Future Resilient Systems Technical Notes Series March 2023 No. 03

A Robust-Stochastic Optimization Approach for a Logistic Planning Problem

Prepared by: Dr Alberto Costa, *Project Coordinator and Senior Researcher, FRS*; Dr Jonas Joerin, *Programme Co-Director, FRS*
Contributors: Mr Chan Wen Feei, *Head (OR – Optimisation), DSTA*; Mr Chen Mingyi Edmund, *Principal Analyst, DSTA*; Mr Tan Bo Wen, *Senior Analyst, DSTA*

(FRS) FUTURE RESILIENT SYSTEMS

→ Planning for training events is a challenging task because coordination is necessary. Training sites need to be identified, and specific equipment should be at each training site during the events. Often, key decisions such as when and where to host training events and how much equipment to send to sites need to be defined much earlier than the starting date of the events and cannot be easily modified later. In real applications, the starting time of each event may deviate from what was scheduled, and the amount of equipment present at each training site when the event starts is also uncertain.

→ Optimization under uncertainty (stochastic and robust) has been employed to devise an equipment redistribution plan between the training sites so that equipment requirements of training events are fulfilled as much as possible, even in the presence of uncertain start times and equipment availability.

→ Results show that the solutions found by considering uncertainty outperform the deterministic ones where uncertainty is not addressed, when evaluating the trade-off between performance, i.e., fulfillment of events, and feasibility.

→ The significance of this work is that a decision-support tool based on optimization under uncertainty can help find solutions more realistic and feasible than those identified with deterministic approaches, thus improving future planning strategies.

The Logistic Planning Problem

We consider a problem in Singapore, where training events are scheduled at some locations, and each event requires specific equipment to be run. At each period, the equipment can be either redistributed between training sites or assembled from components in a warehouse and then shipped to training sites. Such decisions must be taken much earlier than the start of the training events. The objective is to fulfill the equipment requirements of events as much as possible.

The abovementioned problem can be modeled as a Mixed-Integer Linear Program (MILP), i.e., an optimization problem with linear objective function and constraints where some variables are integer and solved directly by available MILP solvers like CPLEX (IBM ILOG, 2020). Such solvers take as input the MILP model representing the problem and return as output the optimal solution. At the same time, some key parameters, such as the starting time of training events and the amount of equipment present at each site when events start, are affected by uncertainty. Therefore, a deterministic approach (where no uncertainty is considered) is not suitable.

A more realistic version of the problem aims at determining the amount of equipment to be redistributed between training sites and to be shipped from the warehouse to training sites when considering uncertainty on the abovementioned key parameters. This means that such decisions, termed *here-and-now*, must be taken before (and should ideally remain feasible after) such parameters are realized. Given some target equipment fulfillment level for each event, i.e., the minimum expected fraction of required equipment that should be present at the training site for the duration of the event, the objective is then to maximize the number of events for which such targets are achieved. Indeed, as uncertainty is involved, one may maximize the expected number of such events in a stochastic optimization fashion (Spall, 2005), where the average performance is optimized over a sample of possible scenarios, or optimize the performance of the worst-case scenario with a robust optimization approach (Ben-Tal et al., 2009). In this work, we decided to use a mixed approach as summarized in the next section.

The main goals of this project are as follows:

- Create a model of the Logistic Planning Problem under uncertainty.
- Develop an optimization approach to solve the problem.

Apply the solution procedure to a test case from DSTA and understand the implications of the results obtained with different settings.

Methodology

We show here how the Logistic Planning Problem has been addressed. The main idea, whose details are explained in the following, is to select a subset of n random scenarios, optimize the performance against them (either with a stochastic or a robust optimization approach), and test the performance of the obtained solutions using another set of $N > n$ random scenarios.

Let the Logistic Planning Problem be formulated as follows:

$$\max_{x \in D^x,\, y \in D^y(x,z)} t(x,y,z),$$

where x and y are the decisions to be taken before and after the realization of the uncertainty z, respectively, and y includes integer variables. The objective function (performance measure) $t(\bullet)$ to be maximized defines the number of events for which the target equipment fulfillment level is achieved. Feasible x and y decisions belong to sets D^x and $D^y(x,z)$, respectively. Constraints and the objective function are assumed to be linear in x and y.

