

Operations Research - A Catalyst for Engineering Grand Challenges



**Respectfully Submitted to the
National Science Foundation
May 2014**



Operations Research - A Catalyst for Engineering Grand Challenges

May 28, 2014

Suvrajeet Sen, Chair
University of Southern California

Cynthia Barnhart, Massachusetts Institute of Technology
John R. Birge, University of Chicago
E. Andrew Boyd, PROS
Michael C. Fu, University of Maryland
Dorit S. Hochbaum, University of California -Berkeley
David P. Morton, University of Texas-Austin
George L. Nemhauser, Georgia Institute of Technology
Barry L. Nelson, Northwestern University
Warren B. Powell, Princeton University
Christine A. Shoemaker, Cornell University
David D. Yao, Columbia University
Stefanos A. Zenios, Stanford University

This report was funded through NSF Grant CMMI-1243182, and we are grateful for the support. The statements and conclusions in this report are not necessarily those of the National Science Foundation.

OR as a Catalyst

Reviewers

Brenda J. Dietrich (IBM)
Peter W. Glynn (Stanford University)
Jong-Shi Pang (University of Illinois and USC)
Andrew J. Schaefer (University of Pittsburgh)
J. Cole Smith (University of Florida)

Acknowledgments

The authors of this report are grateful to NSF for its continued appreciation of the impact of Operations Research, and its growing role in our highly interconnected society. We are also grateful for the tremendous support from many leaders in the OR community, especially, the diligent and conscientious team of reviewers. Their insightful feedback has helped guide the recommendations of this report. Suvrajeet Sen, the chair of the team of authors, is grateful for the opportunity to work with an exceptionally accomplished and collegial team of authors and reviewers. The chair is also grateful to Randolph Hall (VP Research, USC) for his remarks on an earlier version, and to USC's Epstein Institute for its support. Finally, this report would not have taken the shape it has, were it not for the diligence of Siqian Shen (Michigan) who kept great records of discussions/presentations at our meetings, and to Sarah Crabb for her timely editorial support.

OR as a Catalyst

Executive Summary

The growth and success of Operations Research (OR) depends on our ability to transcend disciplinary boundaries and permeate the practices of other disciplines using ideas, tools, and experience of the OR community. This report is intended to continue the tradition of transcending disciplinary boundaries by using the U.S. National Academy of Engineering's (NAE) Engineering Grand Challenges as a source of inspiration for the OR community. Our goal is to view these challenges as an opportunity for the OR community to play the role of a catalyst - utilizing OR to facilitate some pressing technological challenges facing humanity today.

A panel of thought-leaders convened by the NAE (and facilitated by NSF) unveiled its vision of the Engineering Grand Challenges in 2008. Over the past several years, this report has invited (and received) feedback from international leaders and professional organizations, including the Institute for Operations Research and the Management Sciences (INFORMS). As input from the OR community, several past Presidents of INFORMS prepared a white paper, an abbreviated version of which appeared as the President's Column in *OR/MS Today* (April 2008). As predicted, the OR community has been active in many of the thematic areas of the NAE Grand Challenges via publications in topical research areas of our flagship journals, joint major conferences, and other collaborative efforts. The question of whether there are ways to dovetail OR with these challenges is not the issue. Of importance is whether there is a need to introduce greater structure for research and exchange between domain experts in core areas of the engineering Grand Challenges and the OR community.

In order to accelerate the growth, this report recommends a two-pronged approach: (1) An NSF announcement of "Grand Challenge Analytics" as a major EFRI topic, and (2) an NSF sponsored institute for "Multidisciplinary OR and Engineering" which will be dedicated to coalescing a general-purpose theory, as well as building a community to support "Grand Challenge Analytics". Together, these initiatives are likely to unleash a vast array of methodologies onto the engineering Grand Challenges of today. Such an effort could be likened to the manner in which the interface between OR and computer/communications science/engineering has propelled the development of the Internet. Similarly, the long-standing exchanges between the INFORMS and Economics communities has produced deep results, many of which have been honored by the Nobel Prize in Economics. Drawing upon such successes, we propose a new era in which the OR community reaches out to domains that are more directly connected to the NAE Grand Challenges. This more structured approach, driven by NSF sponsorship of research and thematic exchanges (workshops), will result in well-defined outcomes, leading to a strong foundation for the NAE Grand Challenges.

OR as a Catalyst

Operations Research – A Catalyst for Engineering Grand Challenges

Contents

I. Introduction	8
II. OR: A General-Purpose Theory of Analytics	10
III. OR for Sustainability	13
IV. OR for Security	23
V. OR for Health Care	32
VI. OR for Joy of Living.....	40
VII. Active Steps for Catalysis	48
References	50

I. Introduction

Since its inception in the 1940's and 50's, Operations Research (OR) has addressed challenges at the interfaces of traditional disciplines. Many of the awards in this discipline (e.g., the Von Neumann Theory Prize, Lanchester Award) celebrate accomplishments of researchers (e.g., John von Neumann, Fredrick Lanchester) who brought fresh perspectives to some well-established field (i.e., economics, military), transforming the field itself and building the foundations of the new field of OR. Today, we witness the impact of OR in technologies ranging from control systems (interior point methods) and the Internet (distributed dynamic alternative routing) to machine learning (convex optimization) and artificial intelligence (Markov decision processes). In some cases, the penetration of OR has been so fundamental that the OR community has become the leading force in the discipline. In the world of analytics, areas such as portfolio optimization, energy markets, supply chains/logistics, and many others have become so OR-centric that imagining those fields growing without the benefit of OR would be impossible. Consider for instance IBM's vision of building a "Smarter Planet". In its overview of this effort, the industrial giant provides the following situations in daily life,¹ quoted verbatim below.

- *Congested roadways in the U.S. cost \$78 billion in 4.2 billion lost work hours and 2.9 billion gallons of wasted gas annually—and that's not counting the impact on air quality.*
- *Inefficient supply chains cost \$40 billion annually in lost productivity — more than 3% of total sales.*
- *Our healthcare system really isn't a "system." It fails to link diagnoses, drug delivery, healthcare providers, insurers and patients — as costs spiral out of control, threatening both individuals and institutions.*
- *One in five people living today lacks safe drinking water.*
- *And, of course, we've seen what's developed in our financial markets, a system in which institutions could spread risk, but not track it.*

All of these situations call for greater interplay between analytics and systems of various types.

The growth and success of OR depends on our ability to transcend disciplinary boundaries and permeate the practices of other disciplines using ideas and tools pioneered by the OR community. This report is intended to continue the tradition of crossing disciplinary boundaries by using the U.S. National Academy of Engineering's (NAE) Engineering Grand Challenges as a source of inspiration for the OR community. Our goal is to view these challenges as an opportunity for the OR community to play the role of a catalyst - utilizing our ideas and tools to address some of the more pressing technological challenges facing humanity today. *Because of this emphasis, the report will NOT focus on challenges for OR; instead the focus is on "Catalysis."*

¹ http://www.ibm.com/smarterplanet/global/files/us__en_us__general__smarterplanet_overview.pdf

A panel of thought-leaders convened by the NAE (and facilitated by NSF) unveiled its vision of the Engineering Grand Challenges in 2008. Over the past several years, this report has invited (and received) feedback from international leaders and professional organizations, including the Institute for Operations Research and Management Science (INFORMS). As part of the INFORMS input, Barnhart et al (2008) prepared a report on the role that the OR community was likely to play within the context of the challenges. An abbreviated version of that report appeared as the President's Column in *OR/MS Today* (April 2008). As predicted, the OR community has been active in many of the thematic areas of the NAE Grand Challenges via publications in topical research areas of our flagship journals, joint major conferences (e.g., the joint INFORMS-Medical Decision-Making Conference in Phoenix 2012), and several thematic conferences on the Smart Grid, Homeland and Cyber Security, and others. Whether there are ways to dovetail OR with these challenges is not the question; *the question we address is: what benefits might accrue if there was a focused effort to grow OR methodology to address issues which arise from the NAE Grand Challenges?* Because OR brings together a combination of tools from computing, mathematical, and economic sciences, such an effort is likely to unleash a vast array of new approaches onto the engineering grand challenges of today. Such an effort could be likened to the manner in which the interface between OR and computer/communications science/engineering has propelled the development of the Internet towards a decentralized yet stable system based on OR models of fairness, optimality, shadow prices, stability etc. (Kelly et al. 1998). The scalability of these methods has been central to the scalability of the Internet.

This report is organized as follows. The next section provides the backdrop of OR: its tradition as a nexus for computational, mathematical and economic sciences leading towards a general-purpose theory of analytics. The four subsequent sections are devoted to the broad classifications of the NAE Grand Challenges: Sustainability, Security, Human Health, and Joy of Living. Each of these sections also includes a subsection aimed at the interface between the challenge domain and the Social and Economic Sciences. This choice reflects the long-standing presence of OR as a methodology for the design and analysis of socio-technical systems (e.g. transportation) where the success of a technology requires an understanding of human behavior. In addition, such collaboration between OR and Economics has helped build a solid scientific foundation, as evidenced by the Nobel Prize winning works of Arrow, Debreu, Markowitz, Nash, Shapley, V. Smith, and others with strong ties to OR.

Drawing upon such successes, we propose a new era in which the OR community reaches out to domains which are more directly connected to the NAE Grand Challenges. To accomplish this, we recommend two initiatives: one to build research bridges between the grand challenge domains, and analytics, and the other dedicated to coalescing both the knowledge as well as the community through a new Institute under the banner of Multidisciplinary Operations

Research and Engineering (MORE). This vision builds on the OR tradition of outreach to many domains. In fact, as noted throughout this report, the needs identified in the NAE report have already spawned research collaborations between the OR community and several science and engineering domains. However, a more structured approach, consisting of a research plan, coupled with regular thematic exchanges (workshops), will result in well-defined outcomes, leading to more long-term partnerships between the communities. Section VII outlines how this vision may be realized through mechanisms already available within NSF.

II. OR: A General-Purpose Theory of Analytics

OR principles pervade engineering, science and the economic sciences, but so do other computational and mathematical sciences (including statistics). However, OR is unique in its origins and content; it derives almost equally from induction and deduction. For instance, the discovery of the Simplex method was tied closely to planning problems during World War II and this form of OR may be looked upon as inductive. On the other hand, the recent surge of convex optimization as a unifying theme for high dimensional (and sparsity preserving) regression (Bühlmann and van de Geer 2011), machine learning (Xiao 2010), control, signal processing and others (Chiang et al 2007), suggests a deductive approach. Indeed as a field, OR represents the crossroads of computer, mathematical and economic sciences. For this reason, many graduate courses in OR routinely include graduate students from these disciplines. By integrating these fields, OR helps build scalable models and methods for many real-world applications. This integrative feature of OR also uncovers structural similarities between vastly different applications, and provides the foundations for a general-purpose theory of Analytics.

One of the other distinctions between OR and other mathematical sciences is that many of our premier journals are not only devoted to the methodology, but also to industrial strength applications. It is such breadth that prompted George Dantzig to label OR as a “can-do” discipline. This is the likely reason for major OR conferences to cut across both theory and applications, covering mathematical, algorithmic, organizational, and practical implementations of OR. This unique style of discovery may also be traced to the fact that as a field, OR is applied daily to an amorphous body of applications covering a gamut of domains, including the Grand Challenges. In addition, the broad reach of OR also touches industrial innovation. For these reasons, companies that are interested in research at the intersection of business analytics, computing, and mathematical sciences routinely appoint doctorates from the field of OR to lead large research labs (e.g., IBM’s T.J. Watson Labs).

Major users like the Department of Defense are also clamoring for advanced OR tools for trade-off studies (Defense Science Board Report, 2011, especially Chapter 3). Similarly, other

agencies have found great value in OR tools. For instance, one such observation appears in the PCAST report² of 2010, which we quote in the following table (Holdren et al, 2010).

[Progress in Algorithm Beats Moore's Law \(Source: Holdren et al, 2010, pp. 71\)](#)

“Everyone knows Moore’s Law – a prediction made in 1965 by Intel co-founder Gordon Moore that the density of transistors in integrated circuits would continue to double every 1 to 2 years.

Fewer people appreciate the extraordinary innovation that is needed to translate increased transistor density into improved system performance. This effort requires new approaches to integrated circuit design, and new supporting design tools, that allow the design of integrated circuits with hundreds of millions or even billions of transistors, compared to the tens of thousands that were the norm 30 years ago. It requires new processor architectures that take advantage of these transistors, and new system architectures that take advantage of these processors. It requires new approaches for the system software, programming languages, and applications that run on top of this hardware. All of this is the work of computer scientists and computer engineers.

Even more remarkable – and even less widely understood – is that in many areas, performance gains due to improvements in algorithms have vastly exceeded even the dramatic performance gains due to increased processor speed.

The algorithms that we use today for speech recognition, for natural language translation, for chess playing, for logistics planning, have evolved remarkably in the past decade. It’s difficult to quantify the improvement, though, because it is as much in the realm of quality as of execution time.

In the field of numerical algorithms, however, the improvement can be quantified. Here is just one example, provided by Professor Martin Grötschel of Konrad-Zuse-Zentrum für Informationstechnik Berlin. Grötschel, an expert in optimization, observes that a benchmark production planning model solved using linear programming would have taken 82 years to solve in 1988, using the computers and the linear programming algorithms of the day. Fifteen years later – in 2003 – this same model could be solved in roughly 1 minute, an improvement by a factor of roughly 43 million. Of this, a factor of roughly 1,000 was due to increased processor speed, whereas a factor of roughly 43,000 was due to improvements in algorithms! Grötschel also cites an algorithmic improvement of roughly 30,000 for mixed integer programming between 1991 and 2008.”

