A Real-World Network Modeling Project

Nedialko B. Dimitrov, Gerald Brown, Matthew Carlyle
Operations Research Department
Naval Postgraduate School
Monterey, CA, 93943
(ndimitro, ggbrown, mcarlyle@nps.edu)

February 12, 2014

Abstract

The Operations Research Department at the Naval Postgraduate School educates experienced junior military officers in state-of-the-art operations research methods. As part of the educational program all officers go through a standard graduate course on network modeling. Within that course, students complete a class-long network modeling project of a real-world infrastructure system that solidifies their understanding of operations research concepts and moves the learning experience beyond the classroom. Through the project, students abstract their real-world problems into mathematics, repeatedly evaluate the connection between the mathematics and reality, and reason about the model results. The project has educational, research, and practical benefits. On occasion, students make discoveries of such significance that we have contacted the infrastructure system operators with briefings on the student analysis results. Some student projects have eventually influenced Department of Homeland Security and Department of Defense policy. We detail what makes the network modeling project work, and how to implement it in other universities.

1 Introduction

The Operations Research Department at the Naval Postgraduate School (NPS) educates experienced junior military officers, drawn from militaries world-wide, who typically have about ten years of operational experience. Every student in the Operations Analysis Master's of Science degree program is required to take a class on network optimization; in that class, every student completes a class-long, real-world network modeling project. For perspective, the class meets for an hour, four days a week, for 11 weeks, with an enrollment of about 20 graduate military officer students of varied backgrounds, with a quarter of these from international allied countries.

The self-guided, open-ended network modeling project is essential for the students to deeply understand the material they learn in class. The basic steps of the project are: 1) Select a real-world application, 2) Create an abstraction of that application as a network, 3) Ask interesting real-world questions, 4) Perform mathematical analysis to answer those questions, and 5) Interpret the mathematical results for a variety of real-world decision makers. As the students perform these steps, they solidify the concepts taught in every graduate-level network modeling class: network definitions, the meaning of network optimization problems, algorithms for solving network optimization problems, and modeling. This is especially appealing to students with prior real-world, operational experience, such as those at NPS, because it shows them that the concepts they learn in class are not simply abstractions, but can be used to address key real-world questions.

The network modeling project employed at NPS adds palpable benefits for the NPS students, and it would do so for graduate students in mathematical network optimization classes in other universities as well, regardless whether the ultimate career goals of students are academic, industrial, public policy, etc. The goals of the project are to have students experience: 1) Creative, problem-specific, selective abstraction of a real-world problem into mathematics, 2) Critical evaluation of the connection between mathematics and reality, 3) Application of network optimization to a realistic-size problem, and 4) Critical interpretation of the meaning of model results. In achieving these goals, the network modeling project takes students beyond the classroom, to recognize the value of the material they learn. It also often creates key connections with the personal experience of the students that could not be made otherwise, through mere lectures and exams.

In this document, we describe how to implement a similar network modeling project in a graduate-level network optimization class. This project can also be used as a capstone project, at the end of a sequence of optimization courses, even though it is executable simply within a network optimization course (Armacost and Lowe, 2004).

2 Project Context and History

Perhaps the biggest difference from a network optimization class one would take at another university and the one at NPS is selective emphasis on some topics. One topic that receives special emphasis in the NPS network modeling class is network interdiction (Brown et al., 2009; Israeli and Wood, 2002; Pan et al., 2003). Network interdiction models are used to infer the vulnerabilities inherent in an operational network system. Most often, network interdiction models are used in a defensive analysis to infer the vulnerabilities of our own operational networks. The NPS network modeling project began over a decade ago as a *Red Team* modeling project. A defensive analysis, where the defenders temporarily play the part of an intelligent attacker, is often called a *Red Team* exercise in military contexts – with the attackers designated as the Red Team, and the defenders as the Blue Team. The NPS network modeling project began as an application of network interdiction models for such a Red Team exercise. Such an exercise helps our military officer students especially because their jobs after graduating involve creating and evaluating plans for attack on or defense of infrastructure, where the goal may be to achieve some given level of damage, but no more (this is called effects-based targeting), or to create some given level of resilience given an estimate of enemy capabilities (this is called assessment-based planning).

