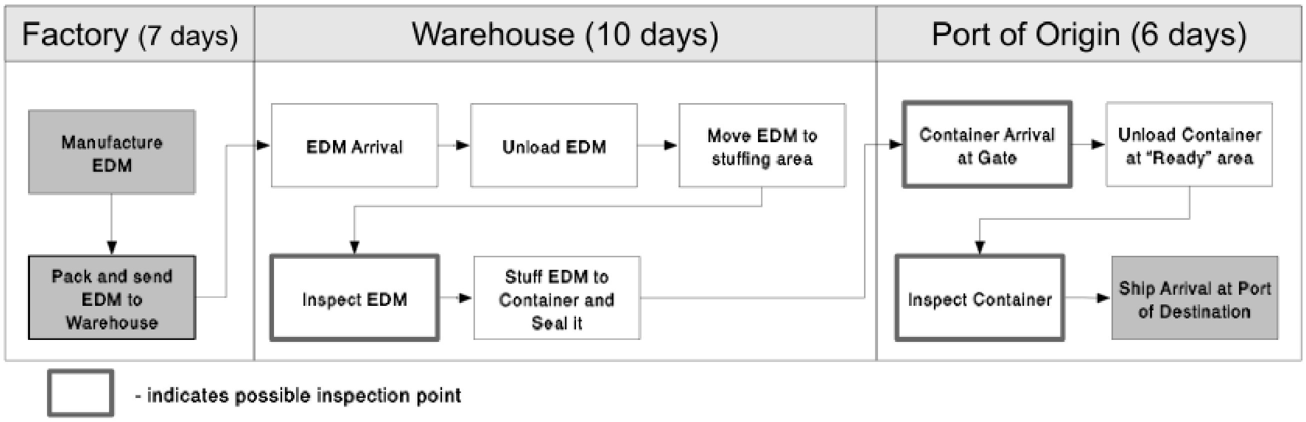

Figure 2: Process Flow of the "As-Is" Case Study Model (Simplified)

be sent to a manual inspection. Both scanning machines for products and containers have the same Type I/II error rates. Manual inspection for both products and containers is assumed to be perfect.

The objective functions of interest, which were described in the previous section, can be characterized as follows:

1. System-wide probability of a Type II error, which can be computed as
$$\mathrm{PFA}_n(\alpha) = E[\mathrm{FA}_3(\alpha)/N],$$
where $\mathrm{FA}_3(\alpha)$ is the number of Type II errors at inspection point 3 under inspection policy α, and N is the total number of pallets.
2. The expected value of negative failure of the containers in the system can be computed as
$$\mathrm{PFR}_n(\alpha) = E\left[\sum_{1\leq j\leq 3} \mathrm{FR}_j(\alpha)/N\right],$$
and N is defined as above, and $\mathrm{FR}_j(\alpha)$ is the number of negative failures at inspection station j, under policy α.
3. The expected cycle time $\mathrm{CT}(\alpha)$ was approximated through RSM using 30 policy samples with 10 replications. Standard warm-up analysis (Welch 1981) was used using 10 replications, resulting in simulation runs of 1,000,000 days.
4. The standard deviation of the cycle time $\mathrm{CTSTD}(\alpha)$ will also be estimated using a simulation/RSM approach with the same number of samples and replications.

We constructed simulation models of the supply-chain with multiple two-step inspections using Extend 6.0 software.

Surprisingly, there was *a negligible response in CT*(α) *or CTSTD*(α) *due to the buffer in the process itself and the low frequency of container arrival, e.g. the system modeled showed load factors far from critical and could accommodate* 100% *inspection rates in all inspection points without any impact to its cycle time*.

3.2 A Critically Loaded System

To test our methodology on a more meaningful scenario, we scaled the system arrival rate up to a critical value that, while maintaining a steady-state accommodating 100% inspection rates, would have a significant increase in its cycle time and its corresponding variance.

To identify the critical arrival rate, we increased the original good arrival frequency in the 100% inspection scenario. To obtain the critical arrival rate, we assumed infinite buffer sizes; the critical arrival rate is defined as the largest observable arrival rate that produces average queue sizes that are within the real process buffer capacities. In this case, we identified a critical arrival rate of 8.39 times of the original arrival frequency ($\frac{3.5days}{8.39}$ or every 0.417 day). Once the arrival frequency became greater than this critical arrival rate, the maximum number of stocks in the warehouse reached average values exceeding its capacity of 18 containers, and its warehouse average cycle time became higher than 18 days with a standard deviation of 396.