² President’s Council of Advisors on Science and Technology

It is said that the benchmark that Grötschel was using was drawn from Bixby (2002), published in the INFORMS flagship journal: *Operations Research*. These advances are not merely new computational benchmarks; they have enormous economic impact as well. For instance in the electricity sector, OR models have transformed market operations with advanced software for scheduling, economic dispatch, power flow and many other functions. A recent FERC report (FERC 2011) attributes savings in the range of \$100 million annually for generator scheduling at just one of the markets (PJM, Pennsylvania, New Jersey, Maryland). Savings of this nature are also reported at other markets, as well as other applications. OR has also orchestrated transformative changes in other sectors of the economy (e.g., airlines, finance, logistics, manufacturing/supply-chains, sports, etc.), and we expect that OR leadership in analytics will further solidify its impact in industry (Nemhauser 2013).

These vignettes, as well as examples throughout this report suggest that we are witnessing a convergence of systems thinking and computational thinking, which has made OR a central enabling technology for a variety of applications. Nevertheless, the level of penetration of OR in some important domains of science and engineering *research* has been somewhat limited. There can be serious consequences of minimal exchanges between OR and some of the Grand Challenge communities. For instance, the recent NOAA report “The State of the Climate in 2012” noted that the year 2012 was one of the 10 hottest years on record (global average), and moreover, a vast majority of the top 10 hottest years have been recorded since the 1990s. In connection with the above report, NOAA’s acting administrator, Kathryn Sullivan, mentions that “many of the planning models for infrastructure rely on the future being statistically a lot like the past, and certainly the data should lead one to question if that will be so.” In other words, future planning models ought to recognize on-going changes due to human activities. In fact, the humanity is not in a position for much trial-and-error policy-making in this arena. For far too long, science and engineering have failed to recognize the need for scalable models and algorithms. With the increasing need for guided scientific exploration, OR approaches for scalable models and algorithms are becoming indispensable. Given the magnitude of the situation, OR modeling should be playing a bigger role in understanding the impact of human choices on the future of the planet. One of the strengths that the OR community can bring to such areas are models and tools that integrate data and decisions. This interplay is the key distinction between OR models and descriptive statistics. Moreover, the former also facilitate risk modeling in a resource constrained setting. The time has come to engage both domain experts as well as OR experts, so that policies/decisions become an integral part of analysis, not an afterthought. Such collaboration has the potential to discover strategies to reverse the ominous trends that have been observed over the past two decades. Through its wide-ranging portfolio of investments in science and engineering *research* NSF is in an ideal position to lead a program of research and dissemination between OR and other *research* communities.

III. OR for Sustainability

Foremost among the challenges are those that must be met to ensure the future itself. The Earth is a planet of finite resources, and its growing population currently consumes them at a rate that cannot be sustained. Utilizing resources (like fusion, wind, and solar power), preserving the integrity of our environment, and providing access to potable water are the first few steps to securing an environmentally sound and energy-efficient future for all of mankind. Analyzing these challenges includes using data and models to choose among alternative management decisions, forecasting the effect of decisions on the future, and quantifying the uncertainty associated with this analysis. While OR methods certainly support many individual features (e.g., nonlinear, dynamic, stochastic or discrete), combinations of these features are often necessary in many of the research questions which arise under the “Sustainability” banner. Combining the features above will require significant extensions of the OR methodology available today.

Make Solar Energy Economical

The sun provides an incredibly powerful and extensive resource for our energy demands. It out-powers anything that human technology could ever produce and produces energy that neither harms our environment nor costs the planet anything. However, our ability to harness solar energy is tempered by our abilities to make it economical. Current technology includes solar photovoltaic (PV) cells and concentrated solar power. These technologies are not considered commercially competitive due to a combination of conversion efficiencies, the relatively high cost of installation, and the challenge of managing the variability and uncertainty of solar energy in a grid that demands absolute reliability. Many consider energy storage to be a critical technology in the use of solar energy. However, cost-effective bulk storage for the grid remains elusive. OR offers systematic set of tools with wide-ranging applicability; some are strategic, and others are somewhat tactical. For example, strategic applications may cover issues in new materials discovery, streamlining installation processes, and designing R&D portfolios. From a tactical point of view, OR algorithms help manage variations of solar energy in a portfolio of energy generation, storage and demand management assets, and many more. At best, current efficiencies are in the range of mid-30% because the systems depend on a number of random variables governed by nature.

A major impediment to scaling solar energy to significant levels of our electricity generation is the volatility due to cloud cover, as well as structural imbalances between solar generation, which peaks mid-day, and energy usage, which peaks in the morning and evening in winter months. Grid operators have to carefully balance generation from steam (which today is derived from nuclear, coal and, increasingly, natural gas) and gas turbines. Steam generation

has to be planned a day in advance, while gas turbines may be planned an hour in advance. These different types of generators have to be managed in combination with energy storage and demand-side management options (“demand response”). Combining advance planning with the uncertainty of cloud cover introduces the need to produce robust solutions to very high-dimensional stochastic, dynamic programs. The OR community is the primary force behind the effort to evaluate policies and technologies to integrate renewable resources into the smart grid (FERC Technical Conference, 2013).

To improve the efficiency of solar panels, the NAE grand challenges report suggests that new discoveries that leverage advances in nanotechnology to almost double the current efficiency may be on the horizon. These new applications of nanotechnology are made possible by the use of optimization to maximize absorption of light in the realm of nano-photonics. Indeed, the advent of new nanotechnology applications for solar cells is establishing optimization as one of the more important tools to change the landscape of solar cells. Until recently, enhanced absorption of light in nano-photon structures was obtained through heuristic design choices that were unable to optimize absorption to maximize efficiency of the nano-photon structure. However, recent research (e.g., Wang and Menon, 2013) is beginning to change that. By integrating numerical simulation with optimization for designing nano-photon devices, engineers are betting on revolutionizing PV system design. We expect advanced OR to play an integral role in meeting this challenge. Indeed, a systematic study has already brought researchers in Aerospace and OR together in formulating the so-called bandgap optimization problem, which has many applications ranging from the design of photonic crystal to fabrication aware optimization (Men et al. 2010; 2011; & 2013). One of the journals publishing this work happens to be *Operations Research*.

Similar advances are being pursued to improve the efficiencies and lifetimes of battery storage. Anodes in battery storage tend to be the weak link in lithium-ion batteries because they absorb the stress of repeated charge-discharge cycles. Nano-materials are being devised which triple the lifetimes of anodes, significantly lowering the lifetime cost of a battery. These advances are no longer expected to come simply from laboratory experiments. In looking ahead to future research, the DOE-sponsored report states, “Simulation-based engineering and science has accelerated progress in scientific understanding and technology development by enabling complex systems such as astrophysical and climate phenomena, aircraft wings, and integrated manufacturing to be explored rapidly and efficiently. This leads to new scientific understanding of systems that are too large for experimental study, reduced time and cost of prototyping, and accelerated deployment of new technologies” (DOE Workshop, 2010). Operations Research is already being used to guide the process of experimentation.

In addition to increasing the efficiency of these cells, the economics of solar power will also rely on innovations to lower their manufacturing and installation costs. The installed cost of a solar panel is over \$5/watt while the cost of the solar panel itself is around \$0.65/watt (and falling). OR has a long history of improving the efficiency of service operations that govern the implementation process. Applying this expertise to improving the efficiency of solar panels will ultimately make solar power economically competitive.

Integrating Storage with Solar Cells

Renewable energy offers the best potential for zero CO₂ emissions, but the lack of control over supply has drawn national attention to the importance of storing energy from electricity. In addition, while there is tremendous attention on breakthrough technologies such as metal-air batteries, these technologies are expensive and require significant research advances before they will become widely adopted. Hundreds of research labs are working on designing new technologies for batteries, but there is also tremendous potential for simply utilizing the resources and technologies that already exist.

Surprisingly, finding optimal policies for managing energy storage resources is beyond the state of the art of modern algorithms for all but the most trivial systems. Real battery systems have to deal with dynamic energy generation from the wind and sun, stochastic prices, time varying loads, and complex battery chemistry that affects the efficiency of energy conversion and battery lifetimes.

We do not yet have the tools to obtain optimal policies for realistic battery systems. Standard practice in engineering is to live with various heuristics, but these can perform well below optimality. Proper battery management, especially for affordable technologies such as lead-acid batteries, can dramatically extend the lifetime of the device. The best policies can outperform “good” policies by 100 percent. Furthermore, it is likely that these algorithmic advances would benefit both existing and future storage technologies.

The next generation of optimal storage control algorithms needs to handle a number of problem characteristics that are beyond the reach of existing tools:

- a. Realistic battery chemistry: Batteries exhibit memory, latency, losses, and aging that can be affected by how they are charged/discharged.
- b. Batteries have to be operated in the context of complex “state of the world” variables including temperature, loads, prices, energy from wind or solar, and the accuracy of forecasts.
- c. Batteries may be coupled with devices such as ultra-capacitors, or may be a part of larger portfolios of storage devices.

- d. Batteries cannot simply charge/discharge at will: Battery operators have to bid into markets and work with existing market mechanisms.

Hundreds of millions of dollars are being invested in advances in energy storage technologies. A 40 to 100 percent improvement in output and lifetimes would be considered a major breakthrough. We can achieve these results simply with advances in optimization algorithms.

Provide Energy from Fusion

Given the perception that Fusion Energy Science (FES) is based on plasma physics this particular challenge might be considered too far-removed for OR to be of any consequence; however, this problem demonstrates the broad reach of OR: OR sits at the intersection of Computational and Mathematical Sciences, and this juxtaposition provides the perspective for modeling and simulation in any modern scientific endeavor. FES is no exception.

Research in FES is a worldwide effort headquartered at ITER (International Thermonuclear Experimental Reactor) located in the south of France. This project, with expected costs to exceed 15 billion Euros, is funded by seven member entities: the European Union, India, Japan, China, Russia, South Korea, and of course, the United States. This truly multi-national project will require multi-faceted collaboration among a variety of scientists and engineers. Because of such demands, the Director of the DoE Office of Science requested the U.S. plasma and fusion research communities to examine opportunities for collaborative science in the area. A highly influential committee of leaders (FESAC) prepared a report on behalf of the community. This report, released on February 28, 2012, suggests a framework based on three major thematic areas of collaboration. Two areas are experimental in nature, and the third is entitled "Understanding the dynamics and stability of the burning plasma state." ITER will test the ability of magnetic confinement to hold the plasma in place at high-enough temperatures and density for the amount of time necessary to complete the fusion reaction. Clearly, modeling and simulation will form an integral part of predicating outcomes of this test. The report recommends:

"Performing predictive modeling in advance of these experiments would provide a stringent test of the models and strengthen the foundation for predicting and optimizing the performance of ITER. ... Optimizing the fusion performance in ITER requires understanding the interplay between the many factors involved in tokamak³ operation and building this into simulation codes. The US has made important strides in

³ Tokamak is a device where the fuels are injected into and confined in a vacuum chamber and heated to temperatures exceeding 100 million degrees.

this area but to be confident of extrapolating to ITER-like conditions it is important to validate the underlying models at the largest available scale.”

The FES community is already in the process of using optimization models and simulations for ITER. Physicists at EURATOM-Suisse in Lausanne, Switzerland recently published research that uses a physics-based nonlinear model of plasma profile dynamics, while minimizing a cost function, under a set of constraints on the actuators and plasma quantities (Felici and Sauter, 2012). The model is, in fact, a PDE-constrained optimal control problem which is discretized and solved using a successive quadratic programming algorithm. While the effort presented in the paper is impressive, numerous opportunities still exist for more advanced OR tools which would allow modelers (such as the plasma physicists in this project) to invoke spatio-temporally driven optimization algorithms for models of this type without having to cast their specific instances into optimization paradigms in which spatial and temporal considerations are only incidental. More modern tools which allow us to incorporate PDE-based simulations into optimization algorithms will be able to provide the kind of infrastructure that will allow physicists to investigate the sensitivity of various parameters, and in turn, may be able to design more powerful reactors than is currently possible.

These complex engineering problems represent a significant algorithmic challenge. Frequently, we are faced with the problem of optimizing 10 or 20 parameters without the benefit of derivatives, often in the presence of measurement noise and, in some cases, without the benefit of convexity. In the worst cases, the optimization has to be accomplished using physical experimentation, which is time consuming and very expensive. While substantial literature has evolved around these problems, we still do not have robust, efficient algorithms that practitioners can use.

Geological Carbon Sequestration

In a series of papers L. Thompson and his team (Thompson et al. 1995; 2000; & 2002) have provided evidence from a variety of geographical locations world-wide that average global temperatures are on the rise. The rate of increase in carbon dioxide emission may be the single largest culprit among all greenhouse gases. It is predicted that if these levels continue to rise, the consequences could be severe: rising sea levels, disruptions in agriculture, and other natural disasters (e.g., hurricanes, tornadoes) may begin striking with greater force and frequency. The NAE Grand Challenges report suggests that a central task for engineers will be to select sites for storing carbon dioxide and to design safe methods for sequestering carbon away from the atmosphere. Some possibilities for Geological Carbon Sequestration (GCS) include storing the gas in deep saltwater aquifers or depleted oil reservoirs.

GCS has tremendous potential because the capacity is so large it could be used to store the global output of carbon for decades while renewable energy technology is being developed by the IPCC (Metz et al. 2005). However there are potential problems, both economic and environmental, with GCS. The environmental risks are that injected CO₂ could increase subsurface pressure to the point that CO₂ or saltwater would move upward through fractures or defective well casings. Such movement could result in pollution (including increased metal concentrations) of freshwater aquifers used for human drinking water or in the escape of CO₂ into the atmosphere, which at worst, could cause deaths and at best represent a waste of the large amount of money spent to sequester the CO₂.