Since its initiation, over 500 students at NPS have participated in a project, and produced approximately 200 analyses of a variety of real-world networks and problems. The project has evolved from a Red Team project, to a project that considers both defensive analysis and network design. The instructors collect the project analysis results, which are often later developed into a full Master's thesis, not necessarily by the same student. For example, author Brown first introduced interdiction models to the core linear programming course, deciding later that interdiction modeling integrated better with the networks course. The first homework from that course ended up being the topic of a Master's thesis, and then developed into a full decision support system for missile defense (Brown et al., 2005). The project also provides an opportunity to begin to formulate a problem that is otherwise of research interest. Faculty in the Operations Research Department, even those who do not teach the network optimization class themselves, often pitch project ideas that can be selected by students. The results of the student analysis can then provide a basis for an alternate formulation of a problem, or even an initial formulation. Throughout the next section, we provide multiple specific examples of project topics analyzed by students and the way in which those analyses are helpful outside of the class context.

3 Project Implementation

The NPS network modeling project is a class-long, small-group project that lasts approximately 11 weeks. The rough schedule for the project is:

Weeks 1-2: Group formation and problem selection;
Weeks 3-4: Network abstraction and initial data development;
Weeks 5-7: Operational modeling of the network;
Weeks 6-8: Example project, done as homework. This step is partially concurrent with the previous step;
Weeks 8-10: Intensive final model development, analysis, and peer review; and
Weeks 11-12: Final presentations and peer reviews.

We describe the process in each part of the schedule in turn, along with advice and examples on implementation.

3.1 Group Formation and Problem Selection

On the first day of class, students are told about the project, which accounts for a third of their class grade. As can be expected, giving substantial weight to the project increases active student participation in the project tasks.

Additionally, on the first day of class, students are asked to pair up to conduct the project. While teams of one student and three students are allowed, students are advised that the optimal team size is two. The reasoning is that in a team of two there can be an active discussion on modeling, and each person has to contribute and understand the project as a whole. In larger teams, like those of three or more, there is a higher likelihood of task division and students not understanding all aspects of the project. A team of one is also suboptimal because there is no partner with whom to critically question modeling ideas. While teams of one are allowed, they are discouraged by the fact that the expectations for a team of one are exactly the same as those for a team of two. A final consideration when forming a team is maximizing diversity in specialty, background, and nationality.

In addition to forming a team, the students have to select a project topic within the first week of class. Surprisingly, the students generally do an excellent job selecting good network problems. An instructor can help with problem selection in the following ways:

- In the first day of class, give an overview of some network modeling applications. In that overview, include some examples of past projects or possible project topics. This is an excellent place to pitch project ideas that are otherwise of research interest.
- Ask students to select problems with plenty of publicly available data–classified or otherwise non-public data is not permitted. The students should have intuitive, English-language answers to the following questions: What are the nodes of the network? What are the edges? What governs operation of the network? In addition, the students should be able to provide web links to publicly available data on the network operation.
- Encourage students to select topics that are personally motivating, currently relevant, involve networks with interesting structure, and are different from topics of other teams.
- Have students submit proposals to the instructor in the first week. In this way, the instructor can turn down proposals that are not really amenable to network analysis, are too similar to suggestions from other student groups, or do not have sufficient data for analysis.

From the perspective of an instructor, one of the greatest worries may be whether the students are able to select a project topic so early in the course. At this point in their curriculum, our students have only completed a linear programming class, which may include some network modeling content. However, students have an intuitive understanding of what is a network and of goals, and possible operations on a network. Combining that intuitive understanding with the examples provided by the instructor typically results in excellent project selection. In the past, at NPS, we have had to turn down only approximately 10% of initial proposals. The students often surprise the instructors with interesting, relevant network problems that the instructor has not considered previously. The quick due date for the initial project proposal is so that groups whose proposals are turned down can refine and finalize a topic in the second week. Recent project topics, influenced by personal student experiences and interests, include: Alabama tornado relief, Amazon.com warehouse positioning, California high-speed rail development, Chesapeake Bay transport, convoy security, cross-border drug movements, cross-border tunnels, Enron communication analysis, humanitarian aid warehouses, Israeli traffic analysis, London Olympic Games transport, medical evacuation, Monterey fire department responsiveness, Monterey Peninsula evacuation, natural gas pipeline transport, rail transport of new cars, ski lifts for Keystone Colorado, small boat attack, truck-based gasoline transport, US oil import security, US potato transport, veteran health network, and Yosemite trails. Many of these are civilian infrastructure systems, and not necessarily those identified by our Department of Homeland Security as critical to national defense. Three-quarters of all our U.S. critical infrastructure systems are owned and operated by non-governmental entities (GAO, 2006). These simple examples, all selected by students, exhibit all the richness of larger, more-critical infrastructure systems.

