REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
05/17/2022	Technical Paper	

4. TITLE AND SUBTITLE

Network Interdiction: A Game Theory Approach

5a. CONTRACT NUMBER

W56KGU-18-D-0004

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Samuel Billingham, Michael Martinez

5d. PROJECT NUMBER

100927.10.100.S240.RA0

5e. TASK NUMBER

W56KGU21F0008

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

The MITRE Corporation
202 Burlington Road
Bedford, MA 01730

8. PERFORMING ORGANIZATION REPORT NUMBER

PRS-22-1141

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

NORAD & USNORTHCOM

10. SPONSOR/MONITOR'S ACRONYM(S)

NORAD & USNORTHCOM

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

We discuss a three-player network interdiction construct that combines optimization and game theory, applied to a use case within the U.S. transportation network. This analysis provides a powerful, robust approach to identify critical infrastructure within the network; simple priority listings may result in overlooked vulnerabilities that are easily exploited.

15. SUBJECT TERMS

Critical Infrastructure Protection; Network Security; Modeling and Simulation; Computer Security; game theory; critical infrastructure; national security; resilience; network interdiction

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE		9	Susan Carpenito
					19b. TELEPHONE NUMBER *(Include area code)* 781-271-7646

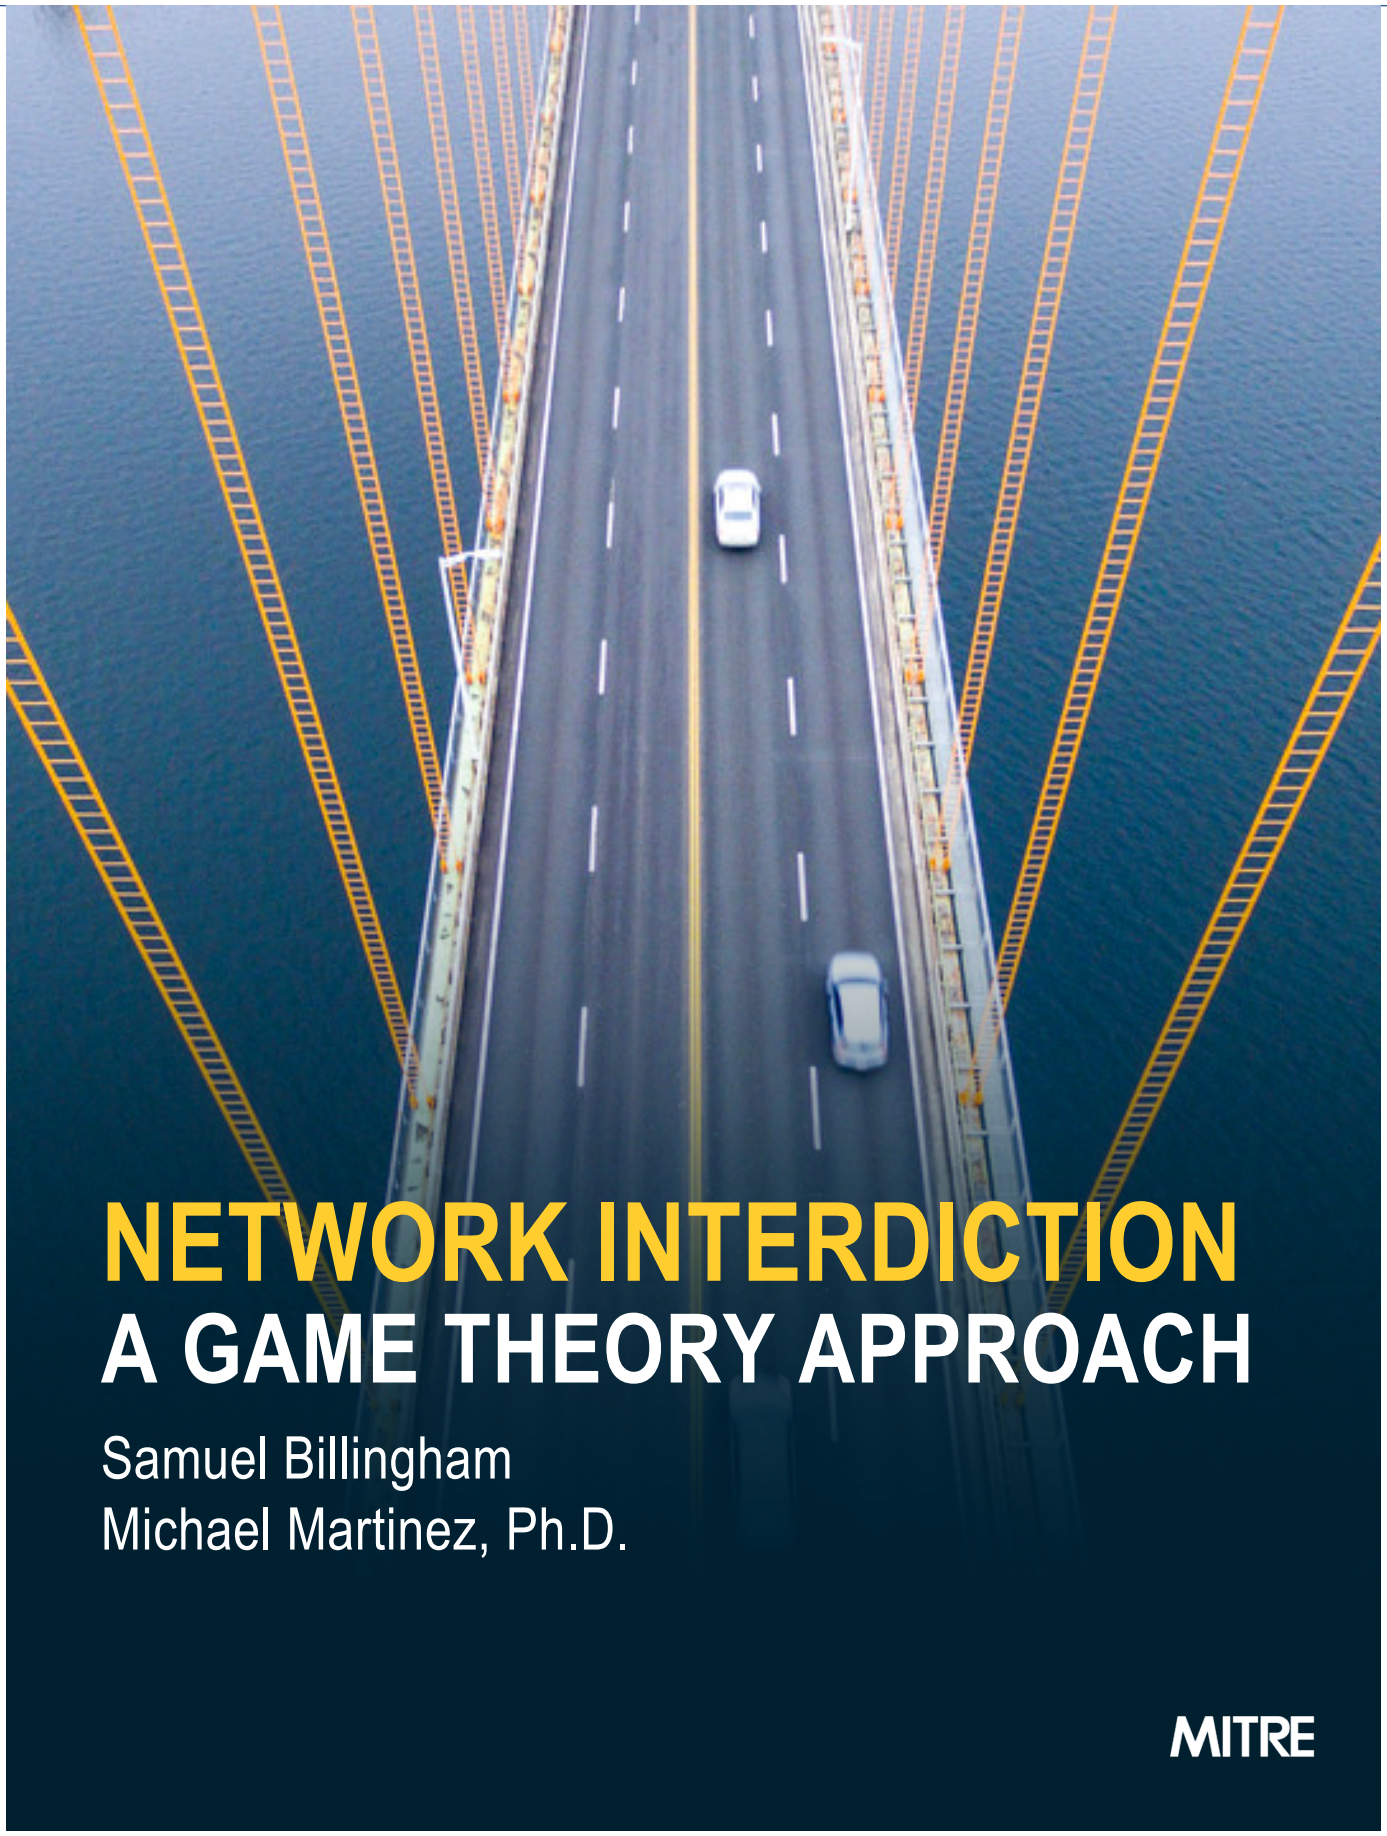

NETWORK INTERDICTION