OR methods can assist in promoting safe geological sequestration of carbon by helping to design monitoring systems and using optimization and statistical methods to estimate the current and future spatial distribution of CO₂ and pressure in the injection region. Note that these injection regions are expected to be at least 1000 meters below the surface. Because the injection site is so deep, the monitoring wells also need to be very deep and are therefore extremely expensive. As a result, a practical monitoring system might consist of only a few *strategically located* monitoring sites, some of which may be in regions outside the CO₂ plume to monitor pressure without increasing the risk of CO₂ escape. Since the CO₂ will take the shortest path upwards, knowing the location of the CO₂ plume is very important, and sensor locations should be chosen to identify leaks reliably. In addition, the increased pressure in CO₂ and saltwater increases the chances of earthquakes (Zabaco & Gorelick, 2012). Injected CO₂ has the potential to spread over hundreds of miles, and the pressure increase can be even more widespread. For this reason, dynamically estimating pressure distribution and narrowing the highest risk areas for CO₂ is an important goal of monitoring.

Espinet and Shoemaker have demonstrated that a combination of a three dimensional simulation model (TOUGH2, multi-phase) of CO₂ can be used with nonlinear optimization methods and with the SOARS uncertainty quantification method to obtain very credible estimates of current spatial distribution of the CO₂ and to get good predictions of its future migration (Espinet & Shoemaker, 2012; Regis & Shoemaker, 2007; Bliznyuk et al, 2011). This is a very difficult problem because the simulation models for a reasonably sized CO₂ geological storage site (and other multiphase subsurface fluid problems) are highly nonlinear and can take many hours -- or days -- for one simulation. Hence, improved optimization and uncertainty quantification methods are required to obtain good solutions with relatively few simulations.

Managing Water Availability and Quality

Approximately 1 in 6 people living today do not have adequate access to water, and more than double that number lack basic sanitation. Indeed, lack of clean water is responsible for more

deaths worldwide than war. Providing access to clean water is an engineering challenge of the highest priority. The goals are to protect water from pollution so that it can be used to meet human and ecological needs and to ensure that the quantity of water is adequate, not excessive. Insufficient quantity of unpolluted water leads to disease, agriculturally and industrially devastating drought, and loss of valuable renewable hydropower. An excess of water leads to floods, loss of life, and destruction of infrastructure and ecosystems. Such questions involving social welfare, especially the interface with technology, has provided fertile grounds for OR applications in the past, and we expect its impact to continue in the context of this world-wide challenge.

Most environmental studies of pollution and ecological health are based on either mathematical or computer simulation models known as “flow and transport” models. One instantiates such models by fitting parameters to data (e.g., soil characteristics, weather, pollutants), and using the output of the model to predict the dynamic and spatial distribution of water and water-borne pollutants. These predictive models can then be used in a variety of ways. With regards to the management of water resources, we can use flow and transport models to study the availability of clean water for human or ecosystem uses. Such studies can identify where one might remediate contaminated groundwater. Further, OR tools allow the study of trade-offs between alternative remediation policies, some of which may cost more but deliver greater quantities of clean water, while others may be less expensive, but deliver lower quantity and quality of water.

While such methodology is not foreign to the water resources community, the large scale and computationally demanding nature of these models make them worthy of collaboration. To give a sense of some of the challenges, consider a case of designing a network of water purification and delivery. Given that such processes can cost over a billion dollars to setup, they must be placed strategically, and the network itself must also be designed to be cost effective. The question of delivering clean water can be posed as a decision (optimization) model to make some discrete choices (e.g., location and network design) so that adequate quantities of clean water (modeled by continuous variables) can be delivered cost effectively. Such models, leading to nonlinear mixed integer programs, can be notoriously difficult, and OR expertise is likely to identify computational bottlenecks which will ultimately reduce costs while increasing the quantity of clean water to make headway into a currently intractable problem.

A major challenge and recurring theme in this setting is the need to solve, design, and control problems in the presence of different forms of uncertainty. Design of water purification technologies, the location of purification facilities, the allocation of sensors and metering devices, and the management of water reservoirs all represent different forms of stochastic

optimization problems which have to be solved in the presence of different forms of uncertainty. General-purpose, reliable algorithms for these problems do not exist.

The Power Grid⁴

Years of stable growth, heavy regulation and marginal profits have produced a power grid that is designed to meet the needs placed on it. Utilities are prepared to handle the modest uncertainties in our ability to forecast temperature, the occasional failure of a generator, and the predictable daily and seasonal peaks. However, strict regulatory oversight has limited the ability of grid operators and utilities to invest in assets without a clear public need. The result is a network of generators and transmission and distribution lines designed to serve these previously stated needs in a robust way, but nothing more.

Electricity is becoming the new money and the way that different sources of power generation can be matched to the wide range of needs in society. Rising electricity demand from the explosion of electronics is being compounded by the anticipated tidal wave of electric vehicles, the only technological option that will allow automakers to meet the mandated 54 MPG CAFE standards. Recharging a single electric vehicle is equivalent to the load of a suburban home, and the existing grid is simply not scaled to handle this load. Rising temperatures are going to further increase loads from air conditioners. These rising demands are pushing the grid closer to capacity and making it more vulnerable to routine variations in generation, transmission, and load.

The robust power grid needs to manage four fundamental elements: generation, with an increasing presence of uncontrollable renewables; the transmission grid that moves power from generators to load serving regions (illustrated above); the distribution grid that delivers power to customers; and the customers themselves who can, in a limited way, be “managed” to adapt their demands to short term capacity shortages. It is only in the last 10 years that we have evolved powerful implementable tools to design and manage a grid *without uncertainty*. Optimizing grids for a single scenario is difficult and can require hours of computer time. Standard modeling practice requires that grid operators maintain a grid capable of sustaining a *single failure*, a problem that might be 1000 times harder than optimizing a grid that assumes no failures. Considering the possibility of two failures would be a million times harder.

Beyond the difficult design problems remain the need for operational controls. All grid operators in the U.S. manage their grid using day-ahead and hour-ahead markets, with additional adjustments being made every five minutes and in real time (regulation). Demand

⁴ This section is not directly mentioned as a Grand Challenge. Nevertheless, the grid is attracting its fair share of attention, and integration with renewable resources remains a challenge today.

response markets will introduce the dimension of planning 2, 4, 6 or more hours in advance as building managers and manufacturers introduce varying notification requirements.

We are not even close to having the OR tools that can handle these large-scale, multistage stochastic optimization problems. Grid operators and utilities have been forced to adopt conservative policies that *over-engineer* the grid for routine uncertainties, yet *underestimate* the impact of severe problems. We lack the tools for even basic problems that require handling rare events, such as controlling a battery storage device to minimize the risk of not covering loads. Sandia National Laboratory is undertaking a major research project to solve the day-ahead stochastic unit commitment problem to handle just *50 scenarios* which is a gross simplification of the many events that might happen. A similar model at Livermore is considering just 6 scenarios. Yet, we cannot even find optimal policies for battery storage devices under realistic conditions. We need a fundamental rethinking of our modeling and algorithmic strategies to solve these large scale problems.

Social and Economic Sciences for Sustainability

Much of the emphasis in sustainability efforts relates to social interactions for which OR models can provide guidance. Basic economic reasoning suggests, for example, that solar energy prices will decline relative to the prices of other energy resources if new solar technologies allow for greater efficiency than currently exists or if the capital cost of producing energy with current solar technologies is reduced. Either of these outcomes requires investment in new conversion technologies or in the devices and manufacturing processes for current technologies. Given these alternatives, solar energy should become more economical if investment increases in developing new technologies of either form. OR methodology can help in assessing the value of policies to encourage such investment. In particular, OR models of interactions between the economy and the energy sector can be used to assess the effectiveness of policy alternatives by evaluating overall social costs measured, for example, in changes in GDP growth, effects on specific industries, or trading imbalances. Such OR models have a long history of informing government policy as in "Waiting for the breeder" (Manne 1994) which determined that the cost of waiting to implement any policy action to encourage development of the breeder reactor was very low. Similarly, equilibrium models for energy production and consumption have a long history from the late 1970's and have evolved into more recent policy questions such as the economic impact of greenhouse gas emissions (Manne, 1979; Manne and Richels, 2001).

The second sustainability challenge of improvement in the power grid can be most effective if individual consumers can learn to modify behavior to shift consumption in response to current conditions. OR models can be used to evaluate responses to policies like dynamic pricing. The

potential effectiveness of fusion power also depends on individual reactions as observed in Japan and Germany's decisions to abandon nuclear power. OR models of social decision making built on underlying social science models of individual preference and political organization can help in assessing these decisions.

Carbon sequestration and management of the nitrogen cycle present additional opportunities to implement OR models in addition to direct analysis of policies promoting investment in these areas. In particular, OR models can evaluate the cost tradeoff between management of, and potential damages from, unintentional carbon releases. They can also analyze the impact of more efficient sequestration on retarding development of non-carbon energy resources. In terms of nitrogen cycle management, OR models can also assess the overall impact of shifting agricultural activity such as even greater rates of urbanization as rural agricultural labor declines.

Lastly, OR models can help provide access to clean water by aligning policies with long-term goals. The allocation of water rights and policies to manage those rights again raises the issue of balancing the preferences of individuals with those of society as a whole. OR models can analyze these decisions through representations of preferences and the mechanism of coordination.

IV. OR for Security

The infrastructure of our nation, and the world, is increasingly interconnected. The ongoing integration of our systems – including transportation, energy, water, communications, finance, and more – has been central to their growth in scale and reach and has facilitated increases in their functional efficiency. At the same time, these interdependencies make our systems more vulnerable to both unintentional hazards and intentional threats. OR researchers and practitioners build operational models of such systems precisely because we know that the system’s performance can depend, often in surprising and subtle ways, on how various components and subsystems interact. As our interconnected systems grow in complexity, having a trusted operational model is even more essential for assessing system vulnerabilities and, in turn, addressing the challenge of how to secure that system. OR is ideally positioned to address the following challenges: formulating operational models of suitable fidelity; understanding vulnerabilities using appropriate models for chance hazards and malicious attacks; and allocating scarce resources to best secure the system.

Restore and Improve Urban Infrastructure

The U.S. Department of Homeland Security identifies our nation’s critical infrastructure as involving the following sectors:

- Energy
- Communications
- Banking and Finance
- Healthcare and Public Health
- Emergency Services
- Water
- Dams
- Transportation Systems
- Postal and Shipping
- Critical Manufacturing
- Food and Agriculture
- Information Technology
- Chemical Industry
- Defense Industrial Base
- Nuclear Reactors, Material, and Waste
- Commercial Facilities
- Government Facilities
- National Monuments and Icons

OR has had a long association with the planning, design, construction, monitoring, and maintenance of several infrastructures. In fact, many Civil Engineering departments are home to OR talent, both among students and faculty. These strong ties are also reflected in editorial boards of journals devoted to infrastructure research, such as *IEEE Transactions on Power Systems*, *Telecommunications Systems*, *Transportation Science*, *Water Resources Research*, and others. Members of the OR community have pioneered many of the advances in these areas, including electricity markets, routing and flow control in communications and transportation, and water supply. More recently, OR models have been instrumental in designing power-aware sensor networks (for structural health monitoring of pipelines, bridges etc.), traffic signal

control, and methodologies for risk assessment and decision-making. Over the years, these systems seemed so reliable that we rarely thought that these could indeed be vulnerable. However, as infrastructures have aged, and geopolitical climate has changed, the path to the future is no longer as predictable as it used to be. Because of this re-orientation, the major focus of this section is on security aspects of infrastructures since this is still an emerging area of research.

We can view a principled OR approach to securing these systems of critical infrastructure as consisting of three steps. First, we must understand how these systems operate. Importantly, we must have a working model, not only of how the system operates under nominal conditions, but also of how it will operate if any subset of its components is degraded or disabled. Given the complexity and scale of our infrastructure, we cannot enumerate a handful of threat scenarios of how a system might be disrupted and expect it to span what we must consider. Rather, we require an operational model that we can systematically query.

The second step in securing our infrastructure requires that we understand its vulnerabilities. To do so, we must model the threats to our systems and distinguish unintentional hazards (e.g., hazards from natural disasters, malfunctions, and human errors) and intentional threats (e.g., threats from vandals, criminals, saboteurs, and terrorists). Operations research literature is replete with probabilistic models appropriate to represent unintentional hazards. And similarly OR tools of game theory and adversarial models can capture the capabilities of an intelligent, informed, and determined attacker.

Given a model, or family of models, from the first two steps, the third step involves allocating scarce resources to best enhance system security. OR is particularly well suited for formulating models which can assess and guide choices of system-level actions.

Key to addressing the challenges in securing infrastructure is recognizing that components in our systems *will* fail and our infrastructure *will* be attacked. We cannot enhance component reliability, and we cannot deter determined adversaries to the point where failures and attacks disappear. The foremost goal in securing infrastructure is that our systems be resilient; i.e., after a system has been degraded, operation of that system will adapt to its new configuration and should be able to do so with minimal loss of system capability.