3.2 Network Abstraction and Initial Data Development

In the third and fourth weeks of class, the students are asked to produce two data files: one for the nodes, and one for the edges of the network. By this time in the course material, students have had some experience with network definitions and know reasonably well what is a node and what kinds of edges networks can have. The exercise of producing data files forces them to begin abstracting their real-world problems into mathematics.

The major stumbling block for the students in this part of the project is the necessity to choose how much reality they should include in their model. For example, they have to begin thinking about what exactly they should make into a node; if their network operation uses roads, they could choose each intersection to be a node for a neighborhood-level model, or should just major street intersections be nodes for a city-level model, or perhaps just highway intersections for a state-level model? Different student teams often take very different approaches to what level of model fidelity they seek. But, each team learns that as they create their network, they must necessarily leave parts of reality outside their model.

A second stumbling block for students is beginning to identify the node data and edge data that is relevant for their problem. Besides simply listing the network nodes and edges, the data files the students create should have relevant data for the network operation. Students begin to critically look at the data sources they identified in weeks one and two, to try to distill the publicly available data into useful modeling input. Often, key pieces of operational data on capacities of edges, or demands for traffic are missing and students have to learn how to create reasonable estimates from alternate publicly available data sets.

Finally, some fraction of the students typically select more abstract network problems and have difficulty formulating the problem as a network. These groups usually require a bit of instructor advice to help them refine and clarify the ideas for their project. Typical examples of this include applications where the network is created simply by gridding a geographic region, as opposed to an obvious road network. Other examples include time-layered networks, for student groups who are set on modeling sequences of events, and full-empty layered networks for student groups modeling applications involving the pickup and delivery of material. A little instructor input to introduce the students to these modeling approaches helps.

From the instructor's perspective, the major points to make to the students at this stage of the project are the following. First, the student groups should create a background narrative for their problem. This helps students ask relevant questions and formulate the problem based on intuition of the narrative. Second, the student groups should be encouraged to start with simple models.

Students can be somewhat disappointed because their initial models do not include all aspects of narrative reality they think are important. It is useful for an instructor to explain that starting with a simple model, and making it more complex over time leads to better understanding of the ultimate results. Finally, the students should understand that their decisions are not set in stone. For example, even though they produce some data files in weeks three and four, essentially all groups edit those data files significantly for their final analysis. At this point of the project, it is simply important to start thinking and working on the modeling abstraction.

3.3 Operational Modeling of the Network

In weeks five through seven, the main goal of the students is to create an optimization model of the operation of the chosen network. Around this point in the course material, the students should receive an overview of several kinds of network models: shortest paths, maximum flows, minimum cost flows, multi-commodity flows, stochastic models, and perhaps multi-objective models. The students can then think about their particular application and decide on the best method of modeling the operation of their network. The student deliverables for this part of the project are 1) a linear program model of the network operation, most importantly the meaning of the linear objective function, and 2) the output of an optimization run of their model on the data they have previously created.

For this part of the project, a common problem is that students tend to select the models that are most familiar. For example, if shortest paths are what have been discussed in most detail in the course, many students tend to model everything with a shortest path objective. An instructor can help here by simply discussing what kind of model is appropriate for what kind of situation. Given a few simple example applications, essentially all students select a relevant optimization model for their chosen problem.

A second common problem for this part of the project is the actual implementation of the optimization model. Students have varying degrees of programming capability. Because of that, an instructor can expect some students to struggle with implementing their optimization model. A simple way to help is for the instructor to post some template code. Template code solves some tiny, unrealistic problem; but, at the same time, the students can edit that template to make it fit their problem of interest. Given starting-point templates, essentially all students are able to implement and run their initial optimization models.

3.4 Example Project

In the sixth through eighth weeks of the course, the students do a sequence of homework assignments that serve as an example project. In this sequence of assignments, the students follow a series of step-by-step questions to create a simple model initially, understand its outputs, and adjust the model to make the outputs more realistic. Doing the example project solidifies student understanding of what is expected from them in their final analysis, and gives them practice with the critical implementation skills they require to be able to complete their analysis.