A GAME THEORY APPROACH

Samuel Billingham
Michael Martinez, Ph.D.

MITRE

ABSTRACT

Ensuring the resilience of critical infrastructure networks, and adequately defending these networks for the purpose of national security, can be a steep challenge, especially in this age of contested environments. Beyond natural disasters and chance accidents, intelligent adversaries may be capable of inflicting serious operational impact with only a few strategically placed attacks. However, simple priority listings may result in overlooked vulnerabilities that are easily exploited. In these cases, careful analysis is required to identify the critical components (nodes and arcs) of the network, for the purpose of focusing efforts to enhance their resiliency. We discuss a three-player network interdiction construct that combines optimization and game theory, applied to a use case within the U.S. transportation network. This analysis provides a powerful, robust approach to identify critical infrastructure within the network.

THE CHALLENGE

The U.S. government operates and oversees a wide variety of networks, and these systems are often classified as critical infrastructure, essential for the health and welfare of the nation. These networks are generally administered with high efficiency, whether the objective is to maximize flow throughout or to send a specific set of items from "source nodes" to "demand nodes" at minimum cost. However, ensuring resilience of these networks, and adequately *defending* these networks for the purpose of national security, can be a steeper challenge, especially in this age of contested environments. Beyond natural disasters and chance accidents, intelligent adversaries may be capable of inflicting serious operational impact with only a few strategically placed attacks (kinetic and/or non-kinetic). In these cases, careful analysis is required to identify the critical components (nodes and arcs) of the network, for the purpose of focusing efforts to enhance their resiliency.

Generally, security planners in the Department of Defense (DoD) and Department of Homeland Security (DHS) are accustomed to creating a priority listing of network components, highlighting which pieces are most important to protect and defend. Depending on the available resources, the priority determines the set order in which to apply resiliency measures. However, it turns out that such a static priority listing may overlook significant gaps, where entirely different sets of components, if compromised, could combine to have a devastating effect on network operations. Priority lists and simple rules of thumb can be effective, but they do not *guarantee* that such a solution is optimal.

A Counterintuitive Example

Pretend we are charged with defending the simple network shown in **Figure 1**. The goal here is to push the maximum flow through the network, from node A to node F. The values shown on the arcs indicate the capacity of that particular link.