The notion of taking a system-level view and of aspiring to resilience in the face of inevitable hazards and threats is well recognized by the U.S. National Security Strategy (US White House, 2010):

As we do everything within our power to prevent these dangers, we also recognize that we will not be able to deter or prevent every single threat. That is why we must also enhance our resilience—the ability to adapt to changing conditions and prepare for, withstand, and rapidly recover from disruption. [p. 18]

...we recognize that the global systems that carry people, goods, and data around the globe also facilitate the movement of dangerous people, goods, and data. Within these systems of transportation and transaction, there are key nodes—for example, points of origin and transfer, or border crossings—that represent opportunities for exploitation and interdiction. [p. 18]

When incidents occur, we must show resilience by maintaining critical operations and functions... [p. 19]

We must, therefore, strengthen public-private partnerships by developing incentives for government and the private sector to design structures and systems that can withstand disruptions and mitigate associated consequences, ensure redundant systems where necessary to maintain the ability to operate, decentralize critical operations to reduce our vulnerability to single points of disruption, develop and test continuity plans to ensure the ability to restore critical capabilities, and invest in improvements and maintenance of existing infrastructure. [p. 19]

OR literature includes significant work for securing infrastructure with a seminal paper by Brown et al. (2006) describing the modeling framework we have sketched and applying that framework to improving the security of the U.S. Strategic Petroleum Reserve; the U.S. Border Patrol at Yuma, Arizona; and an electric power grid. A related survey paper, Brown et al. (2005) includes further applications in securing a network of oil pipelines and the DC-Metro system and improving airport security in Los Angeles (LAX). See Pita et al. (2008) for more on security at LAX, Wein and Liu (2005) for vulnerabilities in the milk supply chain, and Murray et al. (2010) for securing municipal water systems. Further work discusses appropriate roles for adversarial versus probabilistic models (Golany, et al 2009; Brown and Cox 2011; and Cate 2012).

Despite the success of this line of research, our foremost challenges in securing infrastructure remain. Ongoing work involves developing very high-fidelity simulation models of systems of critical infrastructure (Brown, 2007). Modeling the interdependencies of these systems of vast physical scale and detail and of disparate time scales is truly challenging. Additionally, populating such models with appropriate data can be a challenge because the vast majority of the U.S. critical infrastructure is privately owned and operated. Success with such challenges helps build a “trusted model” in the sense of our first step above. However, an all-inclusive operational model typically does not have enough structure to lend itself to rigorous application of our second step to assess system vulnerabilities, let alone to application of the third step to optimize system security. Understanding and communicating to the stakeholders which differences in levels of detail in an operational model make a difference when it comes to understanding system vulnerabilities remains a deep challenge.

Prevent Nuclear Terror

The efforts of terrorist organizations and rogue nations to obtain nuclear material and technology to produce a nuclear weapon are well documented (Bunn, 2010). The United Nations’ International Atomic Energy Agency (IAEA) maintains an Illicit Trafficking Database.

From 1993-2011, over 2000 incidents of unauthorized possession of nuclear and other radioactive material were reported. Sixteen of these incidents involved highly-enriched uranium (HEU) or plutonium (i.e., weapons-grade material). The IAEA reports that in cases where such information is available, the majority of these incidents involved traffickers attempting to sell illicit material. That said, motives of nuclear smugglers will likely change as the material changes hands from its origin to its intended destination.

The Domestic Nuclear Detection Office (DNDO) is part of the U.S. Department of Homeland Security. DNDO is charged with developing the Global Nuclear Detection Architecture (GNDA). Doing so requires coordination across multiple federal agencies including the U.S. Departments of Energy, Defense, State, and the Nuclear Regulatory Commission. As part of the GNDA, the DOE's National Nuclear Security Administration (NNSA) works with other countries to

Deter, detect, and interdict illicit trafficking in nuclear and other radioactive materials across international borders and through the global maritime shipping system. The goal is to reduce the probability of these materials being fashioned into a weapon of mass destruction or a radiological dispersal device ("dirty bomb") to be used against the United States or its key allies and international partners (National Nuclear Security Administration).

Concern with this threat predates 9/11/01. The NNSA program has its origins in 1998, when the U.S., with the Russian State Customs Committee, launched a program that included placing radiation portal monitors (RPMs) at Russian customs checkpoints to deter smuggling of nuclear material out of Russia. That program has since expanded to border crossings and sea ports across the globe.

Much of the work associated with deploying the GNDA has involved NNSA and DNDO installing RPMs at international and domestic seaports, airports, rail, and road border crossings. Customs and Border Protection Officers are equipped with mobile radiation detectors. DNDO is developing other mobile detectors that can be deployed in so-called surge operations informed by intelligence reports. Further initiatives aim to secure cities and address challenges of detecting smuggling attempts between authorized ports of entry including smuggling attempts via small maritime craft and general aviation (USGAO).

Conventional wisdom seems to be that obtaining weapons-grade material may be the most serious hurdle faced by a would-be nuclear terrorist:

It should not be assumed that terrorists or other groups wishing to make nuclear weapons cannot read. – Richard Garwin and Georges Charpak

With modern weapons-grade uranium...terrorists, if they have such material, would have a good chance of setting off a high-yield explosion simply by dropping one half of the material onto the other half... even a high school kid could make a bomb in short order. –Luis Alvarez

For this reason, the NAE report has chosen to focus on nuclear material, outlining the challenges as follows:

1. How to secure material
2. How to detect material, especially at a distance
3. How to render a potential device harmless
4. Emergency response, cleanup, and public communication after explosion
5. Determining responsibility for an attack

1. How to secure the material: Part of securing nuclear material requires a model of the nuclear material supply chain, along with modeling and analysis for physical security and a secure inventory. The IAEA inspects, sometimes randomly, state nuclear programs in an attempt to ensure material is not being misused or diverted, and OR models should play an important role in the timing and nature of these inspections. If it is determined that a rogue nation has developed an illicit nuclear program, the international community has at its disposal various options including diplomacy, embargoes, poaching, sabotage, and military strikes. Understanding how to best interdict such a program relies heavily on OR (Brown et al. 2009; Harney et al. 2006).

2. How to detect material, especially at a distance: Much of the required work in improving detection capability involves technological breakthroughs for better detectors. Shielded HEU is difficult to detect with current technology, and detection at a distance is even more challenging. However, there are important opportunities for better detection algorithms (multi-spectral data analysis of radiation at different energies) and for deploying systems of detectors in multi-layered defense around cities, at border checkpoints (in the U.S. and abroad), and in developing inspection protocols at ports. As indicated above, further challenges exist in systems-level analysis for securing borders between ports of entry. A number of OR researchers have made contributions in these areas (Atkinson et al. 2008; Boros et al. 2009; Dimitrov et al. 2011; Gaukler et al. 2011; Hochbaum and Fishbain, 2011; Madigan et al. 2007; McLay et al. 2011; Morton et al. 2007; Wein et al. 2006 & 2007).

3. How to render a potential device harmless: The U.S. nuclear arsenal has an elaborate system of safety technology designed to prevent accidental and intentional but unauthorized detonations (Elliot, 2005). The U.S. has worked to down-blend, or secure, weapons-grade nuclear material at storage sites in Russia and other former Soviet States, and OR tools could inform how to best prioritize such activities. The U.S. has dismantled its most powerful nuclear weapons, and OR played a role in scheduling the dismantling process at the Pantex site (Asgeirsson et al, 2004).

4. *Emergency response, cleanup, and public communication after explosion.* This challenge has much in common with the discussion in the previous section regarding resilience of critical infrastructure. Here, we require modeling of interconnected systems of critical infrastructure having incurred a massive disruption. This challenge is ideally suited for OR tools for precisely the reasons we sketch in the infrastructure section. Communicating recommended policies to the public, e.g., shelter-in-place, requires OR analysis.

5. *Determining responsibility for an attack.* The foremost challenge in attribution concerns nuclear forensics. The goal is to facilitate identification of the source of the nuclear material and/or the attacker if the device captured is pre-detonation. Failing that, the goal is to identify the attacker by collecting and analyzing radioactive debris post-detonation. There are important opportunities to complement such an analysis using OR tools to carry out “systems-based forensics.” For example, given limited information regarding a captured nuclear smuggler, we may be able to infer (by an “inverse problem”) his/her origin and intended destination. Social network analysis is another example of how OR models may complement material-based forensics.

As indicated by the references given above, the OR community has already made contributions to this area. However, the challenge of preventing nuclear terror has engaged a small number of OR researchers and there are significant opportunities to have greater impact. Moreover, in addition to nuclear and radiological threats, analogous challenges exist regarding chemical and biological attacks.

Secure Cyberspace

Our information and control systems are deeply embedded throughout virtually all of our critical infrastructure systems. These systems have been primarily designed for efficiency, and security is typically included as an afterthought. However, security is increasingly becoming a central concern. This characterization is especially true for the operating systems and applications that run on the computing devices we use daily. We literally “patch” security holes as they are revealed. However, major critical systems, e.g., the power grid, have come into sharp focus, especially as rampant international hacking has become more visible.

The field of OR is in an ideal position to close the gap between efficient operation of infrastructure systems and the ability of these systems to be robust to cyber-attacks. The OR community has been active in a variety of emerging infrastructures including defense, health care, transportation, and of course the internet. This wide reach positions the OR community to provide the kind of understanding of system complexity identified in the NITRD⁵ strategic

⁵ Networking and Information Technology Research and Development

plan of 2012 which calls for a *“new system science ... to provide unified foundations, models, tools, system capabilities and architectures that enable innovation in highly dependable cyber-enabled engineered and natural systems.”* This new science that the report promotes appears to be at the intersection of three perspectives: OR Models and Algorithms, Systems and Controls, and Computer Science. This multi-disciplinary approach has already started to take root, and we anticipate that this convergence will provide the foundations for cybersecurity.

Cybersecurity threats range from individual criminal hackers to organized criminal groups to terrorists to nation states. Adversaries can disrupt our critical infrastructure on a massive scale; for example, China and Russia are said to have “probed” the U.S. electric power grid. Clearly, traditional “perimeter” defenses, like firewalls, are not sufficient. Perimeter defenses are eventually penetrated or otherwise bypassed. They do not deal with denial-of-service attacks, and they fail to deal with adversaries already inside the perimeter. For this reason, a recent report by the Center for Applied Cybersecurity Research at Indiana University suggests that the most practical avenue for research in enhancing cybersecurity is Cyberintelligence (Cate, 2012). As with improving intelligence in the physical world, improved cyberintelligence involves better resources for gathering and processing information, often from non-technical sources.

Given the inherent vulnerabilities of cyber defenses, increased attention has been given to robust design. For example, the power grid has been characterized by large central generators that can economically provide power to large populations over the grid. However, this architecture then becomes vulnerable to cyberattacks (as well as other disruptions such as storms). An alternative strategy is to consider distributed generation which allows micro-grids and regional grids to remain up and running as other parts of the grid recover from disruptions. The proper design of a distributed generation network is precisely the type of network design problem long addressed by the OR community.

Other reports outline four fundamental themes for Cybersecurity R&D (Federal Cybersecurity R&D).

1. Designed-in Security (Reliability-based software design) (Raheja & Gullo, 2012; Singpurwalla & Wilson, 1994)
2. Tailored Trustworthy Spaces (Interactions between negotiating agents and verifying reputation) (Raheja & Gullo, 2012; Ardabili & Liu, 2012)
3. Moving Target (Non-deterministic system design to thwart attacker attempts to learn and predicting attacker behavior) (Moving Target Research Symposium)
4. Cyber Economic Incentives (Creating a cyber-community that provides incentives to adhere to cyber-security standards) (Ardabili & Liu, 2012)

While concepts for software reliability (item 1) have been studied in the OR community for decades [48], the precise needs for designing reliability into systems facing cyberthreats is a relatively uncharted territory. As a result, we expect that collaboration between OR groups and practitioners of cybersecurity will guide the manner in which security is designed into cybersystems. As for items 2 and 4, OR is already providing the basis for new approaches. For example, mechanism design (an economic concept with significant traction in the OR community) is studied in Ardabili & Liu (2012) within the context of creating trustworthy networks. As for item 3, one of the more powerful tools is the framework of Markov Decision Processes (MDP). An optimal randomized policy can be generated via an MDP; this provides a principled approach to item 3 while keeping the adversary guessing. Moreover, we draw attention to one particularly inspirational paragraph of the report related to the complexity of cyberspace (Federal Cybersecurity R&D).

“Deep Understanding of Cyberspace. To operate effectively as a moving target in cyberspace, we must understand our systems state, be aware of our surroundings, know the soundness of the structures on which we rely, and know what is happening around us. Cyberspace is complex, and moving target techniques will increase that complexity. Actions in cyberspace are instantaneous. If we are to manage our moving target capabilities effectively and instantaneously in the face of this complexity, we must greatly enhance our ability to monitor, model, analyze, and understand our own system, the systems in cyberspace with which it interacts, and the threat environment at that point in time. If we are to make these decisions within the tight time constraints of cyber actions, we must greatly enhance the speed of our complex analytics and tighten our feedback loops. Ultimately, we must provide knowledge-driven systems that remove the human from the loop in many system decisions. But for those decisions that do require human decision-making, the combination of high complexity and short processing time strains human cognitive processes, so we must provide novel methods of presenting information, directing attention, and navigating between analytics at different scales. We must also provide capabilities that enable deep, not just comprehensive understanding of cyberspace. Our methods must enable us to view the situation from alternative points of view and to get below surface indicators to determine causes and conditions.”

The type of analytics mentioned above highlights the interplay between data and decisions, calling for greater infusion of OR into cybersecurity. The challenges of cybersecurity rely on the ability to bring a variety of processes and resources together in a timely manner to enhance our ability to deter adversaries, even if they are intent on continuous barrages of attacks. Current statistical approaches for intrusion/anomaly detection often focus on attacks at a network node (see Polunchenko and Tartakovsky 2012, Polunchenko et al 2012). However, denial-of-service

attacks often cascade through a network, and network-wide (i.e. higher dimensional) models become necessary. The OR community, especially the Military OR community, has been a source of many ideas behind such cascading threats. This experience, together with OR's close ties to Computer Science should provide a fertile ground for collaboration among these communities to face this important security challenge.