Currently, we use a fictional escaping drug lord problem as an example project. The narrative background of the example project is that the police have raided the offices of a drug lord, and now he is on the run. It is the job of the students to help the Arizona state police capture the escaping drug lord. The network data for the project consists of a city-and-highway level road network of Arizona, and is provided by the instructor. The basic operational model of the network is–from the drug lord's perspective–a shortest path model to leave the state as quickly as possible. The Arizona state police need help to determine on which roads to place road blocks to stop or maximally delay the escaping drug lord. The final aspect of the narrative is that the drug lord has informants who tell him the positioning of the road blocks, so he can avoid them when they are placed. The narrative creates a shortest path network interdiction problem and introduces a key aspect of the network modeling project–an integer linear program on top of a basic network model to design or interdict the network.

The sequence of homework problems is graduated, with increasing levels of model complication. The first few questions of the homework assignment ask the students to compute the optimal escape route for the drug lord, and look at it on a map to see if it makes sense. The next questions ask the students to manually place a road-block and compute the optimal, responding adjusted escape path for the drug lord. That is followed by a set of questions requiring the students run a sequence of linear integer programs for road block placement, with increasing numbers of roadblocks. In the sequence of questions, the students have to repeatedly alternate between computing results and interpreting these results. The homework assignments then ask the students to adjust the model so that the drug smuggler is not computing the shortest escape route, but rather the maximum reliability escape route. Road block effectiveness is also adjusted in the questions, from fully- to partially-effective and the students analyze the results. The homework questions end by asking the students to incorporate one new aspect of reality into the model, something of their choosing.

Student feedback on the example project has been positive. Most students say they really begin to understand their network modeling project in this phase in the class. The drug lord homework assignments show students, step-by-step, how to add a mixed-integer linear program on top of a network model. Later in the course, for their final project analysis, they are required to add such an interdiction or design feature to their own operational model.

3.5 Intensive Final Model Development, Analysis, and Peer Review

To emphasize the project, the last few weeks of class have no homework other than the project. Problems that would otherwise be given for homework can be completed during class time instead. This slightly lowers the number of lecture hours, but increases the problem solving abilities of the students. However, to make sure students do not leave everything for the last minute, an instructor can hand out several guiding documents and conduct a mid-development peer-review cycle.

Useful guiding documents include a list of deliverables, a sequence of steps to take for final model development, and a schedule of desired progress for the next few weeks. The deliverables at NPS include a 21-minute in-class presentation, followed by four minutes of questions; the slides for the presentation; a three-page executive summary of the analysis meant for a non-expert decision maker; and all of the data and implementation files required to reproduce the project results. Students have already been exposed to most steps for final model development, such as background narrative creation, data development, and solving an operational model. The major new step introduced is a choice either to do a network interdiction Red Team integer linear program, or to do a network design integer linear program on top of the operational model. Many student groups decide to do both, but the students learn a good deal from analyzing which integer program is more useful for their application. The schedule for the progress of a student includes an early draft of the executive summary to be submitted for peer review a week before final project presentations begin.

The peer review is a relatively new addition to the NPS network modeling project. Every student in the class writes a two-way blind review of two other draft executive summaries. The peer review achieves multiple goals. First, it is often easier to see mistakes in presentation and model formulation in the work of others rather than your own. Second, most of our students become managers of analysts when they finish their education, and critical thinking about presentations by others is a key skill required in an effective manager. Third, students get to receive some feedback on their work before the final presentations, and can adjust their model or analysis.

During these final weeks of class, students complete many diverse and fairly sophisticated models. Often student groups require simple answers to a few questions to help them proceed, and this could easily overwhelm the instructor. Accordingly, operations research faculty at NPS fluent in model development have open office hours and expect students not their own to drop in for counseling and advice.

3.6 Final Presentations and Peer Reviews

The last week or two of class consists of the student group project presentations. The required length of the final in-class presentations encourages students to carry out, and describe, significant analyses. It is not possible for a student group to use empty filler for a 21-minute presentation without it being obvious to the audience that the presentation lacks content. At the same time, the required presentation length requires the instructor to devote some class time to these presentations. The class time is justified because students in the audience receive significant benefit from listening to the diversity of topics selected by other groups, and the model of each group tends to be unique.

