SAND2018-10618R

SAND20XX-XXXXR
LDRD PROJECT NUMBER: 210602
LDRD PROJECT TITLE: *Chance-Constrained Optimization for Critical Infrastructure Protection*

PROJECT TEAM MEMBERS:
PI: Bismark Singh, Jean-Paul Watson
PM: Jennifer Troup

ABSTRACT:

Stochastic optimization deals with making highly reliable decisions under uncertainty. Chance constraints are a crucial tool of stochastic optimization to develop mathematical optimization models; they form the backbone of many important national security data science applications. These include critical infrastructure resiliency, cyber security, power system operations, and disaster relief management. However, existing algorithms to solve chance-constrained optimization models are severely limited by problem size and structure. In this investigative study, we (i) develop new algorithms to approximate chance-constrained optimization models, (ii) demonstrate the application of chance-constraints to a national security problem, and (iii) investigate related stochastic optimization problems. We believe our work will pave way for new research is stochastic optimization as well as secure national infrastructures against unforeseen attacks.

INTRODUCTION:

Stochastic optimization is the mathematical science of making reliable decisions under uncertainty. It plays a crucial role in many national security data science applications where decisions need to be made with limited knowledge of the future. Examples include: vulnerability of critical infrastructures [Salmeron, 2005], wind powered unit commitment [Anitescu, 2011], and disaster management [Zabinsky, 2010]. Similar optimization models are also vital to numerous private sector enterprises, such as portfolio selection [Shapiro, 2009], renewable resource management [Pickens, 1991], aircraft sustainability planning [Woodruff, 2010], water resource management [Huang, 2009].

Chance-constrained (CC) models are a subclass of stochastic optimization models where highly-reliable decisions need to be made, while acknowledging that we cannot safeguard ourselves against the entire set of possibilities of the future. Defending against all possibilities is too expensive, and hence a compromise is traded for. CC models answer the following question: how should we make decisions today to defend against *most* of tomorrow's unseen realizations? Thus, CC models serve a risky proposition---safeguarding against a future where high-risk low-likelihood failures could be catastrophic. As an example, a first-stage decision is hardening of certain components of the electrical grid; then, a terrorist strikes, and some components are

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

damaged. The hardening decision must be such that the crucial grid components are secured with a very high probability. Similar problems have been investigated under lightning strikes as well. In most cases, a second-stage decision, known as *recourse*, is made after the uncertainty is realized. In the absence of recourse, recent research has shown that CC models can be solved tractably [Luedtke, 2008]. However, most practical problems have recourse and then CC problems remain intractable even for a modest number of future possibilities. A set of scenarios defines the uncertain possibilities for the future. Typically, the uncertainty is approximated with a finite set of scenarios to achieve computational tractability. Our work has also examined how to construct these scenarios, which are required prior to the optimization.

In this short study, we achieved three aims that we summarize below.

1. First, we developed new algorithms to bound chance-constrained optimization models using classical probability. In this theoretical work, we investigated the strength of these bounds under two different scenario structures (constructed from two sampling schemes). We observed that a larger correlation between the uncertainties results in a more computationally challenging optimization model. This work is currently under second peer review for potential publication in *Optimization Letters*.
2. Second, we developed a mathematical formulation for a stochastic unit commitment model with renewable penetration including chance constraints. Stochastic unit commitment is extensively studied in the optimization community; however CC models haven't been thoroughly investigated. Our model allows thermal generators to operate outside of their prescribed ratings under emergencies (from planned or unplanned outages). This work is currently under review for potential publication in *Computational Management Science*.
3. Third, the insights developed from the investigative work have been useful in our existing projects. We continue our work in stochastic optimization methodologies and applications. One work is currently under peer review with the journal *Networks*, and another is available online at arXiv.

MATHEMATICAL METHODS AND RESULTS:

As mentioned before, in this study we achieved three goals. In this section, we describe the mathematical methods for each of these three aims; for further details, please refer to our SAND reports that correspond to articles that are currently in peer review.

1. **SAND number: 842501**
   Our first task is a theoretical approximation of chance constrained optimization models. To approximate an optimization model, we need to develop both an upper bound and a lower bound for it. The strength of the approximation is given by how close are these two bounds to each other. Below we describe the mathematical details of our modeling.

To this end, we consider a joint-chance constraint (JCC) as a union of events, and approximate this union using bounds from classical probability theory, see Fig 1. When these bounds are used in an optimization model constrained by the JCC, we obtain corresponding upper and lower bounds on the optimal objective function value. In part, we are motivated by existing work [Pinter 1989] where individual chance constraints are approximated using classical inequalities, such as Chebyshev's inequality, Bernstein's inequality and Hoefding's inequality. Our work also approximates chance constraints using classical inequalities, but there are at least two major differences. First, the work in [Pinter 1989] considers only individual chance constraints while we study joint chance constraints. As a result, Chebyshev's inequality, Bernstein's inequality and Hoefding's inequality cannot be directly applied. Second, the work in [Pinter 1989] considers weighted sums of first-stage decision variables and is a stochastic program without recourse. We consider a general two- stage stochastic program with recourse and second-stage costs. In contrast to the existing work, we derive upper and lower deterministic bounds on chance-constrained models rather than the exact deterministic equivalents of the chance constraint itself.