Social and Economic Sciences for Security

Over the past four decades, OR has maintained an influential presence in the area of aviation security by incorporating human behavior within OR models (Salter, 2008). In particular, problems arising in passenger screening have been modeled by incorporating human behavior within stochastic modeling, optimization, and control (McLay et al, 2006; Lee et al, 2009). In other security areas, effective prevention of nuclear terror should rest on a model of terrorist organization and objectives built around OR principles. Cyber security involves similar issues and models of the cost in terms of disrupted individual rights and privacy. All such analyses would be founded on OR's recognition of the complexity of networks and the interconnectivity of all social agents.

The examples above serve to illustrate the importance, and the challenges, of melding operations research and social science approaches to advance progress on the Grand Challenges. Progress critically depends on the decisions of individuals, groups, and organizations that will determine a number of things: the resources devoted to address these challenges; the deployment of those resources in meeting the challenges; the acceptance, use, and spread of the new technologies developed in these efforts; and the impact of the new technologies' use on other aspects of society. Models of decision making and understanding of key social science concepts are essential for analyzing each of these phases of development and implementation. As the science of decision making, OR can play a significant role in making a case for investment in new technology, in assessing and designing policies to encourage the effective allocation of resources consistent with progress on the challenges, and in evaluating the uses of candidate technologies and the consequences of those uses.

V. OR for Human Health

In 2010, total health expenditures in the U.S. reached \$2.6 trillion -- \$8402 per person and 17.9 percent of the U.S. Gross Domestic Product.⁶ By any standard, this is a staggering amount. Expenditures continue to rise both in absolute terms and as a percentage of U.S. spending. Few – if any – believe the health care system is operating efficiently. Both improved treatment and reduced expenditures are possible with proper analysis. In many cases, this requires studying the system or parts of the system as a whole. In other cases, individual procedures, treatments, and medicines can be improved through analytical insights aided by data. In either case, the discipline of OR brings ideal tools to aid experts in traditional domains of the nation’s health care system.

Advanced Health Informatics

Coupling health informatics and analytics can enhance the quality and efficiency of medical care as well as the response to widespread public health emergencies. Comprised of a collection of old technologies (paper) and emerging ones (digital), today’s medical records are in transition. Even with computerized records, standards for storage and retrieval are incompatible between databases, software, and platforms. The lack of interoperability greatly complicates the sharing of information. Engineering future systems that will share data seamlessly, discover patterns, and offer relevant decision support constitute the holy grail of health informatics.

Another aspect of health informatics is concerned with preparedness for pandemics and chemical and biological attacks. Worldwide surveillance and detection monitoring for early signs of attack or onset of an infectious disease promises to lead to good early warning systems that will allow timely response. Rapid diagnosis will require a system that can analyze and identify the agent of harm and track its location and that can efficiently and economically monitor the environment to locate the source of the attack or disease. In its segment on pandemics, the NAE report on Grand Challenges notes that “the usefulness of these approaches depends on numerous variables — how infectious and how deadly the virus is, the availability of antiviral drugs and vaccines, and the degree of public compliance with quarantines or travel restrictions. Understanding the mathematics of networks will come into play as response systems must take into account how people interact.” In some cases, countermeasures may also call for rapid production of antidotes and vaccines. Planning for such large-scale on-demand manufacturing, logistics, and health care delivery is not an activity that can be successful without support from the OR community.

⁶ U.S. Centers for Medicare and Medicaid Services.

Engineer Better Medicines

Engineering better medicines will require collaboration between engineers, biologists, and health professionals. Ideally, we will be able to diagnose and treat people based on their individual differences, a concept commonly referred to as "personalized medicine." Many exciting opportunities for Operations Research lie in the development and optimization of therapeutics and diagnostics, especially in collaboration with scholars in the field of biomedical engineering. There is a growing trend in the development of smart medical "devices": sensors gather physiological data from patients; algorithms calculate treatment decisions, and medical devices actuate these decisions. Devices that utilize this paradigm are already in the market. OR techniques can be used to optimize such medical devices as well as the systems that will adopt them. At the micro level, exciting research questions may include identifying the best ways to gather and transmit the data and developing treatment algorithms. At the macro level, research questions include: determining rules for accessing such systems, designing reimbursement systems that reward providers for appropriate use of such devices, and determining how such devices will affect capacity planning and resource planning for health care systems.

There are also opportunities for optimizing individual diagnostics or treatments. Medical diagnostics typically involve gathering physiological data and translating the data into a treatment diagnosis. With the explosion of sources for medical data (e.g., genomics data, data from electronic medical records, imaging data) OR can contribute to the development of optimization techniques that synthesize data from diverse sources and reach diagnostic decisions that are more sensitive and specific than existing alternatives.

There is also a growing trend of treatment modalities in which the provider uses real time (or almost real-time) imaging data to design the treatment protocol and approach. Illustrative examples come from radiology, cardiology, nephrology, and many other disciplines. OR techniques can be used to optimize the treatment approach based on advanced imaging data.

Reverse-Engineering the Brain

The brain is a powerful, intricate, remarkable organ, interpreting the many signals we receive through our senses and controlling how our bodies act in response to those signals. As important as the brain is, our understanding of how it operates is limited. We can set a broken bone, but when a brain is damaged we can rarely fix the problem. Computers can perform many useful tasks, but they fall far short of what the human mind can achieve. We easily engage in casual conversation, but computers have trouble parsing the words in a spoken sentence, much less assigning meaning to those words.

The brain poses a special challenge in that it does not operate like existing computers. The brain's architecture of neurons communicating via chemical and electrical signals is very different from the transistors etched into a microprocessor. Understanding how the brain operates will make it possible to develop new treatments for damaged brains and will open the door to building computational devices with extraordinary new capabilities.

Brain machine interface: Recent work has demonstrated that electrical signals from the brain can be used to control artificial limbs. For example, Velliste (2008) connected electronic leads to the brain of a monkey who, *by thinking* quickly and repeatedly, manipulated an artificial arm to reach marble-size pellets of food and place them in its mouth.

At the root of this technology are algorithms for interpreting brain signals and converting them to particular movements of an artificial limb. Analyzing data, searching for patterns, and developing optimal responses to brain signals is the domain of OR. Interestingly, learning is a two way process. Algorithms learn from the firings of the brain while the brain learns that certain neuronal firings lead to particular actions. ("When I think *this*, I observe the hand moving to the left, and therefore thinking *this* leads to the desired action.") OR can be used to analyze how neuronal firings adapt in the learning process.

Powerful new computers: Although artificial intelligence (AI) has made reasonable progress in specific domains – game playing and speech recognition, for example – the early promise of AI has failed to materialize. We have yet to create an information-processing machine that comes close to emulating the capabilities of the human brain.

Expectations are that understanding how the brain processes information will allow us to reproduce what it does in a non-biological medium (presumably, algorithms on a computer, though the entire concept of "computer" may be radically different than those of today). Working under the theory that the brain operates as a system and cannot be understood simply by looking at individual neurons, considerable research emphasis is being placed on mapping neural connections throughout the brain (see, for example, the website of the Human Connectome Project, www.humanconnectomeproject.org). Recently, the NIH has launched the Brain Activity Mapping project to record virtually every spike of every neuron in the hope that scientists will be able to discover correlations between spikes, and discover emergent properties of something as complex as the brain (Alivisatos et al, 2013).

Improving Quality while Reducing Costs⁷

At a high level, every health care system is characterized by two performance metrics: cost and quality. OR can help characterize the efficient frontier of health care delivery by seeking answers to questions like: which health care systems lie on the efficient frontier? how do they make the trade-off between quality and costs? Is there a universal efficient frontier for all diseases and treatments, or does there exist a different efficient frontier for every disease and treatment? What are the characteristics of systems on the efficient frontier and what makes them efficient? Are they efficient in the inputs they use (i.e., the medical technology) or how they transform the inputs into outputs (i.e., their operations), or both? Does the answer depend on where these systems lie on the efficient frontier and can practices adopted by systems at one part of the frontier be adopted by systems at other part of the frontier and essentially shift the frontier upwards?

Healthcare delivery in today's world typically takes the form of a large and complex system involving expensive resources such as clinics, hospitals, and pharmacies; highly skilled workers such as doctors, nurses, and other medical professionals; and insurers and administrators. Such a system fundamentally differs from traditional engineering systems in its *human centric characteristics*, with ultimate goals of providing high-quality service and customer satisfaction at affordable cost (i.e., the system must be *quality-and-efficiency* driven).

One of the most significant problems facing the health care system is keeping costs under control while providing high levels of service. Doing so requires a careful analysis of costs and benefits, but as Kaplan and Porter (2011) argue, "The biggest problem with health care is that we're measuring the wrong things the wrong way."

At the core of their argument is that health care facilities – hospitals in particular – place their emphasis on departments and individual line items rather than on the patient. For example, if the cost of staffing intensive care is expensive, a hospital might cut back wages in the Intensive Care Unit or limit future hiring. The authors point out that such myopic decisions may lead to increased costs elsewhere in the system that outweigh the benefits derived from a well-funded, properly staffed Intensive Care Unit. Such interactions are the classic mistakes that are made when focusing on components rather than systems. The situation is comparable to traffic engineers in the 1960's fixing a bottleneck by retiming one traffic signal, only to create a new bottleneck downstream.

⁷ The cost of health care in the U.S. is one of the main factors driving changes in the system. Even though, this area was not identified within the Engineering Grand Challenges, this presents an unprecedented opportunity for OR to make a monumental contribution to this area.

Additionally, Kaplan and Porter (2011) propose tracking costs specifically as they relate to patients and procedures. For example, an appendectomy requires time in the operating room, recovery, and follow-up. Patient based costing focuses on the cost of an appendectomy, not on services by department.

The idea is not entirely new. Fetter (1991) presented a classification system for procedures called *Diagnosis Related Groups* (DRGs). The original DRG classification system became the basis for Medicare reimbursement in 1983 (AHIMA, 2006). DRGs allowed the government to exercise some degree of control over rising health care expenditures which until that time was based on a billable list of individual charges. With billable lists of individual charges, hospitals performed whatever work they deemed necessary and sent Congress the bill. DRGs allowed Congress to say that “a well-run hospital should be able to complete an appendectomy at a particular cost, so when an appendectomy is being performed, that’s the amount the hospital will receive.” Fetter’s work on DRGs led him to win INFORMS’ Edelman Prize in 1990 for estimated government savings of over \$50 billion in the years it had been operational (US Centers for Medicare & Medicaid Services).

However, while the U.S. government continues to reap benefits from the use of DRGs, few hospitals have adopted DRG-type methods for tracking costs internally. Kaplan and Porter (2011) are advocates for this type of internal record keeping. OR can help in a number of important ways.

Patient Classification Methods: The number of classification groups in Fetter’s initial set of DRGs came to roughly 500. Since that time, a number of more refined classification systems have been developed by many different sources. Such systems are constantly being refined to reflect advancements in technology and the patient population. OR can be used to define, refine, and analyze classification systems, problems which can naturally be cast as optimization problems.

Variability: Patient classification methods are based on averages (e.g., what are the average resources used by a patient receiving an appendectomy?) A hospital tracking its costs needs to be concerned not just with averages, but with actual patient encounters. Suppose a hospital performs 300 appendectomies. If 90 percent of them are at or slightly above “fair and reasonable average cost” determined by the government, but 10 percent are far above the average due to complications, the hospital will not cover its costs. OR can deal with this variability in tracking costs and predicting future cost performance.

The Resource Management Perspective: Notwithstanding the negative connotation associated with healthcare as a resource management problem, it is important to bring a management perspective into the health care system. Today’s healthcare delivery systems --- hospitals being

primary examples --- have yet to utilize and benefit from the kind of tools and technologies that run modern manufacturing facilities (such as a semiconductor wafer fab), let alone all the research and innovations that have gone into enterprise supply chain management over the last couple of decades. This is an area where OR can most effectively help, in collaborations with researchers and practitioners in public health, medicine, and other related disciplines. Some areas to consider are presented below.

- *Hierarchical planning in hospital resource management:* At an operational level, a well-managed 21st century hospital should be run at a level of sophistication no less than a multi-billion dollar wafer fab. In particular, operational decisions should be structured with three levels of planning:
 1. Aggregated (capacity) planning (with a monthly or quarterly time frame)
 2. Detailed planning/scheduling (on a daily basis)
 3. Real-time execution and disruption recovery
- *Revenue optimization and resource control:* A central issue in healthcare is how to allocate resources (such as physicians and surgeons, expensive diagnostic facilities/procedures, wards and beds, ICU) among patients of different categories so as to achieve the optimal tradeoff among various, possibly conflicting, objectives: to maximize the value of care/service and to deliver care/service in the most efficient manner while maintaining costs at an affordable level. Research in airline revenue management and call-center staffing has amply demonstrated the effectiveness in using sophisticated analytical techniques to strike the right balance among quality of service, revenue goals, and equipment/staffing costs. It will be useful to explore the application of similar ideas, principles, and approaches to healthcare delivery (e.g., operating room scheduling, bed allocation, nurse staffing).
- *Ideas/principles/tools in supply chain management* can be applied to managing the *value-chain* of healthcare delivery:
 1. Coordination among patients (consumers), healthcare providers and insurers;
 2. Outcome-based pricing (“performance contracts”);
 3. Studies of integrated practice units (IPUs) and integrated care networks;
 4. Inventory control and logistics management in large hospitals.
- *Network Effects:* While the government and other payers are interested in determining a “fair and reasonable payment” for various procedures, hospitals and other service providers are interested in the actual cost of services at their institutions. This means tracking the actual cost of the procedures and analyzing their performance, not just individually, but in

how they interact with one another in the competition for limited hospital resources. Network analysis of this type is well understood in the field of OR.