An instructor can do several things to make the presentations more effective:

- In the grading guidelines, an instructor can specifically create a score penalty for presentations that are too short. This forces student groups to create analysis content.
- In the grading guidelines, an instructor can stress analysis and creativity, which motivates unique analysis content for each group.
- An instructor can hand out writing and presentation style advice. This shows the students some basic guides they can follow to prepare effective presentations.
- Students can peer review the presentations of groups whose draft executive summaries they reviewed. The peer reviews force the audience to focus on presentations, and leads to critical thinking on presentation styles. Furthermore, it allows the reviewers to see progress in the analysis project of the presenter.
- An instructor can invite outside guests including senior officers to observe the student presentations. This increases pressure on the students to perform well because some of their audience has never attended class with them. Also, outside guests tend to ask unique and interesting questions following the presentations.

Basic criteria for evaluating final projects include the thoughtfulness of analysis, innovation in problem selection and execution, technical complexity, presentation delivery, and executive summary. Perhaps the most important criterion is the thoughtfulness of analysis, which is often evident in the final presentations. Some good questions to ask groups at the end of their presentations include: 1) Would a real-world decision maker really care about the questions you asked? 2) What other questions would a real-world decision maker ask, in addition to the ones you addressed? 3) What are the major interest groups in the problem you analyzed, and what are the English-language take-aways for each of them based on your analysis? Groups with lower quality projects often do not have good answers to these questions. Groups with good answers to these questions have completed an effective cycle of analysis, interpretation, and re-analysis.

4 Reasons that the Network Modeling Project Works

A number of factors make the class-long modeling project possible in a network optimization class, though it may not be possible in other classes. First, most students have a strong intuitive understanding at the beginning of the course of what makes a network, so they can select a project topic quickly. Second, we are able to equip the students with modeling language prototypes for a wide variety of network models. This significantly reduces implementation time for the students, maintains their exposure to important modeling language syntax, and allows them to focus on the actual modeling problem. Third, a small number of basic mathematical network models are applicable to a large variety of real-world problems. This allows for a rich variety of analysis at the end of class, while maintaining mathematical similarity so that the student audience can understand how each group derived its results.

The project is especially effective for students with real-world operational experience, such as those at NPS. However, even students without extensive real-world experience benefit from the connections the project creates between mathematics and analysis questions that matter.

The final presentations are often quite impressive to the student audience because they illustrate the wide variety of applications that can be addressed by the mathematical tools gained in class. On occasion, students discover things about the systems they model of such significance, we have contacted the operators, owners, and/or authorities, offering briefings on the student insights.

The insight and directions of the projects range from military and defense applications, to public sector, to corporate operations. Figures 1 through 3 show key final slides from several in-class presentations. The project of Figure 1 concerns a network interdiction model to quickly detect contamination in the Washington, DC, Metro system. In that project, the students' operational model is a multi-commodity flow accounting for movement of passengers from station to station, as well as additional variables to account for contamination conveyance. The integer programming model on top selects the best place to locate contamination detectors, and the best place for an attacker to introduce contamination.

The project in Figure 2 concerns the Obama administration's plan for a national high-speed rail network. The students' operational model considers moving passengers between large metropolitan centers using rail lines. The integer program on top considers budget-limited construction of new high-speed rail lines. The plan output by the model is remarkably close to the plan suggested by the Obama administration.

Finally, the project in Figure 3 optimizes the fun of a ski route in Big Sky Resort, Montana. The students' model is based on a time-layered network consisting of 96 layers, each accounting for movement between a five minute period. The model includes aspects of reality such as decreasing fun value for taking the same ski run twice; different speeds for different slopes; different slope preferences for advanced, intermediate, and beginner skiers; and the construction of a contiguous run starting at the resort, lifts taken, runs taken, and returning to the resort at the end of the day. The students' operational model is an integer program itself, to make sure they receive a single route as opposed to a randomization between different routes. The students also performed sensitivity analysis, considering how the route would change if certain lifts are not operational, or certain slopes lack snow.

For instructors wishing to implement a similar project in their classes, all of the material for recent offerings of the networks course at NPS is available online (Dimitrov, 2012), including project handouts for each step of the project, the example project, template code, guides for written and oral presentations, and all of the final student presentations and executive summaries.