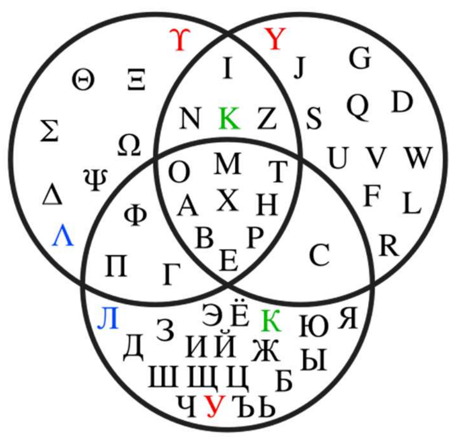

Figure 1 A joint chance constraint is a union of events

We can envision a JCC as a union of “failures”, at different times, to be less than a threshold. If we approximate this union using a quantity larger than the union, then the feasible region representing an optimization model with the JCC is restricted, and thus the approximate model gives a lower bound under maximization. Conversely, if we approximate the union using a quantity smaller than the union, then the feasible region representing an optimization model with the JCC is enlarged, and thus the approximate model gives an upper bound under maximization. Motivated by this fact, we study six well-known approximations of the union. We investigate the computational strength of these approximations, pitted against each other using two different sampling schemes.

Built on
LDRD
Laboratory Directed Research and Development

2. **SAND number: 865314**

   Our second task is a practical demonstration of how chance constrained models are applicable to national security and securing critical infrastructure. As an example, we choose a unit commitment model-how to schedule generators and at what levels to meet power demand while minimizing cost.

   To this end, we develop a novel chance-constrained unit commitment formulation under renewables production uncertainty. Stochastic unit commitment is the mathematical technique for determining which generators should be on, and at what levels, to meet a certain demand. The uncertainty can arise from a stochastic demand, or from uncertain production levels. Optimization methods are frequently employed for the solution of unit commitment models; however our work is novel in the following sense.

   For most thermal generators, their underlying technical constraints are often based on engineering judgements, and system operators in the US operate generators outside these limits in emergency situations. To the best of our knowledge, previous work has ignored this important practical consideration. We develop a CC formulation that incorporates these emergency requirements in a stochastic unit commitment model. Specifically, we allow generators to run outside of their rated limits for a small fraction of times. We demonstrate the tradeoff between costs saved with the percent of allowed emergencies. We also show that the problem is quite tractable for a medium-sized system. We demonstrate how billions of dollars can be potentially saved by choosing a more flexible modeling choice. Our application also serves to demonstrate an example of how chance constraints can be used to model a critical infrastructure problem.

3. **SAND number: SAND2018-10030 J and 842541**

   Third, the results gleamed from this exercise have also been helpful in our investigation of other topics in stochastic optimization that are closely tied to chance constraints. Previously, we were working on minimum *s-t* cut problem. This is a two-stage stochastic optimization model, which is computationally hard to solve. This work is improved as a result. And, we also developed a method to forecast hourly scenarios for solar power using autoregressive models. These scenarios can be used in both CC and two-stage stochastic optimization models. We intend to use these developed solar power scenarios in future stochastic optimization and/or chance-constrained optimization models.

The following articles resulted directly or indirectly from this project:

1. Bismark Singh, Jean-Paul Watson. "*Approximating Two-Stage Chance-Constrained Programs with Classical Probability Bounds*". Under second revision (Optimization Letters). SAND number: 842501

2. Bismark Singh, Bernard Knueven, Jean-Paul Watson. "*A Chance Constrained Model for Stochastic Unit Commitment*". In review. SAND number: 865314
3. Bismark Singh, Steffen Rebennack, and Oleg Prokopyev. "*Two-Stage Stochastic Minimum s-t Cut Problems: Formulations and Complexity*". Under revision (Networks). SAND number: 842541
4. Bismark Singh, David Pozo. "*A Guide to Solar Power Forecasting using ARMA Models*". Available on ArXiv. SAND number: SAND2018-10030 J

DISCUSSIONS AND CONCLUSIONS:

Stochastic optimization is the mathematical science of making decisions under uncertainty. It finds vital applications in national security applications, where an uncertain demand needs to be met with an uncertain supply. Examples include: allocation of electricity with increasing penetration of renewables into the energy grid, installation of sensors to detect cross-border nuclear smuggling, distribution of critical medical resources during a biological disaster. These seemingly unconnected examples are crucially connected with the fact that allocation decisions need to be met before we know about tomorrow's uncertainty. Stochastic optimization offers a natural mathematical ploy to model such challenges mathematically. It answers the following question: how should we make trustworthy decisions today without knowing tomorrow's unseen realizations?