- *Business analytics (BA)*: The ideas and tools of BA can be applied to:
 1. Online market places for competing healthcare plans, allowing customers to do comparison shopping, as well as analyzing consumer preferences;
 2. Appointments scheduling for clinical visits, medical diagnostic facilities;
 3. Hospital resource management: operating room scheduling, bed/length of stay management;
 4. Healthcare informatics to improve medical efficiency/reduce costs;
 5. Measuring costs: time-driven activity-based costing;
 6. Measuring value: beyond measuring individual outcomes.

BA can be used to unlock new sources of economic value in healthcare delivery, and provide fresh insights into better management of the healthcare value chain. Yet, grand challenges remain in how to process, in an efficient and cost-effective manner while preserving a high level of security and privacy protection, the ever-growing amount of patient-care data and information, and how to best utilize these for better decisions.

Social and Economic Sciences for Human Health

Health challenges involve many of the same issues that arise in other areas such as designing and assessing policies to promote investment. In addition, OR can address general issues such as the impact of rationing care through prices, waiting lists or other priority systems. This is particularly relevant with advanced health informatics that enables both customized care and individualized payment or priority. The cost and impact of potential losses of privacy with increased data availability can also be assessed with OR. In addition, OR can help measure the effect of increasingly expensive custom-engineered medicines on other outputs of society. Finally, models of social interactions and individual incentives are essential for reverse engineering the brain because as psychologists have observed human behavior cannot be isolated to an individual and must always be considered in the relevant social context. One aspect of social science which might be in line with “reverse engineering the brain” is discussed below.

One of the most fascinating hypotheses in cognitive science is that some decision-making tasks can be well modeled using Ornstein-Uhlenbeck processes, with certain thresholds providing reasonable estimates of decision-times. These models, which are collectively referred to as Decision-Field Theory, raise the possibility of answering an important question: how does the brain make decisions? (Busemeyer & Diederich, 2002). Given that neurons are modeled using

diffusion processes, the brain can be modeled as an enormous network of diffusion processes which collectively implement decision-making as a continuous time stochastic process. Whether this conjecture contains any validity is unclear. But scientists like Henry Markram at EPFL (Lausanne, Switzerland) and others at IBM who are supplying a Blue Gene machine for the Blue Brain project should be able to allow OR researchers to perform experiments which shed light on such conjectures.

VI. OR for Joy of Living

Of the four broad categories of challenges, this area is clearly the most nebulous for scientists and engineers. In connection with the NAE Grand Challenges, we have interpreted “Joy of Living” as consisting of three core problem areas: Advanced Personalized Learning, Enhancing Virtual Reality, and Engineering the Tools of Scientific Discovery. However, one recognizes that “Joy of Living” encompasses a much broader class of challenges dealing with improving quality of life on a daily basis. For example, reducing traffic congestion in urban areas, improving response times of first-responders, designing smart, energy efficient homes, and others raise many novel OR questions. One such example is an application related to predicting movie recommendations associated with the so-called “Netflix Prize” problem. This problem is concerned with a matrix whose columns represent “user names”, and the rows represent “movie names.” The problem of predicting which movies should be recommended to a user can be formalized as a “matrix completion” problem of inferring some entries, based on partial data about movie likes/dislikes of some users. Among the more widely cited approaches to this problem is a method by Candes and Tao (2005) where the term “Dantzig-selector”⁸ was coined for a convex (linear) programming formulation of this problem. Other “Machine Learning” approaches routinely draw upon optimization as a core technology for inference. Other “joys of life,” such as sports, have also seen many applications of analytics; in addition to the well-publicized baseball movie “Moneyball,” there is Major League Baseball scheduling which is done routinely using OR models. In this sense, OR casts such a wide net in the “Joy of Living” area, that the following subsections (pertaining only to the NAE Grand Challenges) explicitly discuss only a small subset of applications for “Joy of Living.”

Advanced Personalized Learning

Personal learning approaches range from modules that students can master at their own pace to computer programs designed to match content presentation with a learner's personality. Thus far, systems have been designed for storing instructional content, delivering it to students, and facilitating the interaction between instructors and learners. Two of the best known tutoring approaches are classified as Cognitive Tutors and Constraint-Based Modeling (Corbett and Anderson, 1995; Mitrovic, 2012). The former (CT) considers a skill to be correctly applied by a student when a rule matches the student's actions. In contrast, a CBM system considers the skill to have been acquired when a predicate is matched correctly. We expect that engineers will be developing more intelligent web-based education systems and teaching methods that optimize learning given the diversity of individual preferences and the complexity of each human brain. Such optimization may call upon joint representations using both integer

⁸ In honor of George B. Dantzig, the father of Linear Programming

programming and constraint programming as surveyed in Bockmayr and Hooker (2005). For a recent survey of methods to model skills, see Desmarais and Baker (2012).

Engineers -- in collaboration with experts from other fields -- should strive to design “recommender” systems that guide individual learning using web-based resources and creating algorithms that adjust recommendations to aid student progress. Recommender systems will optimize the order of presentation using algorithms which eliminate unsuccessful presentation sequences and modify successful ones for a new round of tests in which the least successful are again eliminated and the best are modified once more. By combining machine learning with concepts of discrete optimization, it may be possible to tune recommendations to specific students. For a survey on recommender systems, we refer to Shishehchi et al (2011).

A first-rate education is far more than just a collection of courses completed. The focus of personalized learning has been on adaptive electronic instruction and open access to courses; while these are important, the widespread and cost-effective dissemination of personalized, affordable, and effective education, particularly at the secondary and post-secondary levels, requires the solution of difficult and large-scale systems engineering problems.

Consider this plausible vision of a Bachelor’s degree program in the next decade: While many students continue to desire a degree in a traditional discipline, an increasing number want to create custom degrees without sacrificing rigor or acceptance. In either case, their education includes a combination of on-campus, off-campus, international, online, and experience-based education integrated in a seamless manner and customized to their interests and objectives. This will likely include both hands-on and simulated labs, electronic discussion threads and substantial individual tutoring, and courses that are drawn from their home university and the best that the world has to offer. This is clearly education, not just training.

Allocation of instructional resources and the sequencing of course delivery in a traditional university are relatively straightforward: faculties are hired and assigned courses based on their expertise and on the needs of well-defined degrees in a largely repeating schedule. Even for such a straightforward program, the cost of a college or university education is beyond the means of many families with low or modest incomes. Delivery of the integrated and customized program described above would be cost infeasible without viewing it as the kind of complex sequencing, scheduling and resource allocation problem at which OR excels.

The systems problems inherent in personalized education are even more acute at the secondary level, where major cities struggle to provide their students with an educational foundation suitable for life in the information age while operating under severe budget constraints. For example, any proposal for greater electronic delivery of instruction implies the need for greater access to electronic tools (e.g., laptops, tablets, and internet), access that

cannot be assumed for all urban populations and certainly cannot be provided in abundance by cash-strapped school systems. Effective use of available resources, while providing equity and a high level of customization, will provide some of the most challenging optimization problems that operations researchers have tackled to date.

Enhancing Virtual Reality

Virtual reality (VR) is becoming a powerful new tool for training practitioners, for treating patients and for various forms of entertainment. Virtual reality attempts to re-create an actual experience, combining vision, sound, touch, and feelings of motion engineered to give the brain a realistic set of sensations. Practical applications include working with those who suffer from certain phobias and post-traumatic stress disorders. Another challenge is improving the technology so that dispersed people could seamlessly see, hear, and touch each other and share real objects and equipment; this would be particularly useful for the military and emergency response teams. Already, engineers are creating entire cars and airplanes “virtually” To test design principles, ergonomics, safety schemes, access for maintenance, and more. Engineering challenges are centered on replicating human senses and motion in the virtual world so that participants accept the illusion and respond to virtual events and simulations in realistic ways. The resolution of the video display needs to be high enough with fast enough refresh and update rates for scenes to change like they do in real life and the field of view must be wide enough with realistic lighting and shadow. Touch poses arguably the most formidable challenge in virtual reality; with today's technology, one cannot feel an accidental bump against a virtual piece of furniture. For some uses, gloves containing sensors can record the movements of a user's hand and provide tactile feedback, but somewhat crudely.

Military OR Application: OR specialists may be able to partner with VR experts to enhance a user's experience within a VR system. In addition, by combining OR and machine learning, one may be able to enhance the capabilities of OR models themselves. Consider for example, a training exercise in which military personnel are using VR to employ resources within a synthetic battle environment (e.g. STOW 1997⁹). This might involve dynamic decision-making in requiring deployment of a variety of resources including mobile sensors for assessing threats, deploying personnel to mitigate threats, assigning personnel to specific tasks, etc. By fusing OR tools within such VR environments, the user's learning experience can be far more effective than may be possible without OR support. Alternatively, we may also use OR tools (e.g., inverse problems for learning policies) to learn human-like strategies by observing responses of the military decision maker.

⁹ STOW '97 was a United States Department of Defense (DoD) Advanced Concept Technology Demonstration that was integrated with the United Endeavor 98-1 (UE 98-1) training exercise

Energy Economics Application: In the previous application, OR models were considered as part of the tools supplied within a VR environment. In this application, we present a case in which an OR model learns from the VR environment. Suppose we are interested in surveying the population's preferences for electricity use at peak hours. The Brattle Group has developed a model for demand response by using responses from thousands of consumers in various parts of the U.S (Faruqui and Sergicil, 2009). Such surveys could be conducted in an immersive environment where summer-time temperature and humidity are emulated within a lab/home where consumer choices (e.g. HVAC and other appliance usage) can be monitored within a VR environment. Given that such data is a central piece of OR modeling, VR has the potential to transform the manner in which surveys are conducted for operational models.

Engineering the Tools of Scientific Discovery

A wide range of problems in science and engineering require expensive trial and error experimentation. The goal of this experimentation might be to find a new material to solve the problem of energy storage or solar conversion; it might be to find the best molecular structure to cure cancer; or it might be to identify the best design of a new device.

These experiments can be time-consuming and expensive, and the outcomes can exhibit noise. A single experiment can take a month, and a new initiative may require a year to get started. A scientist may face the situation of tuning the concentration of a half dozen additives, the temperature of a process, and the time required for a material to cure.





With its expertise in optimization and the efficient management of resources, OR offers an exciting technology for managing resources within the laboratory. The knowledge gradient is a strategy that guides experimental design to maximize the rate at which we learn. It uses a Bayesian learning model that takes advantage of the high levels of domain expertise possessed by the scientists who lead this research. Such optimal learning approaches offer a principled scientific method for belief extraction. The applications of this concept are exceptionally broad, with potential for guidance in areas ranging from materials science and drug discovery, to process and device design (Powell and Frazier 2008, Powell and Ryzhov 2012).

Fairness and Privacy in the Age of Big Data¹⁰

With the ever-increasing reach of information and sensor technology, fairness and privacy can become easily compromised. This calls for increasing attention to fairness in decision making for society, legislation, policy, and voting. There is a classic body of research on ordinal ranking/voting for which Arrow's impossibility theorem applies.

¹⁰ As in other sections, we identify issues that were not identified in the Grand Challenges report. This is one more.

However, it is possible to achieve desirable properties by using cardinal decision making. In that context, it is possible to measure fairness in terms of deviation from each individual's score, as well as from each pairwise comparison. The latter is important as pairwise comparisons are insensitive to the individuals' scales and often represent the “true” preferences of the individuals. The common “simple” approach to cardinal ratings is to take the averages. Consider for instance an NSF panel in which 4 individuals participate and provide their ratings as in the following table.

					Total
100001		4.5	4.0	3.5	12.00
100002	5	4	3.75		12.75
100003	5	4.25		3.25	12.25
100004	5	3	3		11.00

Observe that each of the individuals who rated those pairs prefer the first proposal to the second, and the first proposal to the third, yet the average score of the second proposal is highest. Proposal three is the second highest and the first, most preferred proposal, comes in third. This is in clear reversal of the revealed preferences of the reviewers.

This type of “fairness” model that minimizes the total penalties for violating the pairwise preferences of the individuals and the total deviation from their individual scores has been studied for any type of functions in Hochbaum and Levin (2006) and in Hochbaum (2010). However, this is only one type of fairness, and others in other contexts should be modeled and addressed as well (see Balinski M. and R. Laraki 2010).

For instance, the concept of “price of anarchy” has been studied, mostly in the Computer Science discipline. This type of research considers protocols under which individual agents optimize their own utility while maintaining relatively high societal utility. The factor loss in societal utility is roughly interpreted as the price of anarchy. This area of devising incentives -- or in other words, mechanisms under which the actions of optimizing individuals will also benefit society -- is important to investigate, and devise new concepts and models for such issues.

Another issue of growing concern today is the need for “privacy” in data-rich environments. One work along this line was presented by Gusfield (1990) in the limited context of data

presented in 2-D tables. In that paper, the issue was what type of information can be presented in the aggregate so that deducing individual private information is within intervals that are too wide, or NP-hard.

OR researchers are presented with major challenges concerning these two issues of fairness and privacy. The challenge is to devise a quantification scheme and to establish criteria to measure the relative extent of fairness or privacy in decisions and policy making.