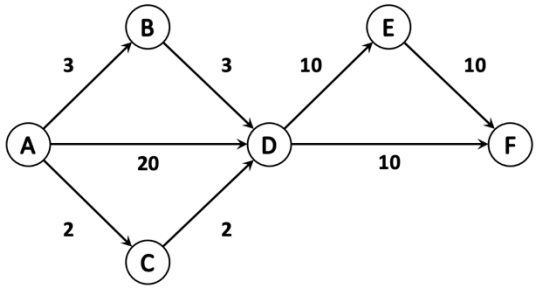

Figure 1. Example Network for Maximum Flow. Source: Naval Postgraduate School

The maximum flow through the network initially is 20 units. An intelligent attacker would clearly snip arc A→D, cutting the maximum flow down to only 5. *So a defender would naturally look to protect that "most-vital" arc.* However, what if the attacker is able to inflict *two* attacks instead of one? In this case, he has two optimal choices: snip D→E and D→F, or snip D→F and E→F. In either case, the flow is reduced completely to zero—and notice that arc A→D is no longer in the discussion! The attacker's strategy changed entirely, when a different budget, or capability "quiver" is considered. This quick and simple example defies a straightforward, prioritized list of critical network components.

Thus more scrutiny must be given when analyzing network flows. *Simple priority listings may result in overlooked vulnerabilities that are easily exploited.* Instead of creating static priority menus, an optimization approach to network interdiction should be adopted.

OUR APPROACH

In our work for combatant commands within the DoD, we employ optimization theory to this problem of network interdiction to ensure the U.S. government (USG) doesn't "overlook" an even stronger, counter-intuitive solution that may be hidden just beneath the obvious vulnerable nodes and links. This classic approach is rooted in mathematics and has been applied to various network flow problems for decades. Indeed, the RAND Corporation studied the Soviet rail system back in the 1950s for network vulnerabilities[1].

Many times, optimizing network flows is actually a simple problem, allowing most operators to run their systems at maximum efficiency. However, in light of hostile intentions, the situation becomes much more complex and the solutions not so obvious. In order to account for both sides, the optimization can be combined with Game Theory—specifically, the Stackelberg model. In this paradigm, the players take turns sequentially and have full, open knowledge of each other's moves, as in chess. The players' objectives are diametrically opposed, each seeking to counter the other's moves. Combining both optimization and game theory in this fashion provides a powerful approach to this class of network planning and interdiction problems.

The Naval Postgraduate School has researched these "Attacker-Operator" designs for over ten years, and has developed a robust library of applied models[2]. In fact, they have extended the construct to include a third player, a *Defender*, who seeks to assist the Operator by emplacing defenses at various components of the network. This

COMBINING BOTH OPTIMIZATION AND GAME THEORY IN THIS FASHION PROVIDES A POWERFUL APPROACH TO THIS CLASS OF NETWORK PLANNING AND INTERDICTION PROBLEMS.

interdiction game can be termed "Defender-Attacker-Operator" since that is the order in which the moves occur. As we saw with the sample network, critical network components are not so easily identified, so adding a third player allows for a much more rigorous cross-check of network performance.

A Solution Methodology

Moving first but possessing a limited budget, the Defender decides where to place defenses within the network. Seeing where these defenses are laid, the Attacker next chooses which components to attack, likely some combination of arcs and nodes. (One assumption here is that the Attacker will avoid defended components entirely and instead concentrate his strikes elsewhere.) Finally, the Operator must flow commodities across this network that has now been both defended and attacked—considering that defended components work as intended, while attacked components experience some measure of degradation, or perhaps temporary or entire failure.

As when exploring for a Nash Equilibrium within Game Theory, the players take turns optimizing from their perspective, while the resulting network performance metrics begin to converge. If minimum time (shortest path) is the primary objective for the game, then the Defender and Operator are choosing defenses and routes that will allow for

[1] See Alderson *et al.,* 2013.

[2] See Alderson *et al.,* 2011.

minimum overall transit time of the "cargo" in the network. The Attacker, however, intends to delay as much cargo as possible by as long as possible. As the game iterates, the players "learn" how to counter the opposing side, and may find better and better strategies to respond to their opponent while maintaining strong performance on their part. Eventually, the performance metrics balance, and no player can improve their strategy without sacrificing performance elsewhere—the solutions have converged.

With three players, the solution algorithm is fairly complicated, requiring two separate loops and passing solutions between players. The algorithm is outlined in **Figure 2**. Note that at each stage of the game, the current solution for various players is "locked" into place, while the opponent seeks the optimum strategy to counter the current status of the board.