Despite being perhaps the most natural way of mathematically modeling resiliency problems, stochastic optimization has seen limited attention in problems of national security. A number of reasons have contributed to this unfortunate fact, some of which we highlight below. Through this short project, we have attempted to highlight the importance of stochastic optimization techniques for modeling infrastructure resiliency. Although we have not been successful in creating a generic and tractable solution method for such methods, we believe we have created a positive dent in this direction. Below we highlight some applications where stochastic optimization finds a critical role, as well as reasons such models have seen limited attention. These examples serve to highlight the crucial role chance-constrained models play in national security examples

As a first national security example, we commit to a reliable power supply for tomorrow without knowing the exact availability of renewable sources (such as solar and wind). First, we make an energy promise decision for tomorrow, later we observe the availability of renewables. The promise decision must be such that even under very low renewable availability, we provide a reliable dispatch. Here a subclass of stochastic optimization models (known as chance-constrained programming) is used which assures a reliable power supply with a very high probability. In this study, we focused primarily on chance-constrained programming, and two of our works are under peer review with optimization journals, see SAND reports: 842501 and 865314.

As a second national security example, we install detectors across borders to prevent smuggling of nuclear materials. First, we make a decision to install a, expensive and possibly unmovable, detector at a particular location. Then, a smuggler chooses his path to travel across the border. The install detector must be such that we maximize our chances of catching the smuggler. Here a subclass of stochastic optimization models (known as robust optimization) is used which maximizes the chances of the smuggler being detected irrespective of his actual path.

As a third national security example, we distribute antivirals during a pandemic without knowing the number of people who would fall sick at different times and locations. Allocated antivirals cannot be returned, and thus an excess allocation results in wastage while a shortage results in possible fatalities. The allocation decision must be robust to account for varying levels of the severity of a pandemic. Here a subclass of stochastic optimization models (known as two-stage stochastic programs) is used which maximizes the likelihood of the antivirals being allocated to the right locations at the right time. Extensions of our project led to new results in two-stage stochastic programming, see SAND report number: 842541.

To summarize, chance-constrained optimization is a relatively new mathematical technique for modeling reliability problems, and the applications to critical infrastructure are even newer. The practical reason has been the computational intractability of the general CC optimization formulation. In our project, we demonstrate an application of chance constraints to a stochastic unit commitment model see SAND report number 865314. We demonstrate the computational tractability of a medium-sized problem, and also demonstrate the potential savings with a CC formulation as opposed to a traditional formulation. The theoretical reason for the difficulty of solving CC optimization models is that in general they are NP-hard. In a byproduct of our project, we investigate a two-stage stochastic program and prove the corresponding model is NP-hard, see SAND report number: 842541.

ANTICIPATED OUTCOMES AND IMPACTS:

Current state of the art algorithms to solve general stochastic programs are limited by the size of the problem. Modelers rely on sampling a finite set of scenarios (realizations) of the uncertainty to achieve a tractable formulation. Hence, current research has relied on finding effective ways to approximate stochastic programs such that the approximation is as close as possible to the original formulation -one direction of our project has also focused on this aspect. Consequently, having scenarios which accurately represent the unknown uncertainty is also crucial. A byproduct of our study was a scheme to efficiently sample solar power scenarios, addressed via SAND report number: SAND2018-10030 J.

Built on

LDRD

Laboratory Directed Research and Development

Several of our existing projects in 1464, have focused on the ultimate aim of developing new methods to solving CC models in a tractable manner, while not compromising on the set of scenarios considered. This is a challenging task, as the CC optimization is a known NP-hard problem and is also computationally intractable.

In this project, we took steps towards this direction. We completed a theoretical work to develop approximations to CC optimization, and an applied work to demonstrate the functionality of CC optimization to infrastructure security. We seek to continue work in this general spirit. We are also motivated by work we have completed before this exercise. As part of his PhD dissertation at The University of Texas at Austin, Bismark Singh developed a novel algorithm to heuristically achieve high quality solutions for CC problems and demonstrated their viability on a hybrid wind-diesel generator. In future work, we aim to extend this algorithm's capability to electrical grid sized problems and embed it within Pyomo, the open-source optimization language developed at Sandia by Jean-Paul Watson. Current state-of-the-art CC models can only solve models up to a few thousand scenarios, while we seek to extend it to tens of thousands. Successful completion of this work would open new avenues of academic research, as well as provide planners oversight in designing reliable infrastructure.

We believe our work would have a broad impact on research and development related to power grid planning and resiliency, in several projects funded by agencies ranging from DOE/OE, DOE/EERE, DOE/FE, and DOE/ARPA-E. Extensions of our work could be used to rigorously secure electrical grids, natural gas networks, and transportation infrastructure. Finally, by leveraging Sandia's high performance computing (HPC) facilities, the project will continue to establish Sandia's leadership in the HPC community.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.