Using this critical arrival rate, we randomly selected 39 inspection policies, together with two extreme policies (no inspection at any point, 100% inspection at every inspection point). Ten replications were generated for each selected inspection policy by simulation. We then used 410 simulation samples to find the meta-models of $\mathrm{CT}(\alpha)$ and

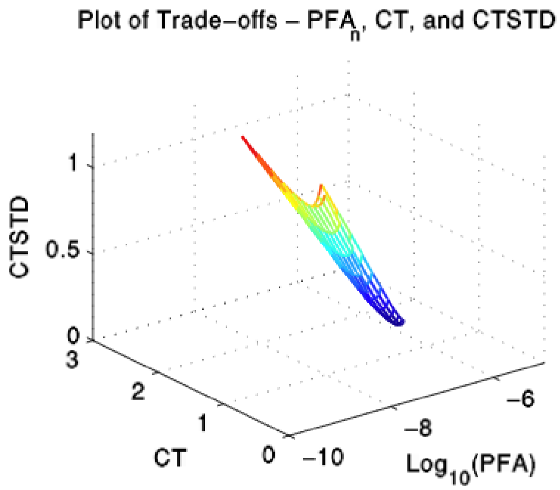

Figure 3: Pareto Efficient Frontier for the Critically Loaded System Described in Section 3.2 for Objectives PFA$_n(\alpha)$ and CT(α) and PFA$_n$ and CTSTD(α)

CTSTD(α). In addition, twenty inspection policies were randomly selected, and ten replications were also generated in the same manner by simulation. Using these additional 200 simulation samples, we further validated the robustness and accuracy of the metamodels of CT(α) and CTSTD(α).

3.3 Results of the Case Study Model Under Critical Load

In general, the more one inspects in the warehouse, the longer the cycle time and its standard deviation become, as expected. The 3rd inspection (2nd inspection at port) has no impact because any delay caused by the 3rd inspection would be absorbed during the 72 hour waiting period required by the local port authority. The probability PFA$_n(\alpha)$ responds nonlinearly to inspection rates; the risk levels decrease slowly when inspection rates are relatively low, but drop down to approximately $1/50$ of the inspection free risk level (10^{-7}), *once the average inspection rate, i.e.,* $\overline{\alpha} = (\alpha_1 + \alpha_2 + \alpha_3)/3$*, is greater than* 80%.

Additionally, PFR$_n(\alpha)$ is almost linear to inspection rates increase because of its very structure; the more one inspects, the more one rejects the containers with no hazardous materials (system-wide Type I error). However, in this model, perfect inspection was assumed in the manual inspection, thus this objective function value became zero, which indicates that it is indifferent to trade-offs. For a given objective function value, e.g., CT$(\alpha) =$ CT*, the trade-offs were recognized in PFA$_3(\alpha)$ and CT(α), PFA$_3(\alpha)$ and CTSTD(α) but not in CT(α) and CTSTD(α). For CT(α) and CTSTD(α), the metamodels indicated the strong influence from the first inspection point (α_1). However, PFA$_3(\alpha)$ is symmetric in $\alpha_1, \alpha_2, \alpha_3$ (i.e., policies $[0.3, 0.5, 0.8]$ and $[0.8, 0.5, 0.3]$ have the same PFA$_3$).

The approximation of the cycle time CT(α) was constructed by the 1st order of linear combination of 1st and 2nd inspection rates using 39 randomly selected inspection policies and two extreme inspection policies $(0, 0, 0)$, and $(1, 1, 1)$. The 3rd inspection rate was removed due to the statistical insignificance and did not add a significant accuracy of the model or a better fit even for additional 20 randomly selected inspection policies, which were used to test the robustness of the approximation of CT(α).

For CTSTD(α), the same inspection policies and procedure were taken to collect the simulation samples. The metamodel of CTSTD(α) is a nonlinear design that uses inspection rate α_1 as its only input. This is due to the fact that the proportion of warehouse cycle time standard deviation to the total cycle time standard deviation was more than 99% and there was very little impact from inspections at port because of 72 hours waiting period.

Both CT(α) and CTSTD(α) metamodels had the coefficients to be statistically significant with a white noise, i.e., where independent errors are normally distributed with mean 0, and variance $\sigma^2 > 0$. The influence of inspection rates to both CT(α) and CTSTD(α) was significant, and all coefficient estimates were proved to be statistically significant with a 5% error level using a t-test. The metamodels constructed through RSM for CT(α) and CTSTD(α) are given by

$$\begin{aligned} \widehat{CT}(\alpha_1, \alpha_2, \alpha_3) &= 11.52 + 2.22\alpha_1 + 0.19\alpha_2, \\ \widehat{\mathrm{CTSTD}}(\alpha_1, \alpha_2, \alpha_3) &= 1.77 * 1.03^{\alpha_1}. \end{aligned}$$

3.4 Summary of the Case Study Model Experiments

As mentioned before, we found that the original process could actually handle 100% inspection rate at all inspection points with no impact to either CT(α) or CTSTD(α). However, once the system approached a critical arrival rate we saw the impact and trade-offs from the inspection(s) on CT(α) and CTSTD(α).