Although our perspective is military, the network modeling project could be implemented and bring significant educational results in a civilian university. One fundamental question addressed by the project, especially Red Team models, is analysis of damage inflicted intentionally as opposed to random failures in operation. There is much to be gained from stress-testing civilian and commercial infrastructure systems to find key vulnerabilities. Such analysis can be both motivating to and carried out by graduate students without significant military or business backgrounds. In addition to the brief listing of specific project topics above, most of which are civilian in nature, instructors could suggest topics focused on corporate operation. A "just in time" supply chain makes a wonderful military target, but is also worthy of analysis to see where just a little more inventory capacity might lend a lot of resilience. Though most efficient logistics systems employ sole-sourcing, a little analysis of dual- or back-up-sourcing can lend a lot of insight. Some might find rapacious competitors to be quite avid adversaries, and wonder whether and where either party

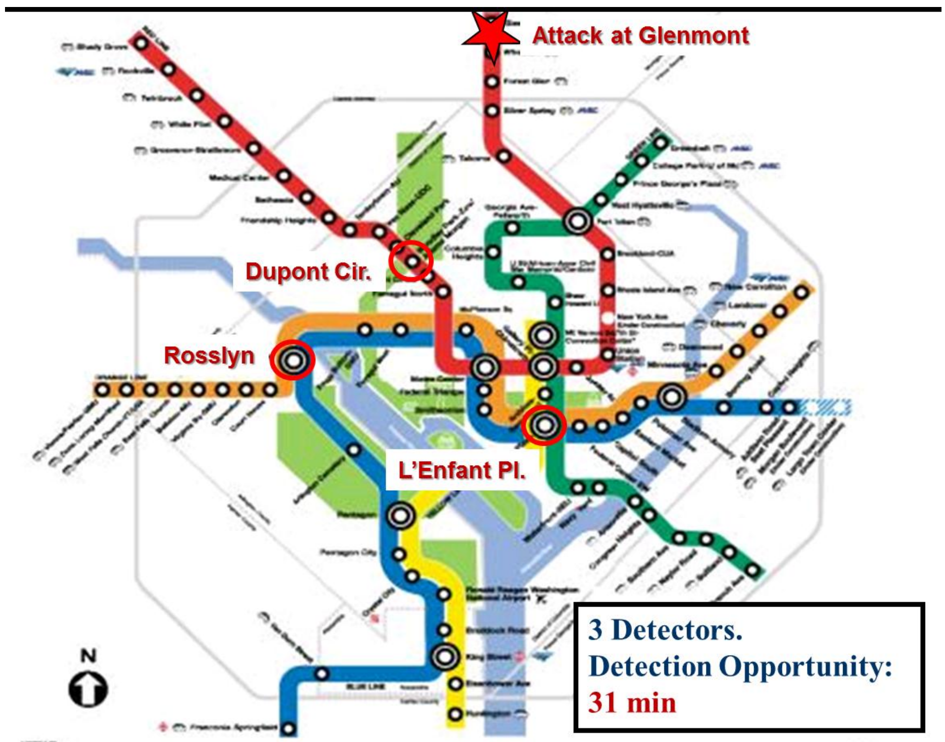

Figure 1: Placement of chem-bio detectors in the Washington, DC, Metro system. The map shows the optimal placement of three chem-bio detectors and the location of an optimal attack that would maximize the time to first detection, given that the locations of these large, expensive, visible detectors are known to the attacker. The underlying multi-commodity network model accounts for train conveyance of the contaminants and cross-contamination of passengers at stations.