In the deterministic version, z is known and the problem is an MILP that can be solved with CPLEX. As explained earlier, we address the case where z is random. Let $T^*(x,z)$ be defined as:

$$\max_{y \in D^y(x,z)} t(x,y,z),$$

i.e., the optimal solution value of the problem to solve after the uncertainty z is realized and the here-and-now decisions x are defined. A stochastic-optimization approach aims at solving the following problem:

$$\max_{x \in D^x} \mathop{E}_{z \in U}[T^*(x,z)],$$

where U is the set of possible realizations of the uncertain parameters z, and $E[\,]$ is the operator representing the expectation, which is often approximated with the average over a subset of U. On the other hand, a robust optimization approach would solve the following problem:

$$\max_{x \in D^x} \min_{z \in U} T^*(x,z),$$

where the objective is to optimize against the worst-case scenario. Unfortunately, finding $T^*(x,z)$ involves an optimization over (integer) variables y and requires solving an MILP. Therefore, a sample average approximation approach to solve the stochastic optimization version of the problem would be computationally expensive when the number of samples is large. Solving the robust problem also requires some iterative approach since duality cannot be used to obtain a single-level model due to the integrality of variables y. As a consequence, the robust solution, which is often conservative, may not be computationally cheap to find.

The approach we implemented to address the problem takes elements from both the stochastic and robust methodologies and consists of the following steps:

1. Identify n random samples $z_1, \dots, z_n$ from the set U.
2. Solve the *average-case maximization* (ACM*) problem*:
$$\max_{x \in D^x} \frac{1}{n} \sum_{i=1}^{n} T^*(x, z_i).$$
3. Solve the *worst-case maximization* (WCM) *problem*:
$$\max_{x \in D^x} \min_{i \in \{1,\dots,n\}} T^*(x, z_i),$$
that can be reformulated as follows:
$$\max_{x \in D^x} \tau$$
$$s.t.\ \forall i \in \{1,\dots,n\} \quad \tau \le T^*(x, z_i).$$
4. Compare the solutions x obtained at steps 2 and 3 against a set of $N > n$ random samples taken from set U in terms of performance and feasibility.

In both cases, the duplication of the problem structure is needed. The option of point 2 is the sample average approach often used in stochastic optimization. Therefore, the number of samples n cannot be too large. At the same time, the robust version of point 3 could benefit from a limited number of samples and provide solutions less conservative, albeit potentially less feasible.

Results and Insights

We choose $N = 1000$ and $n = 25$ as numbers of samples for testing and training the model, respectively. We first look at the results obtained from solving the ACM problem. Figure 1 shows the feasibility curve of the solution when tested against the N samples.

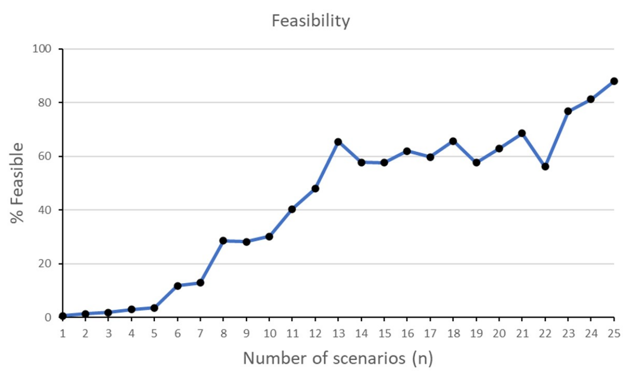

Figure 1 Feasibility of the ACM solution across the n samples.

We can observe an overall increase in feasibility as the problem captures a larger number of scenario samples (n).

Figure 2 below shows the performance curves obtained from plotting the best (max), average, and worst (min) performance of the ACM solution when tested against the N samples (infeasible instances excluded from performance computation).

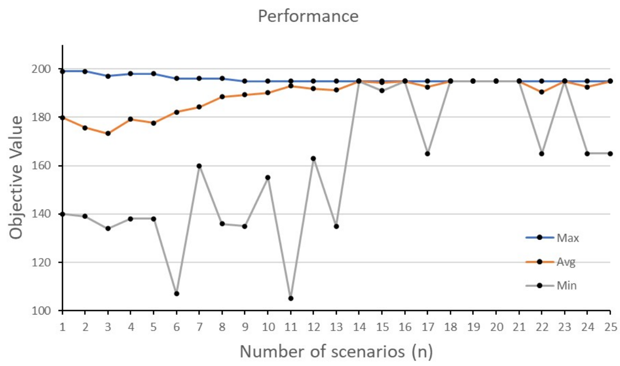

Figure 2 Best, average, and worst performance of the ACM solution across the n samples without infeasible solutions.