In reality, the Defender moves first, followed by the Attacker and then the

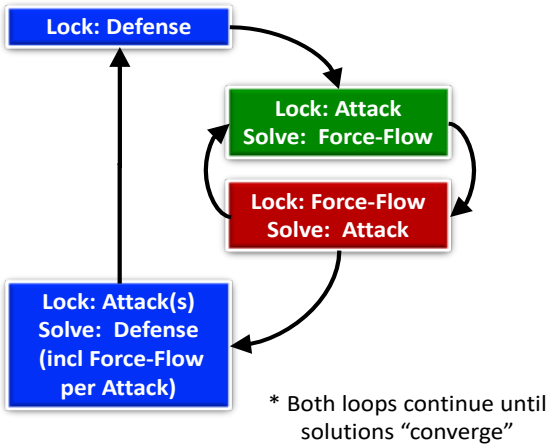

Figure 2. Solution Algorithm
Source: MITRE

Operator. However, the solution algorithm works in reverse, having the Operator and the Attacker engage first, seen in the green and red procedures; this inner loop iterates until their two solutions converge. At that point, the resulting attack plan is passed to and locked into place by the Defender, who tests all possible defenses against it, while simultaneously examining the subsequent Operator flow plan that would accompany each possible defense option. Once that solution is determined, that defense plan is locked into place and passed back to the inner loop, where the Attacker and Operator again iterate. As the attack plans are resolved, these are collected into a list, so the Defender must account for each of them—which places him in a "worst-case" scenario, unsure which attack plan the Attacker may utilize. This outer loop must converge in order for the overall solution to be decided. At that point, the solution yields an ideal defense strategy, a "worst-case" attack plan, and a corresponding optimal flow routing to mitigate the attack effects, while likely making use of defended infrastructure.

USE CASE

In 2019, NORAD & USNORTHCOM tasked MITRE to conduct a Resiliency Analysis of infrastructure critical to the force projection mission within the Continental United States (CONUS). The effort sought to identify transportation network assets critical to timely force projection[3] of military equipment and ammunition from bases through Ports of Embarkation over three transportation modes: rail, road, and oceanic shipping. The effort was limited to identifying only strategic infrastructure (roads, rail lines, intersections, ports, and shipping lanes), so it does not consider individual transportation vehicles (trucks, locomotives, railcars, vessels).

[3] A unit's expected arrival date in theater determined based on theater mission objectives.