Social Sciences, Decision-Making, and Human Behavior

An underlying goal in addressing the Grand Challenges is to increase value for individuals and for society as a whole, both now and in the future. Therefore, Grand Challenge solutions involving technologies, processes, and policies must be evaluated in the context of how individuals and organizations respond to these innovations and what value is ascribed to them. Fundamental to achieving this is a capability to model the behavior, decision processes, and utility functions of individuals and organizations.

Advanced personalized learning and personalized medicine (both described in this report) are examples of strategies to advance progress on the Grand Challenges through the provision of individualized or tailored services. Developing these services is a significant challenge. Massive amounts of data must be harnessed; decision models and approaches that integrate knowledge in science, social science, and management must be developed; and advanced techniques must be invented to understand and predict human preferences, choices, and behavior. By meeting the challenge of developing these capabilities and integrating them within operations research-based decision frameworks, individuals will have advanced capabilities for informed and effective decision-making. These individualized learning and medical services act as personal decision assistants that can produce greater utility for individuals, and simultaneously increasing welfare for society as a whole.

Unfortunately, optimizing value for each individual can sometimes lead to non-optimal results for society. This is observed when multiple stakeholders (individuals or organizations) compete for or are expected to share constrained resources. Take, for example, the case of personalized mobility. With today's information, computing and operations research technologies using models that capture the traveler's transportation preferences, and the congestion levels and corresponding travel-times across the transportation network, personalized utility-maximizing routes and trip departure times can be determined for a traveler. In generating these personalized recommendations, however, if the behaviors of other travelers using the transportation network are ignored, the results will be far from optimal. The issue is that transportation networks, like many systems, are dynamic, with individuals adjusting their

routes as they obtain information about travel times with travel times along paths changing as re-routing decisions are made. Optimizing an individual's routing decision requires knowledge of the behavior of the individual whose choice is being optimized, as well as the knowledge of the behavior of all others in the transportation system at the same time.

This example serves to illustrate the general concept that models of decision making and understanding of key social science concepts are critical for analyzing and improving systems that involve humans. In the context of decision making, social science focuses on the choices that individuals make; the processes that groups use to combine individual preferences to a collective action; and the ways in which organizations at levels from the small firm to the whole of society delegate and exercise decision making authority. OR models add to this understanding by representing the choices or preferences of individuals and the collective value of different outcomes to groups in formal models of the decision making optimizing their own objectives. OR-based models (such as those of Barnhart et al. 2012) can be developed to find solutions on the efficient frontier (as introduced in the OR for Healthcare Section) that identify efficient operating strategies and balance the competing objectives of the system stakeholders and the system itself.

Another example of the conflict between the objectives of society as a whole and those of individuals can be illustrated with a different mobility example. Each traveler makes decisions that maximize his or her utility, resulting in a "user optimal equilibrium" in which no traveler can be better off by changing his/ her path. This solution, while optimal for each individual, is not optimal for society as a whole—it has associated greater total travel time than that of a "system optimal" solution that can be achieved with a central authority mandating the path followed by each traveler. The increase in total travel time of the user optimal solution over that of the system optimal solution is referred to as the "price of anarchy" (as introduced in the Joy of Living section of this report.)

OR models are ideally suited to characterize and quantify the "price of anarchy welfare gap" resulting when system users have the freedom to optimize their individual behavior. In systems for which this gap is significant, OR has also been used to create and analyze policies designed to increase social welfare by requiring or incenting changes in behavior. These policies might be in the form of administrative mandates or laws that force behavioral change, or they might be in the form of mechanisms designed to elicit behavioral changes as users optimize their own utilities. For example, returning to the case of mobility, pricing schemes that dynamically charge for use of a roadway segment based on levels of congestion work to incentivize some travelers to adjust their routes and/ or travel times to those corresponding to lower levels of congestion. As another example, in the case of constrained runway capacity, operations research-based models of market mechanisms, such as auctions and exchanges, have been

used to deduce which airlines derive the most value from the available capacity, allowing efficient capacity allocations to be determined (Ball et al 2006).

OR has long played a role in optimizing dynamic prices in systems with many customers or stakeholders, such as in revenue management or congestion pricing. In systems with fewer stakeholders, some receiving relatively large portions of the available resources (such as in the case of airlines utilizing capacity in an air transportation system), additional complexities arise. In such systems, the formulation and analysis of effective policies is complicated by gaming behavior by the stakeholders, whereby they might provide untruthful or strategic information to advance their own interests. The challenge—which is significant—is to recognize this behavior and to adjust policies accordingly. While significant challenges persist, operations researchers, such as Vaze and Barnhart (2012), have begun to develop and apply game-theoretic, operations research-based approaches to optimize equity-efficiency trade-offs in such constrained resource systems with fewer stakeholders.

Group decision-making provides an example of how OR's modeling capabilities can assist in analyzing and designing effective policies. A challenge in these environments is to reflect each individual's preferences in determining a single action for the group as a whole. The action might be based on an overall group preference ranking of a set of alternatives (e.g., proposals to accept for funding or students to select for admission using partial preferences (i.e., only ordering some of the choices) from the individuals within the group. As described in the Joy of Living section of this report, if the group collectively can agree that it is "fair" to follow a ranking that violates in aggregate the fewest individual preferences, then an OR optimization model can be applied to find that ranking. This approach can be applied to many more examples and other types of fairness in which groups must make collective decisions that reflect the interests of all individuals.

VII. Active Steps for Catalysis

Over the years, NSF has championed many new science and engineering research directions. Based on this study, we see an opportunity for NSF to guide several communities in directions which will ultimately catalyze the Engineering Grand Challenges. This section recommends two specific Actions, and followed by a discussion of how these actions play into the growth of “Grand Challenge Analytics.”

Recommended Actions

Action 1. NSF should announce an EFRI (Emerging Frontiers in Research and Innovation) topic for “Grand Challenge Analytics”. These proposals should not only be judged according to the impact on a Grand Challenge problem, but also on the novel methodology that will be developed as a result of the research. EFRI is a well-established program within NSF, and given the ground work of this report, we believe that NSF program officers will find it relatively straightforward to craft a RFP on this topic.

Action 2. Concurrently with Action 1, we recommend the formation of an Institute which will invite both EFRI-funded researchers as well as others from the field to participate in workshops which will explore common themes resulting from “Grand Challenge Analytics” projects. These workshops will not only help cross-fertilization between projects, but also help develop a general-purpose theory of analytics.

Discussion

What will NSF resources accomplish that is not being done today?

OR has a long history of bringing analytics to real-world problems. The recommended actions will accelerate dissemination between engineering and the computational, mathematical and economic sciences. With the role of OR as a nexus for these disciplines, the Operations Research program is ideally suited to guide this investment. The recommended actions will help build bridges between OR and a wider range of science and engineering communities, while continuing its strong presence in current high-impact areas such as energy, health, manufacturing, sustainability, and transportation.

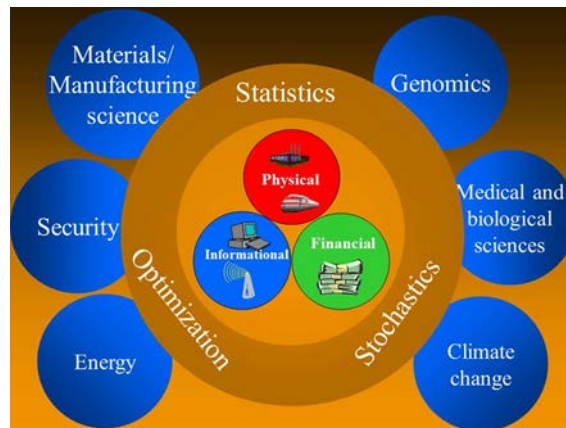
Why Now?

Trial-and-error experimentation is giving way to greater focus on modeling, optimization and simulation as a methodology for seeking new discoveries in energy, materials, nano-technology,

medicine, and of course, manufacturing systems and supply chains. In addition, the analytics boom propelled by OR is taking business enterprises to new levels of competitiveness. This combination of innovation and competitiveness bodes well for U.S. industry, and can only be sustained by greater cross fertilization between OR, and the more traditional areas of science and engineering.

How will this investment transform OR?

The growth of OR resulting from NSF actions may be visualized in the form of the graphic below. The innermost (light brown) circle in the following figure holds the early disciplines which motivated OR. These include computing/communications, finance, military, transportation, and others. As a result of its support of these “traditional applications,” the field of OR has created its own “ring” of concepts which integrate several fundamental pillars of OR knowledge¹¹. We expect that the next phase of OR growth will result from greater exchanges with domains associated with the Grand Challenges (which appear on the outer edges of the following figure). Over time, these challenges will add new dimensions to OR, which will then be represented by other new “rings” of OR concepts. If the recommendations of this report are adopted, we not only expect new science and engineering knowledge, but new transformative technologies, and new OR as well.



¹¹ For the sake of brevity, we have only listed a subset of OR applications and fundamentals.

References

- [1] Alivisatos, A.P., Chun, M., Church, G.M., Greenspan, R.J., Roukes, M.L., and Yuste, R. (2013). "The Brain Activity Map and Functional Connectomics." *NIH Project*.
- [2] American Health Information Management Association (2006). The Evolution of DRGs. Retrieved from http://library.ahima.org/xpedio/groups/public/documents/ahima/bok1_047260.hcsp?dDocName=bok1_047260
- [3] Ardabili, P.N. and Liu, M. (2012). "Establishing network reputation using mechanism design." *International Conference on Game Theory for Networks (GameNets)*. Vancouver, Canada.
- [4] Asgeirsson, E., Berry, J., Phillips, C.A., Phillips, D.J., Stein, C., and Wein, J. (2004). "Scheduling an industrial production facility." *Proceedings of IPCO X*, pp. 116–131.
- [5] Atkinson, M.P., Cao, Z., and Wein, L.M. (2008). "Optimal stopping analysis of a radiation detection system to protect cities from a nuclear terrorist attack." *Risk Analysis*, 28, 353-371.
- [6] Balinski M. and R. Laraki (2010). *Majority Judgement: Measuring Ranking and Electing*, MIT Press.
- [7] Ball M.O, Donohue G, Hoffman K (2006) "Auctions for the safe, efficient and equitable allocation of airspace system resources," In: Cramton P, Shoham Y, Steinberg R (eds) *Combinatorial Auctions*, MIT Press, Cambridge, 507–538.
- [8] Barnhart, C., Daskin, M.S., Dietrich, B., Kaplan, E., and Larson, R. (2008). "Grand challenges in engineering." *INFORMS (Institute for Operations research and the Management Sciences) White Paper*, submitted as response to NAE Grand Challenges.
- [9] Barnhart, C., Bertsimas, D., Caramanis, C., and Fearing, D.S. (2012) "Equitable and Efficient Coordination in Traffic Flow Management." *Transportation Science* (published online before print February 8, 2012, doi:10.1287/trsc.1110.0393).
- [10] Bixby, R.E. (2002). "Solving real-world linear programs: A decade and more of progress." *Operations Research*, 50, 3-15.
- [11] Bliznyuk, N., Ruppert, D., Shoemaker, C.A. (2011). "Bayesian inference using efficient interpolation of computationally expensive densities with variable parameter costs." *Journal of Computational and Graphical Statistics*, 20, 636-655.
- [12] Bockmayr, A. and Hooker, J. N. (2005). "Constraint Programming." In K. Aardal, G. Nemhauser, and R. Weismantel (Eds.), *Handbooks in Operations Research and Management Science 12: Discrete Optimization* (pp. 559 - 600). Elsevier.

- [13] Boros, E., Fedzhora, L., Kantor, P.B., Saeger, K.J., and Stroud, P. (2009). "Large scale LP model for finding optimal container inspection strategies." *Naval Research Logistics*, 56, 404-420.
- [14] Boyd, E.A. (2012). "Health care pricing in the U.S.: The case of hospitals." In O. Özer & R. Phillips (Eds.), *The Oxford Handbook of Pricing Management* (pp. 99-105). Oxford: Oxford University Press.
- [15] Bühlmann, P. and van de Geer, S. (2011), *Statistics for high-dimensional data*, Springer-Verlag, Berlin.
- [16] Brown, G.G. and Cox, L. (2011). "How probabilistic risk assessment can mislead terrorism risk analysts." *Risk Analysis*, **31**, 196-204.
- [17] Brown, G.G., Carlyle, W.M., Harney, R., Skroch, E., and Wood, R.K. (2009). "Interdicting a nuclear weapons project." *Operations Research*, 57, 866-877.
- [18] Brown, G.G., Carlyle, W.M., Salmerón, J., and Wood, R.K. (2005). "Analyzing the vulnerability of critical infrastructure to attack, and planning defenses." In H.J. Greenberg and J.C. Smith (Eds.), *Tutorials in Operations Research: Emerging Theory, Methods, and Applications*. Hanover, MD: INFORMS.
- [19] Brown, G.G., Carlyle, W.M., Salmerón, J. and Wood, R.K. (2006). "Defending critical infrastructure." *Interfaces*, **36**, 530-544.
- [20] Brown, T. (2007). "Multiple modeling approaches and insights for critical infrastructure protection." In D. Skanata and D.M. Byrd (Eds.), *Computational Models of Risks to Infrastructure*. IOS Press.
- [21] Bunn, M. (2010). *Securing the Bomb 2010*. Cambridge, Massachusetts and Washington, D.C.: Project on Managing the Atom, Belfer Center for Science and International Affairs, Harvard Kennedy School and Nuclear Threat Initiative.
- [22] Busemeyer, J.R. and Diederich, A. (2002). "Survey of decision field theory." *Mathematical Social Sciences*, 43, 345-370.
- [23] Candes, E.J. and Tao, T. (2005). "Decoding by linear programming." *IEEE Trans. Information Theory*, 51, 4203-4215.
- [24] Cate, F.H. (2012). "Roundtable on cyber threats, objectives and responses," Center for Applied Cybersecurity Research Report.
- [25] Chiang, M, Low, S.H., Calderbank, A.R. and Doyle, J.C. (2007), "Layering as optimization decomposition: a mathematical theory of network architectures," *Proceedings of the IEEE*, Jan. 2007, pp. 255-312.