We can observe small fluctuations with signs of convergence of all 3 curves.

If we include the infeasible[1] instances in performance computation, the average performance of the ACM solution when tested against the N samples will have a similar trend to the feasibility curve in Figure 1. The best and worst performances are omitted as they will not be affected by the infeasible instances.

We next look at the results obtained from solving the worst-case maximization problem. Figures 3 and 4 are the curves analogous to Figures 1 and 2, respectively. In Figure 3, similar to Figure 1, we can observe an overall increasing trend. In Figure 4, we can observe relatively

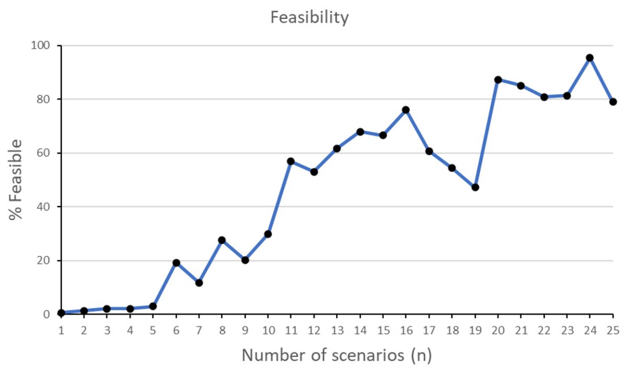

Figure 3 Feasibility of the WCM solution across the n samples.

larger fluctuations with signs of convergence of only 2 curves, i.e., best-case and average-case.

Similarly, if we include the infeasible instances in performance computation, the average performance of the WCM solution when tested against the N samples will have a largely similar trend to the feasibility curve in Figure 3, except for one instance between $n = 11$ and $n = 12$ where the trend differs due to a huge spike in the average performance in Figure 4.

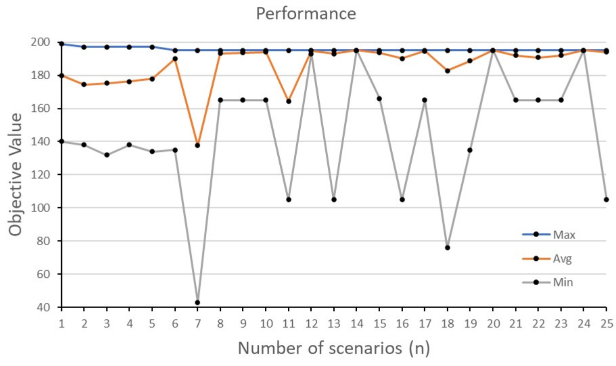

Figure 4 Best-case, average-case, and worst-case performance of the WCM solution across the n samples without infeasible solutions.

In this problem, the final selected solution comes from solving the WCM problem with $n = 24$. It produces the highest feasibility, i.e., 95.4%, with an acceptable performance (i.e., objective function value) of 195. Though the highest performance among all the solutions is 199, this solution only achieves feasibility of less than 30%. Thus, our final selected robust solution is chosen to give the best trade-off between feasibility and performance overall.

Suppose we have chosen n to be 15 instead of 25. The solution that provides the best trade-off between feasibility and performance will come from solving the ACM problem with $n = 13$. This solution is, in fact, poorer than our selected WCM solution at $n = 24$, highlighting the importance of extending the number of samples (n) if computing resources allow, to reduce the likelihood of missing out on good solutions.

We next perform a comparison between the robust and deterministic solutions to highlight the advantages of the former. First, we exhaustively compute the optimal solution for each of the N scenarios. These solutions represent an ideal world where we are able to solve our problem *just-in-time* when the uncertainty unfolds. However, in reality, our decisions can only be made *here-and-now* before uncertainty unfolds. We shall illustrate that the goodness of our robust *here-and-now* solutions just slightly falls short of the *just-in-time* ones.

For the comparisons, we shall be using the deterministic solution and two other robust solutions; our selected robust solution (WCM solution at $n = 24$), and an intermediate robust solution (WCM solution at $n = 18$). Table 1 below shows the comparisons between the deterministic solution and the two robust solutions when using the *just-in-time* solutions as a benchmark.