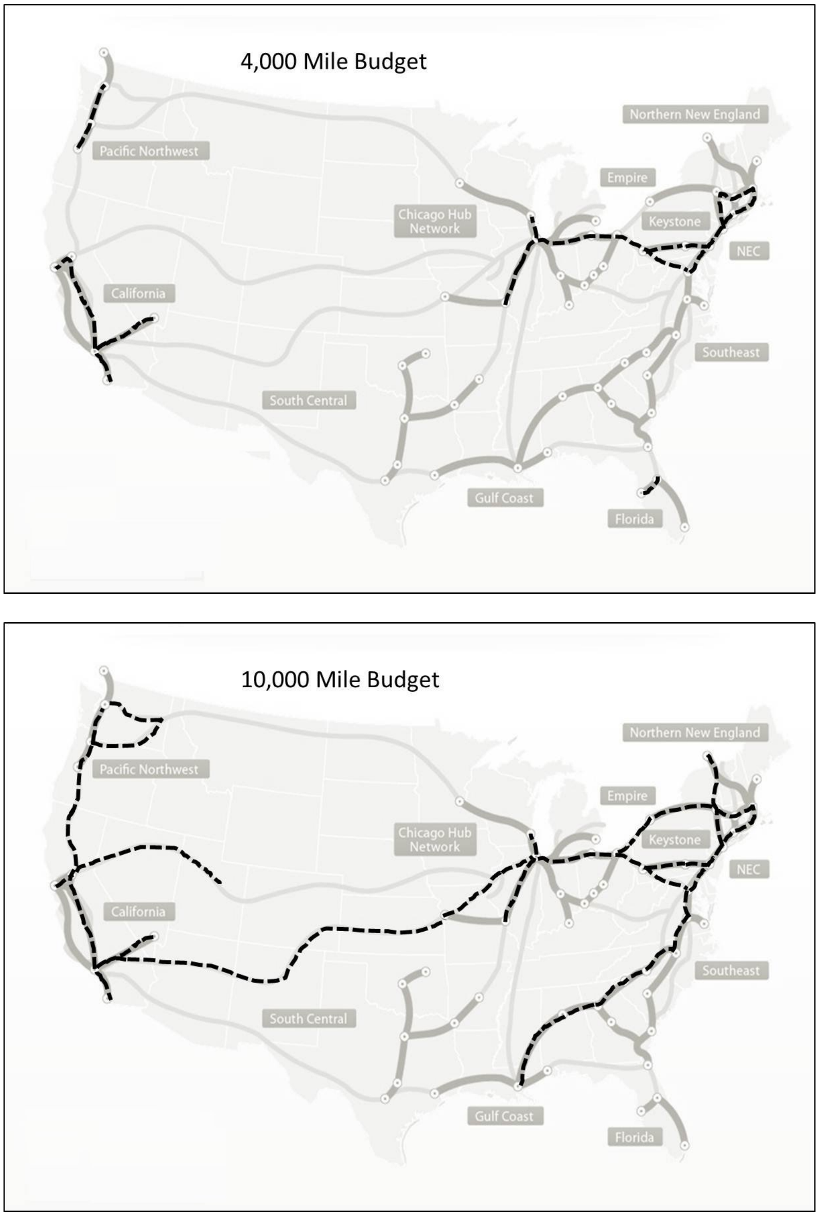

Figure 2: Construction of a high-speed rail system in the United States. The maps show the Obama administration's high-speed rail plan compared to a plan created by the student group's mixed integer programming model. The Obama administration's plan is: thick gray lines designating high speed rail corridors and thin gray lines designating other rail routes. The students' plan is overlayed in dashed black lines. The operational model captures passenger transport between major population centers using either regular transport or high-speed rail. The mixed integer programming model on top captures budget-limited construction of high-speed rail routes. For low values of budget, 4,000 miles, the students' model output agrees remarkably with Obama's proposed plan. For higher budget values, 10,000 miles, there is some disagreement. The students postulated that the disagreement is due to Obama's objective to build political consensus for a high-speed rail that is not captured in their model.