4 CONCLUSIONS

The results showed the relationship between input variables (e.g., inspection rate at each inspection point) and objective functions. Also, the approximations of objective functions obtained by simulation and response surface method (RSM) were relatively straightforward (e.g., first order model) due to the simplified structure of the research model. For the case study model, CT(α) and CTSTD(α) were also functions of inspection rates yet not all inspection rates and the form of CTSTD(α) was highly adapted the model ruling such as 72

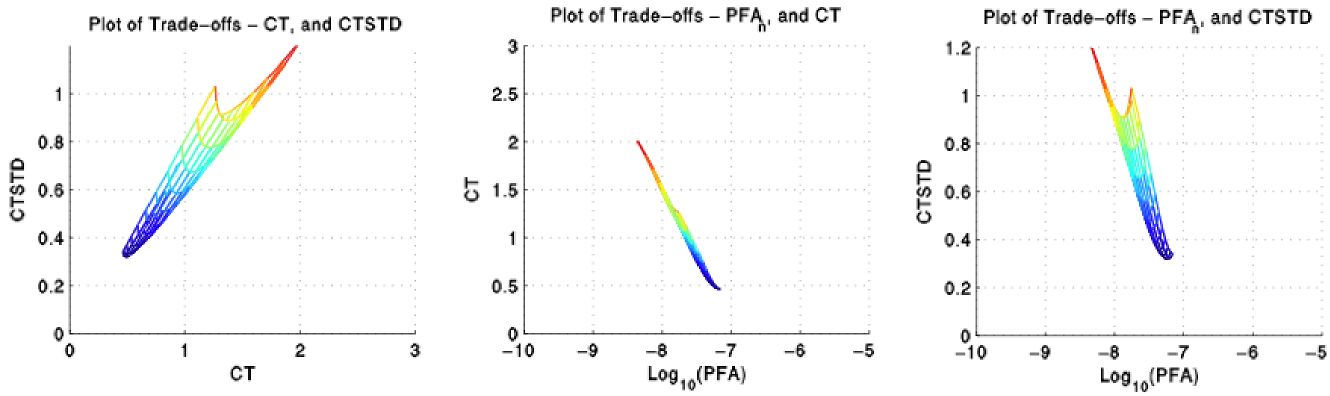

Figure 4: Projections of the Pareto Efficient Frontier for the Critically Loaded System Described in Section 3.2 for Objectives $\mathrm{PFA}_n(\alpha)$ and $\mathrm{CT}(\alpha)$ and PFA_n and $\mathrm{CTSTD}(\alpha)$

hour requirement of container arrival in port prior to loading. This shows that actual business rules drive a major impact of the model behaviors, and for a decision maker, it would be required to understand the impact from current business rules and/or the robustness of Pareto optimal policies.

Our future research will focus on testing the methodology of other objective functions such as inspection implementation cost, considering a different inspection policy and/or assumptions (e.g., more inspection points, dynamic hazard rate in supply-chain).

REFERENCES

Bier, V., Y. Haimes, J. Lambert, N. Matalas, and R. Zimmerman. 1999. A survey of approaches for assessing and managing the risk of extremes. *Risk Analysis* 19 (1): 83–94.

Duffuaa, S., and H. Al-Najjar. 1997. A general inspection plan for critical multicharacteristic components. *International Journal of Production Research* 35 (10): 2723 – 2736.

Duffuaa, S., and H. Al-Najjar. 1995. An optimal complete inspection plan for critical multicharacteristic components. *The Journal of the Operational Research Society* 46 (8): 930.

Giddings, A. P., T. Bailey, and M. J. T. 2001. Optimality analysis for facility location problems using response surface methodology. *International Journal of Physical Distribution & Logistics* 31 (1): 38.

Harrald, J. 2005. Sea trade and security: An assessment of the post-9/11 reaction. *Journal of International Affairs* 59 (FALL/WINTER 2005): 157 – 178.

Harrald, J., S. Hugh, and R. vanDorp. 2004. A framework for sustainable port security. *Journal of Homeland Security and Emergency Management* 1 (2): Article 12.