From Table 1, we observe that the deterministic solution has a high probability (99.3%) of infeasibility when uncertainty is considered. The intermediate robust solution has the next lowest feasibility and the selected

[1] Objective values of infeasible instances were set to zero during performance computation.

robust solution has the highest feasibility, as expected.

	Deterministic	Robust (n = 18)	Robust (n = 24)
% Infeasible	99.3%	45.6%	4.6%
% Underperform	99.7%	100%	100%
Max deviation from optimal	28.93%	61.42%	2.01%
Min deviation from optimal	0%	1.02%	1.02%
Avg deviation from optimal	8.85%	7.56%	1.42%

Table 1 Comparison between the deterministic solution and the robust solutions against the *just-in-time* benchmark.

To examine the optimality of the deterministic and robust solutions as compared to the *just-in-time* solutions, we have excluded the infeasible instances from the analysis, otherwise, the presence of infeasible instances for all 3 solutions will trivialize certain comparisons. From Table 1, it appears that the robust solutions underperform all $(N = 1000)$ *just-in-time* solutions, while the deterministic solution manages to match the optimality of the *just-in-time* solutions in 0.3% of the scenarios. This observation reinforces the trade-off between feasibility and performance in the robust optimization method.

Although our selected robust solution underperforms all the *just-in-time* solutions, the margin of underperformance was not significant. The largest margin was 2.01%, while in some scenarios the value was only 1.02%. On average, the margin of underperformance was only 1.42%. When looking at the intermediate robust solution, the margin of underperformance ranges between 1.02% – 61.42%, with an average of 7.56%.

While the deterministic solution has a few scenarios (0.3%) where its optimality matches that of the *just-in-time* solutions, in some other scenarios it actually underperforms by as much as 28.93%. On average, the underperformance margin is at 8.85%. Incidentally, the deterministic solution has the highest (average) margin of underperformance compared to the two robust solutions. With also a high probability of infeasibility, we can clearly see the benefits of the robust optimization approach when dealing with uncertainty in our problem.

Conclusions

Based on the findings in the previous section, we can conclude that adopting the robust optimization approach is preferred for this particular problem as it provides a high probability of feasibility with some trade-off in performance. Solving the *worst-case maximization* problem with $n = 24$ samples yields a better result for the considered test case. However, we also observe that solving the *average-case maximization* problem produces relatively more stable results in terms of feasibility and performance. Moreover, depending on the choice of n, solving the *average-case maximization* problem can sometimes produce a better result. Therefore, it is recommended to solve both problems if computing resources permit.

We also acknowledge certain drawbacks of the robust optimization approach. Since the duplication of the problem structure is required, it can be computationally expensive especially when the problem becomes more complex. For this problem, with our choice of sample size $n = 25$, the time taken to solve each of the problems (ACM and WCM) is approximately 12 hours. We can expect this timing to increase exponentially as the size and complexity of the problem increases. Nevertheless, given the nature of the problem where there is a low tolerance for infeasibility, it is worthwhile to invest the time and effort to generate a robust solution, as it often requires a much larger time and effort to deal with the infeasibility downstream in the real world.

It is worth mentioning that the approach described here could be employed in other applications where planning under uncertainty is involved, for example, parcel shipping and freight carried by trucks, air, trains, and ports.

List of References

- Ben-Tal, Aharon, Laurent El Ghaoui, and Arkadi Nemirovski. (2009). *Robust optimization*. Vol. 28. Princeton university press.
- IBM ILOG. (2020). *IBM ILOG CPLEX 12.6 User's Manual.*
- Spall, James C. (2005). *Introduction to stochastic search and optimization: estimation, simulation, and control*. John Wiley & Sons.

Acknowledgment

This work is an outcome of the Future Resilient Systems project at the Singapore-ETH Centre (SEC) supported by the National Research Foundation, Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme. The authors acknowledge the support by the Defence Science and Technology Agency (DSTA) for their active participation in the implementation and experimentation of the Robust Optimisation methodology.

Contact

Future Resilient Systems
Singapore-ETH Centre
1 Create Way
CREATE Tower, #06–01
Singapore 138602

info@frs.ethz.ch
www.frs.ethz.ch