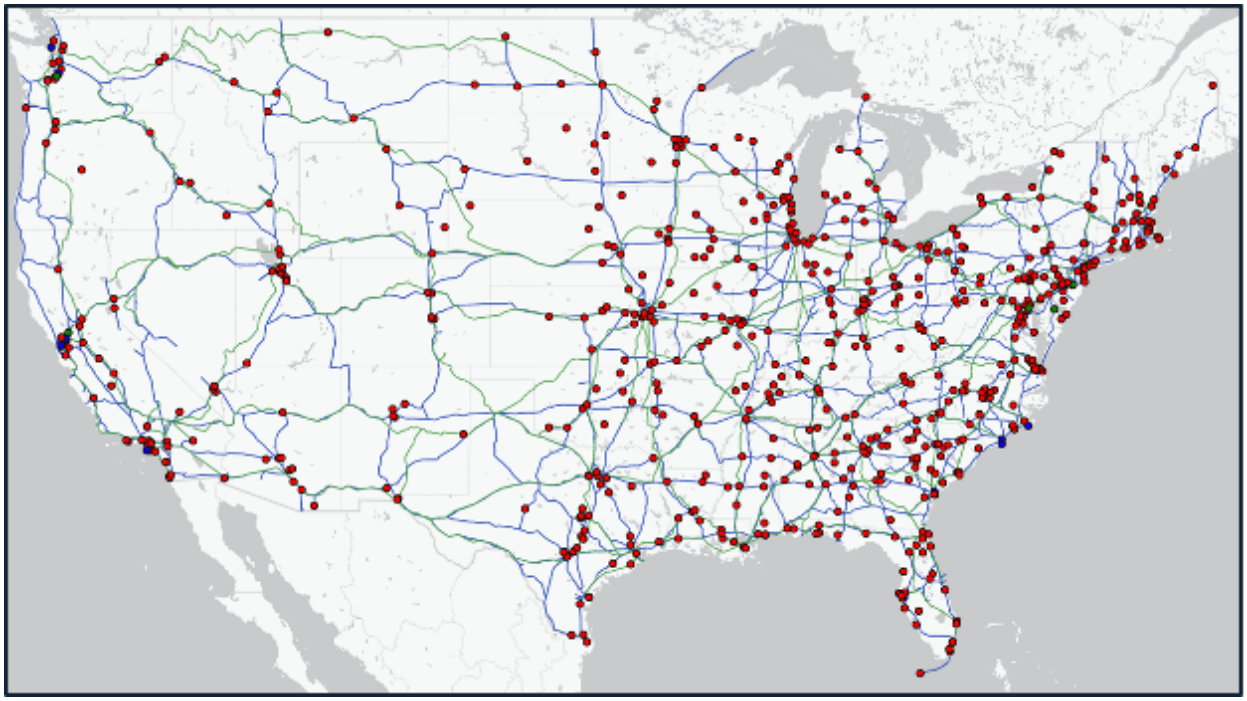

Figure 3. CONUS Transportation System with origins and ports (road - blue, rail - green)
Source: MITRE

We began with the network infrastructure. **Figure 3** shows the designated road and rail networks used to move equipment and ammunition throughout CONUS, as well as the origins of the equipment and ammunition. This network contains approximately 3,800 nodes and 5,200 arcs. To identify critical infrastructure, we assumed a minimum time objective, where an operator would flow network "cargo" items in as timely a manner as possible, an adversary injects delays by attacking the infrastructure, and a defender attempts to thwart these delaying attacks by defending the infrastructure. Along with the network, four types of data are required. First, the Time Phased Force Deployment Data (TPFDD) for an Operational Plan (OPLAN) provides the locations and amounts (in short tons) of planned equipment and ammunition to be deployed. Approximately 2,000+ unique items are modeled and moved across the infrastructure. Second, transit times (in hours) are needed along the various arcs (rail, road, shipping links) in the network. Third, processing times (in short tons per hour) are required for the various handling nodes within the network. Lastly, expected delay times for arcs (in hours) or nodes (in short tons per hour), should they be attacked, are required.

A model based on the Defender-Attacker-Operator (DAO) construct was used to identify critical infrastructure under various attacker and defender budget combinations along with flow times for the cargo. While detailed results of our analysis are sensitive, we provided a list of critical infrastructure to NORAD & USNORTHCOM, as well as details of what cargo is being delayed and the amount of delay under various attacker/defender budget (and thus, capability) combinations. Note that adjusting the budget combinations allowed us to identify any fundamental changes to the Attacker or Defender strategies, as shown in the initial example.

Infrastructure that is defended consistently across budget scenarios is considered more

critical to timely force projection than infrastructure that is routinely attacked, which in turn is considered more critical than infrastructure that is typically neglected. The critical infrastructure lists allow for focusing of resources to conduct in-depth vulnerablity assessments at the identified infrastructure. Also, theater delay results support mission risk assessments due to late arrivals in the receiving theater. In general, we found that an attacker can inflict enough delay into the force projection network to place theater mission objectives at risk with very small budgets. We also found (to our relief!) that mitigating much of the delay is possible with the application of even fewer defenses.

ADVANTAGES & DISADVANTAGES

The advantage of the DAO methodology is its extensibility. While the use case demonstrated its applicability to a transportation network, it is not limited to transportation. As long as the underlying system is a network (e.g, communications, command & control) the methodology can be used to identify where potentially dangerous vulnerabilities exist in the system. Likewise, while the use case demonstrates a minimum-timeliness example, the methodology also supports other (minimum-cost, maximum-flow) objectives.