Hwang, C., and K. Yoon. 1981. *Multiple attribute decision making, methods and applications: a state-of-the-art survey*. Springer-Verlag.

Kaio, N., and S. Osaki. 1989. Comparison of inspection policies. *The Journal of the Operational Research Society* 40 (5): 499–503.

Khuri, A., and J. Cornell. 1996. *Response surfaces: design and analyses*. 2nd Edition ed. New Yourk: Marcel Dekker, Inc.

Lahdelma, R., S. Makkonen, and P. Salminen. 2002, July, 31. Modelling dependent uncertainties. *Stochastic Multicriteria Acceptability Analysis*.

Lambert, J., N. Matalas, C. Ling, Y. Haimes, and D. Li. 1994. Selection of probability distributions in characterizing risk of extreme events. *Risk Analysis* 14 (5): 731–742.

Lindley, D. V. 1994. *Making decisions*. Second Edition ed. John Wiley & Sons Ltd.

Looney, R. 2002, August. Economic costs to the United States stemming from the 9/11 attacks. Technical Report, the Center for Contemporary Conflict (CCC).

McCarter, M. 2005. The programs of 2005. *HSToday* 2 (1): 20–23.

Nakayama, M. 1980. Nash equilibria and Pareto optimal income redistribution. *Econometrica* 48 (5): 1257–1264.

Qi, C., K. Tang, and A. Sivakumar. 2002. Simulation based cause and effect analysis of cycle time and WIP in semiconductor wafer fabrication. In *Winter Simulation Conference 2002*. San Diego, CA.

Raouf, A., J. Jain, and P. Sathe. 1983. A cost-minimization model for multicharacteristic component inspection. *IIE Transactions* 15 (3): 187.

Reese, A. 2003, June-July. Building the secure supply chain. *iSource Business* (site accessed on April 2, 2005).

Rosoff, H., and D. von Winterfeldt. 2005, October. A risk and economic analysis of dirty bomb attacks on the ports of Los Angeles and Long Beach. Technical Report N00014-05-0630, Center for Risk and Economic Analysis of Terrorism Events (CREATE) University of Southern California.

Secretariat-UNCTAD 2004. UNCTAD report - container security: Major initiatives and related international developments. Technical Report ISO/TC N 997.

Shang, J., S. Li, and P. Tadikamalla. 2004. Operational design of a supply chain system using Taguchi method, response surface methodology, simulation, and optimization. *International Journal of Production Research* 42 (18): 3823 – 3849.

Soland, R. 1979. Multicriteria optimization: A general characterization of efficient solutions. *Decision Science* 10: 26–38.

Srivastava, A., K. Hacker, and K. Lewis. 1999. Development of a kriging based surrogate approximation method for large scale systems. *Design Optimization: International Journal for Product and Process Improvement*.

Stana, R. 2004. Summary of challenges faced in targeting oceangoing cargo containers for inspection. Technical Report 17, U.S. General Accounting Office.

Tulsiani, V., Y. Haimes, and D. Li. 1990. Distribution analyzer and risk evaluator (dare) using fault trees. *Risk Analysis* 10 (4): 521–538.

U.S. Custom and Border Protection 2004, November. Securing the global supply chain - custom-trade partnership against terrorism (C-TPAT) strategic plan. Technical Report, U.S. Custom and Border Protection.

U.S. Customs and Border Protection 2006. Fact sheet: When was the container security initiative developed and why?

U.S. Department of Transportation 1999, September. An assessment of the marine transportation system: A report to congress. Technical Report, U.S. Department of Transportation.

Welch, P. D. 1981. On the problem of the initial transient in steady-state simulation. *IBM Watson Research Center*.

Wilson, R. 1967. A Pareto-Optimal dividend policy. *Management Science* 13 (9): 756–764.

World Trade Organization 2004. World trade report 2004. Technical Report, World Trade Organization.

AUTHOR BIOGRAPHIES

JUNKO SEKINE is a doctoral candidate at Department of Engineering Management and Systems Engineering. Her e-mail address is `<jsekine@gwu.edu>`.

ENRIQUE CAMPOS-NÁÑEZ is an assistant professor of systems engineering in the George Washington University. He is a member of INFORMS. His e-mail address is `<ecamposn@gwu.edu>`.

JOHN HARRALD is professor of engineering management and systems engineering at the George Washington University, where he is also director of the Institute for Crisis Disaster and Risk Management. His e-mail address is `<harrald@gwu.edu>`.

HERNAN ABELEDO is an associate professor of engineering management and systems engineering at the George Washington University. His e-mail address is `<abeledo@gwu.edu>`.