- [26] Committee on Methodological Improvements to the Department of Homeland Security's Biological Agent Risk Analysis, National Research Council. *Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change*. Washington, DC: The National Academies Press, 2008.
- [27] Corbett, A.T. and Anderson, J.R. (1995). "Knowledge tracing: modeling the acquisition of procedural knowledge." *User Model. User-Adapt. Interact.*, 4, 253–278.
- [28] Defense Science Board Report (2011). "Enhancing Adaptability of U.S. Military Forces," p.31
- [29] Department of Energy. (2010). Workshop on Computational Materials Science and Chemistry for Innovation.
- [30] Desmarais, M.C. and Baker, R.S.J. (2012). "A review of recent advances in learner and skill modeling in intelligent learning environments." *User Modeling and User-Adapted Interaction*, 22, 9–38.
- [31] Dimitrov, N.B., Michalopoulos, D., Morton, D.P., Nehme, M.V., Pan, F., Popova, E., Schneider, E.A., and Thoreson, G.G. (2011). "Network deployment of radiation detectors with physics-based detection probability calculations." *Annals of Operations Research*, 187, 207-228.
- [32] Elliot, G. (2005). "US nuclear weapon safety and control." *MIT Program in Science, Technology, and Society*. Retrieved from <http://web.mit.edu/gelliott/Public/sts.072/paper.pdf>
- [33] Espinet, A. and Shoemaker, C.A. (2012). "Comparison of optimization algorithms for parameter estimation of multiphase flow models with application to geological carbon sequestration." *Advances in Water Resources*, 49.
- [34] Faruqui, A., Sergicil, S., (2009). "Household response to dynamic pricing of electricity – A survey of the experimental evidence." *The Brattle Group Report*.
- [35] Federal Cybersecurity Research and Development Program. (2012).
- [36] Felici, F. and Sauter, O. (2012). "Non-linear model-based optimization of actuator trajectories for tokamak plasma profile control." *Plasma Physics and Controlled Fusion*, 54.
- [37] FERC Staff Report (2011), "Recent ISO software enhancements and future software and modeling plans."
- [38] FERC Technical Conference (2013), "Increasing Real-time and Day-Ahead Market Efficiency Through Improved Software."
- [39] Fetter, R.B. (1991). "Diagnosis related groups: Understanding hospital performance," *Interfaces* 21(1), 6-26.

- [40] Gaukler, G.M., Li, C., Cannaday, R., Chirayath, S.S., and Ding, Y. (2011). "Detecting nuclear materials smuggling: using radiography to improve container inspection policies." *Annals of Operations Research*, 187, 65-87.
- [41] Golany, B., Kaplan, E.H., Marmur, A., and Rothblum, U.G. (2009). "Nature plays with dice – terrorists do not: allocating resources to counter strategic versus probabilistic risks." *European Journal of Operational Research*, 192, 198-208.
- [42] Gusfield, D. (1990) "A little knowledge goes a long way: Fast detection of compromised data in 2-d tables." *Proceedings of The IEEE Symposium on Research in Security and Privacy*, pp. 86-94.
- [43] Harney, R., Brown, G.G., Carlyle, W.M., Skroch, E., and Wood, R.K. (2006). "Anatomy of a project to produce a first nuclear weapon." *Science and Global Security*, 14, 163-182.
- [44] Hochbaum, D.S. (2010) "The separation, and separation-deviation methodology for group decision making and aggregate ranking." In J. J. Hasenbein (Ed.), *TutORials in Operations Research*, 7, Hanover, MD: INFORM. pp. 116—141.
- [45] Hochbaum, D.S., and Fishbain, B. (2011). "Nuclear threat detection with mobile distributed sensor networks." *Annals of Operations Research*, 187, 45-638.
- [46] Hochbaum, D.S. and Levin, A. (2006). "Methodologies for the group rankings decision." *Management Science*, 52(9) 1394--1408.
- [47] Holdren, J.P., Lander, E. and Varmus, H. (2010). "Report to the President and Congress – Designing a Digital Future: Federally Funded Research and Development in Networking and Information Technology." President's Council of Advisors on Science and Technology.
- [48] International Atomic Energy Agency. Fact sheet: IAEA Illicit Trafficking Database (ITDB). Retrieved from <http://www-ns.iaea.org/security/itdb.asp>
- [49] Kaplan, R.S. and Porter, M.E. (2011). "How to solve the cost crisis in health care." *Harvard Business Review*.
- [50] Kelly, F.P., Maulloo, A.K., Tan, D.K.H. (1998). "Rate control for communication networks: shadow prices, proportional fairness and stability." *Journal of the Operational Research Society*, 49, 237-252.
- [51] Lee, A. J. and McLay, L. A. and Jacobson, S. H. (2009). "Designing aviation security passenger screening systems using nonlinear control." *SIAM Journal on Control and Optimization*, 48(4), 2085-2105.
- [52] Madigan, D., Mittal, S., and Roberts, F., (2007). "Sequential decision making algorithms for port of entry inspection: Overcoming computational challenges." *In Proceedings of IEEE International Conference on Intelligence and Security Informatics*, 1-7.

- [53] Manne, A. (1974). "Waiting for the breeder." *Review of Economic Studies, Symposium* 47-65.
- [54] Manne, A. (1979). "ETA-MACRO: A model for energy economy interactions," in R. Pindyck (Ed.), *Advances in the economics of energy and resources*, Vol. 2, Greenwich, CT: JAI Press.
- [55] Manne, A. and Richels, RG. (2001). "An alternative approach to establishing trade-offs among greenhouse gases." *Nature*, 410, 675-677.
- [56] McLay, L. A. and Jacobson, S. H. and Kobza, J. E. (2006). "A multilevel passenger screening problem for aviation security." *Naval Research Logistics*, 53, 183-197.
- [57] McLay, L.A, Lloyd, J.D, and Niman, E. (2011). "Interdicting nuclear material on cargo containers using knapsack problem models." *Annals of Operations Research*, 187, 185-205.
- [58] Men, H., Freund, R.M., Nguyen, N.C., Saa-Seoane, J., Peraire, J. (2013). "Fabrication-adaptive optimization, with an application to photonic crystal design." *Operations Research*, to appear.
- [59] Men, H., Nguyen, N.C., Freund, R.M., Lim, K.M., Parrilo, P.A., Peraire, J. (2011). "Design of photonic crystals with multiple and combined band gaps." *Physical Review E* 83(4) 046703.
- [60] Men, H., Nguyen, N.C., Freund, R.M., Parrilo, P.A., Peraire, J. (2010). "Bandgap optimization of two-dimensional photonic crystals using semidefinite programming and subspace methods." *Journal of Computational Physics*, 229, 3706-3725.
- [61] Metz, B., Davidson, O., de Connick, H., Loos, M., Meyer, L. (2005) (Eds.) "Carbon Dioxide Capture and Storage." *Intergovernmental Panel on Climate Change Report*. Retrieved from http://www.ipcc.ch/pdf/special-reports/srccs/srccs_wholereport.pdf
- [62] Mitrovic, A. (2012). "Fifteen years of constraint-based tutors: what we have achieved and where we are going." *User Modeling and User-Adapted Interaction*, 22, 39-72.
- [63] Morton, D.P., Pan, F., and Saeger, K.J. (2007). "Models for nuclear smuggling interdiction." *IIE Transactions on Operations Engineering*, 38, 3-14.
- [64] Moving Target Research Symposium (2012). Retrieved from <http://cps-vo.org/group/mtrs/program>
- [65] Murray, R., Haxton, T., Janke, R., Hart, W.E., Berry, J., and Phillips, C. (2010). "Sensor network design for drinking water contamination warning systems: A compendium of research results and case studies using the TEVA-SPOT software." (Technical Report EPA/600/R-09/141) *National Homeland Security Research Center*, Office of Research and Development, U.S. Environmental Protection Agency.
- [66] National Nuclear Security Administration. Fact sheet: NNSA's Second Line of Defense Program. Retrieved from:

<https://nnsa.energy.gov/mediaroom/factsheets/nnsassecondlineofdefenseprogram>

- [67] National Science and Technology Council. (2012). The Networking and Information Technology Research and Development (NITRD) Program 2012 Strategic Plan.
- [68] Nemhauser, G.L. (2013). "Integer Programming: Global Impact," Plenary Speech, EURO Conference, Rome, Italy.
- [69] Pita, J., Jain, M., Western, C., Paruchuri, P., Marecki, J., Tambe, M., Ordonez, F., and Kraus, S. (2008). "Security via randomization: A game-theoretic model and its application to the Los Angeles International Airport." *Proceedings of IEEE Conference on Technologies for Homeland Security*.
- [70] Polunchenko, A. and Tartakovsky, A.G. (2012). "State-of-the-Art in Sequential Change-Point Detection," *Methodology and Computing in Applied Probability*, vol. 14, no. 3, pp. 649-684.
- [71] Polunchenko, A. Tartakovsky, A.G. and Mukhopadhyay, N. (2012) "Nearly Optimal Change-Point Detection with an Application to Cybersecurity," *Sequential Analysis*, vol. 31, no. 3, pp. 409-435.
- [72] Powell, W. and Frazier, P. (2008). "Optimal Learning." TutORials in Operations Research: State-of-the-Art Decision-Making Tools in the Information-Intensive Age (J.C. Smith, ed.) pp. 213-246.
- [73] Powell, W. and Ryzhov, I. (2012). *Optimal Learning*. Wiley.
- [74] Raheja, D. and Gullo, L.J. (2012). *Design for Reliability* Hoboken, NJ: John Wiley and Sons.
- [75] Regis, R.G., Shoemaker, C.A. (2007). "A stochastic radial basis function method for the global optimization of expensive functions." *INFORMS Journal on Computing*, 19, 497-509.
- [76] Salter, M. B (2008). "Imagining numbers: risk, quantification, and aviation security." *Security Dialogue*, 39(2-3), 243-266.
- [77] Shishehchi, S., Banihashem, S.Y., Zin, N.A.M., Noah, S.A.M. (2011). "Review of personalized recommendation techniques for learners in e-learning systems." *International Conference on Semantic Technology and Information Retrieval*. Putrajaya, Malaysia.
- [78] Singpurwalla, N.D. and Wilson, S.P. (1994). "Software Reliability Modeling," *International Statistical Review*, 289-317.
- [79] Thompson L.G., Mosley-Thompson E., Davis M.E., Henderson K.A., Brecher H.H., Zagorodnov ,V.S., Mashiotta T.A., Lin P.N., Mikhalenko V.N., Hardy D.R., Beer, J. (2002). "Kilimanjaro ice core records: Evidence of Holocene climate change in tropical Africa." *Science*, 298(5593), 589-593.

- [80] Thompson L.G., Mosley-Thompson E., Davis M.E., Lin P.N., Henderson K.A., Cole-Dai J., Bolzan J.F., Liu K.B. (1995). "Late glacial stage and holocene tropical ice core records from Huascarán, Peru" *Science*, 269(5220) 46-50.
- [81] Thompson, L.G. (2000). "Ice core evidence for climate change in the Tropics: implications for our future." *Quaternary Science Reviews*, 19(1–5), 19–35.
- [82] U.S. Centers for Medicare and Medicaid Services. (2010). National Health Expenditures 2010 Highlights. Retrieved from www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/downloads/highlights.pdf
- [83] U.S. White House. (2010). *National Security Strategy*. Washington, D.C.: Government Printing Office.
- [84] United States General Accounting Office (GAO). (2011) "Combating nuclear smuggling: DHS has developed a strategic plan for its global nuclear detection architecture, but gaps remain." (Technical Report GAO-11-869T, 2011) *Testimony Before the Subcommittee on Cybersecurity, Infrastructure Protection, and Security Technologies*, U.S. House of Representatives.
- [85] Vaze, V. and Barnhart, C., (2012). "Modeling Airline Frequency Competition for Airport Congestion Mitigation." *Transportation Science*, to appear (published online before print May 3, 2012, doi: 10.1287/trsc.1120.0412)
- [86] Velliste, M. et. al. (2008). "Cortical control of a prosthetic arm for self-feeding." *Nature*, 453, 1098-1101.
- [87] Wang, P. and Menon, R. (2013). "Optimization of periodic nanostructures for enhanced light-trapping in ultra thin photovoltaics." *Optics Express*, 21(5), 6274-6285.
- [88] Wein, L.M., Liu, Y., Cao, Z., and Flynn, S.E. (2007). "The optimal spatiotemporal deployment of radiation portal monitors can improve nuclear detection at overseas ports." *Science and Global Security*, 15, 211-233.
- [89] Wein, L.M., Wilkins, A.H., Baveja, M., and Flynn, S.E. (2006). "Preventing the importation of illicit nuclear materials in shipping containers." *Risk Analysis*, 26, 1377-1393.
- [90] Wein, L.M. and Liu, Y. (2005). "Analyzing a bioterror attack on the food supply: the case of botulinum toxin in milk." *Proceedings of the National Academy of Science*, 102, 9984-9989.
- [91] Xiao, L. (2010), "Dual averaging methods for regularized stochastic learning and online optimization," *Journal of Machine Learning Research*, pp. 2543 – 2596.
- [92] Zabaco, M. and Gorelick, S. (2012). *Proceedings of the National Academy of Sciences*. (Boyd 2014)

OR as a Catalyst