Another advantage of this methodology is its practicality. Because it is simply searching for elements where vulnerabilites exist, and not describing these vulnerabilites at an engineering level, it allows decision makers, via a broad look at a complex system, to readily target where risks to their missions could exist when facing an intelligent adversary. It also allows for the quick identification of where they need to focus future efforts and resources for risk analyses and vulnerability assements. Also, while not demonstrated in the use case, the method could be adapted for two practical extensions: distinctions between attack types (e.g. kinetic versus non-kinetic attacks) and the impacts of area attacks/defenses (i.e., attacks/defenses that, if employed, affect multiple components of the network simultaneously). While this adds some complexity to the analysis, it is not insurmountable.

Another advantage of this approach is that it is not as sensitive to erroneous input data as others. That is, utilizing the best available data, one can readily identify vulnerabilities.

WHILE NO METHODOLOGY COMPLETELY SURVIVES "GARBAGE IN," THIS METHOD IS FAIRLY ROBUST TO PERTURBATIONS IN THE INPUT DATA.

Much of the input data used for the handling and movement times in the use case can be found in open sources. Likewise, since attack impacts are less known, subject matter expert estimates were used and applied consistently across the network; vulnerable nodes are still identified. Sensitivity analyses for these rough estimates were conducted, showing that the results hold up to moderate changes in the original inputs. If leaders (or analysts) are concerned that the input data is not precise enough to trust the initial results, then the inputs may still be used to identify and narrow the focus for follow-on data collection efforts, streamlining the refinement task.

A couple of disadvantages of this method are the requirement of specialized software and high-powered computing to solve these problems on a massive scale such as our

example. While optimization solvers range anywhere from the simple engine found in MS Excel to powerful commercial solvers (e.g., Gurobi) and anywhere in between, the scale of these problems often drives the need for more expensive options. There has been advancement in open-source solvers over the years and some could be used to solve smaller versions of these network problems. The computing challenge centers around the need for a capable, single machine to solve the problem. That is, the "solve" operation for these DAO problems is not parallelizable, the way most modern simulation efforts or gaming applications are. Thus, a single powerful processor is needed. Luckily for our use case, we were able to utilize an on-hand computer at NORAD & USNORTHCOM.

CONCLUDING THOUGHTS

Network interdiction is becoming an important aspect of national security planning, since both conventional and (oftentimes simplistic) asymmetric threats can be readily employed to cripple critical networks, for discrete windows of time. For example, the crash of the Hong Kong-flagged ship, the *Ever Given*, dramatically affected international shipping by blocking the Suez Canal for six days in 2021. Our "Forts-to-Ports" application allowed security planners to focus future resiliency enhancing efforts, and requests for resources where they are needed most, by readily identifying critical network components. By knowing where to look, the effects of intelligent attacks can more readily, rapidly, and to the maximum extent possible, be mitigated.

ABOUT THE AUTHORS

Samuel Billingham is an Operations Research Analyst with The MITRE Corporation. He holds a M.S. in Operations Research from The University of Alabama in Huntsville.

Michael Martinez is an Operations Research Analyst with The MITRE Corporation. He holds a Ph.D. in Operations Research from the Colorado School of Mines, and served on active duty with the U.S. Air Force before retiring as a Lieutenant Colonel in 2020.

ABOUT MITRE

MITRE's mission-driven teams are dedicated to solving problems for a safer world. Through our public-private partnerships and federally funded R&D centers, we work across government and in partnership with industry to tackle challenges to the safety, stability, and well-being of our nation.

The views, opinions, and/or findings contained herein are those of the authors and should not be construed as an official government position, policy, or decision unless designated by other documentation.

REFERENCES

Alderson, D. L., Brown, G. G., Carlyle, W. M., & Wood, R. K. (2011). *Solving Defender-Attacker-Defender Models for Infrastructure Defense*. in ICS 2011, 12th INFORMS Computing Society Conference, pp 28–49. https://doi.org/10.1287/ics.2011.0047

Alderson, D. L., Brown, G. G., Carlyle, W. M., & Cox, L. A. (2013). *Sometimes There is No 'Most-Vital' Arc: Assessing & Improving the Operational Resilience of Systems.* Military Operations Research, Vol. 18 No. 1, pp 21–37. doi: 10.5711/1082598318121