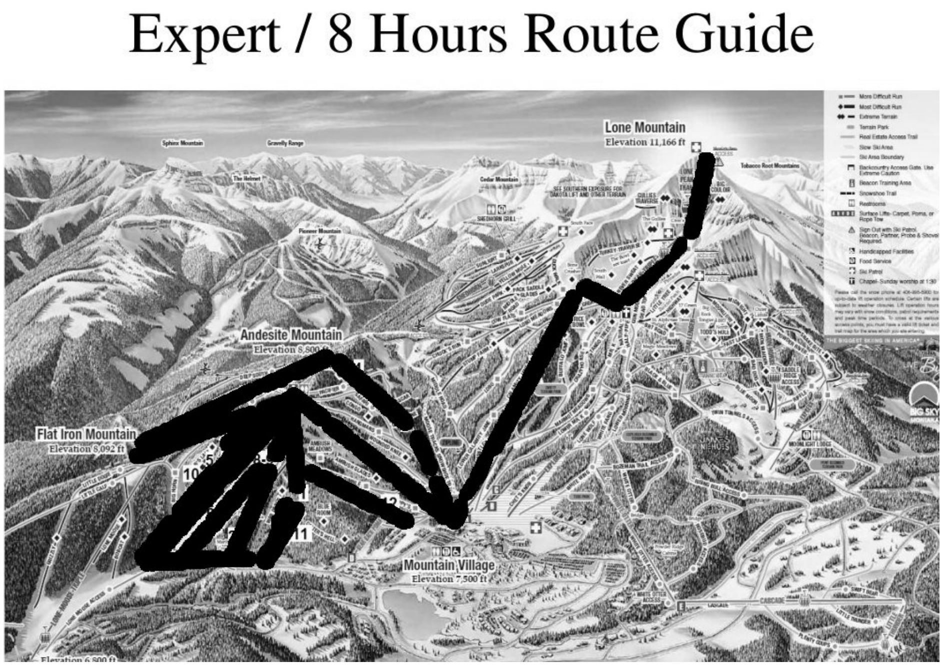

Figure 3: Maximizing fun for skiing in Big Sky, Montana. The map shows an optimized eight hour skiing route for an advanced skier in Big Sky Resort, Montana. The thick black edges represent the runs selected in the route. The students' model includes aspects of reality such as an advanced skier's preference for more difficult runs, and the fact that doing the same run twice results in lower fun the second time around. Further, the route is contiguous, mapping a full eight hours starting from the bottom of the mountain, lifts taken, runs taken, and returning home. Not pictured is additional sensitivity analysis by the students on route variations resulting from broken lifts or lack of snow. In some sense, this project is fun and humorous, but in another it represents a useful marketing tool for a resort operator.

has vulnerabilities. Others might view threats of operational interference from organized labor to be just as worthy of analysis. In addition, businesses may not understand their vulnerability to correlated failures in manufacturing. For example, in 2011 flooding in Thailand crippled hard drive manufactures, sending hard drive prices soaring (Fuller, 2011). The Red Team and design models of the network modeling project can help students in civilian universities understand how to identify and design their way past such vulnerabilities.

Beyond the classroom, these project exercises have led to a number of theses, decision support systems, and research innovations. For instance, we have often had to explain to our students why a prioritized list of attacks or investments does not necessarily lead to the best allocation of resources. We finally realized we needed to formalize this presentation in Alderson et al. (2013). We believe similar research outcomes will result in civilian universities.

We are gratified that our network modeling project has achieved outside recognition. Our Master's-level networks course and its network modeling project have been suggested to influence Department of Homeland Security policy (National Research Council, 2010, pp.70-71). We are also aware from many sources within the Department of Defense that we have already influenced defense policy, and we have published a limited number of these (e.g., Brown et al., 2005).

References

D. L. Alderson, G. Brown, M. Carlyle, and A. Cox. Sometimes there is no 'most-vital' arc: Assessing and improving the operational resilience of systems. *Military Operations Research (to appear)*, 2013.

A. P. Armacost and J. K. Lowe. Operations research capstone course: A project-based process of discovery and application. *INFORMS Transactions on Education*, 3:1–25, 2004.

G. Brown, M. Carlyle, D. Diehl, J. Kline, and K. Wood. A Two-Sided Optimization for Theater Ballistic Missile Defense. *Operations Research*, 53(5):745–763, 2005.

G. G. Brown, W. M. Carlyle, R. Harney, E. Skroch, and R. K. Wood. Interdicting a nuclear-weapons project. *Operations Research*, 57:866–877, 2009.

N. D. Dimitrov. Network Flows and Graphs (Spring 2012). http://neddimitrov.org/teaching/201202NFG.html, June 2012. Accessed on 2012-08-02.

T. Fuller. Thailand flooding cripples hard-drive suppliers. *The New York Times*, November 6, 2011.

United States Government Accountability Office GAO. Critical infrastructure protection: Progress coordinating government and private sector efforts varies by sectors' characteristics. Technical Report GAO-07-39, United States Government Accountability Office (GAO), 2006.

E. Israeli and R. K. Wood. Shortest-path network interdiction. *Networks*, 40:97–111, 2002.

National Research Council. *Review of the Department of Homeland Security's Approach to Risk Analysis*. National Academies Press, Washington, DC, 2010.

F. Pan, W. Charlton, and D. P. Morton. Interdicting smuggled nuclear material. In D.L. Woodruff, editor, *Network Interdiction and Stochastic Integer Programming*, pages 1–20. Kluwer Academic Publishers, Boston, 2003.