Deterministic Formulations of Probabilistic Set Covering Problems

Shabbir Ahmed
Georgia Tech

We consider set covering problems with unreliable coverage. Such problems arise in various stochastic location models, such as sensor placement. These problems can be formulated as probabilistically constrained optimization models involving Bernoulli random variables. By exploiting the probabilistic structure of the underlying uncertainty, we develop novel deterministic mixed integer linear programming formulations for these problems.

Quadratic and TGP models within the MADS Algorithm for Blackbox Optimization

Charles Audet
École Polytechnique de Montréal

The surrogate management framework was introduced back in 1999 for expensive blackbox optimization. A surrogate model is not necessarily a good approximation of a blackbox function, but needs to share some similarities, and its main feature is to be cheaper to evaluate than the true function. The present talk discusses strategies to use and build surrogate models within the mesh adaptive direct search (MADS) algorithm. A first method is inspired by work in Derivative-Free Optimization, and consists of building quadratic models, and a second strategy exploits the recent Treed Gaussian Process statistical tool, resulting in global models. These dynamic surrogate methods are applied to build models of both objective and each constraint functions defining the blackbox optimization problems. Results of numerical experiments with the Nomad software on constrained problems are presented to quantify the added value of this approach.

Evaluating Mission Assurance for Andersen AFB, Guam

Gerald Brown
Naval Postgraduate School

We show how we are applying optimization and a Defender-Attacker-Operator model of infrastructure on Guam, American Marianas Islands, to evaluate mission resilience of Anderson AFB. Although we model electricity, water and telecommunications for the entire island, it turns out that transport and storage of JP8 jet fuel is the critical infrastructure.

Finding GMA Orthogonal Arrays by Enumeration

Dursun Bulutoglu
Air Force Institute of Technology

Factorial designs are used extensively in a wide range of scientific and industrial investigations for extracting as much information as possible at a fixed cost. Orthogonal arrays (OAs) that sequentially minimize the generalized wordlength pattern (GWP) are called generalized minimum aberration (GMA) OAs. GMA OAs are the most efficient factorial designs for certain statistical models. The GMA property of OAsis preserved under design isomorphism, hence classifying OAs up to isomorphism allows a GMA OA to be found. The most efficient algorithms in the literature extend all non-isomorphic OAs by columns via integer linear programming with isomorphism rejection or constraint programming with isomorphism rejection. However, both algorithms suffer from the exponential growth in the number of OAs that need to be enumerated with increasing number of runs (rows). New transformations based on the equivalence of orthogonal designs (ODs) are proposed to overcome this growth. When applied to enumerating all non-isomorphic OAs in our test bed of problems, these new transformations decreased the number of intermediate designs that needed to be enumerated by a factor of 10 to 15. This, in turn, reduced the overall solution times by the same factor.

Clique Relaxation Models in Network Analysis

Sergiy Butenko
Texas A&M University

Increasing interest in studying *community structures*, or *clusters* in complex networks arising in various applications, in particular, in the analysis of social and biological networks, has led to a large and diverse body of literature introducing numerous graph-theoretic models relaxing certain characteristics of the classical clique concept. This talk analyzes the elementary clique-defining properties inherently exploited in the available *clique relaxation* models and proposes a taxonomic framework that not only allows to classify the existing models in a systematic fashion, but also yields new clique relaxations of potential practical interest. The relationships between the cohesiveness properties of different clique relaxation structures are established, and practical implications of choosing one model over another are discussed. In addition, effective scale-reduction techniques based on new clique relaxation structures are proposed that allow to solve very large-scale, sparse instances of the maximum clique problem to optimality.

A Dynamic Optimization Approach for Persistent Search Problems

Christos G. Cassandras
Boston University

In persistent search (or "persistent monitoring") problems the objective is to control the movement of multiple cooperating agents to minimize an uncertainty metric in a given mission space over some given mission time. In a one-dimensional mission space, we show that the optimal solution is a set of "oscillators" where each agent moves at maximal speed from one switching point to the next, possibly waiting some time at each point before reversing its direction. Thus, the solution is reduced to a simpler parametric optimization problem: determining a sequence of switching locations and associated waiting times at these switching points for each agent. This amounts to a hybrid dynamic system which we analyze using Infinitesimal Perturbation Analysis (IPA) to obtain on-line gradients of the uncertainty metric with respect to the decision variables, i.e., the switching point locations and associated waiting times. Using this information, we obtain a complete on-line solution through a gradient-based algorithm. An unusual feature of this algorithm is that the dimensionality of the parameter vector is a priori unknown and must be determined along with the actual values of the decision variables. Exploiting previously established robustness properties of IPA gradients, we also show that the solution we obtain is independent of the uncertainty model used. This implies that the same approach can be used in a setting where the uncertainty model is stochastic with unknown characterization. This work establishes the basis for extending the approach to a two-dimensional mission space by decomposing it into regions and using the one-dimensional optimal solution within each such region. The persistent search problem is an instance of a dynamic optimization problem which can be reduced to a parametric optimization problem where one can exploit methodologies such as IPA and obtain complete on-line solutions.

Coherent Risk-Adjusted Decisions over Time: a Bilevel Programming Approach

Jonathan Eckstein, Andrzej Ruszczyński
Rutgers

This project concerns a bilevel optimization approach to resolving the time-inconsistency and incomposability difficulties inherent in applying coherent risk measures to stochastic optimization problems in which decisions and information evolve together over time.

First, we consider the setting of extended two-stage stochastic programming, and describe the drawbacks of prior approaches, showing why a multi-level optimization technique might be attractive.

Second, we present a negative (but expected) recent result, showing that risk-averse bilevel stochastic programming models of the kind we propose are NP-hard for the common mean-semideviation and average/conditional-value-at-risk measures.

Third, we discuss some recent results on constructing upper and lower approximations of time-inconsistent problems by tractable time-consistent models, and developing a dynamic-programming-like solution procedure.

Finally, we investigate related theoretical and practical issues in dynamic programming approaches when compositions of measures of risk are employed. We compare deterministic and randomized policies, and we illustrate our results on an optimal stopping problem and a stochastic optimal path problem.

Solving IPs and Improving Bounds by Restricting and Relaxing Formulations

Ergun Ozlem, George Nemhauser
Georgia Tech

We present initial results on incorporating iterative strategies in developing better and faster optimization algorithms. We will first discuss iteratively solving efficient but invalid formulations of a multi-period network expansion problem. In our preliminary work in this area, we solve different versions of a formulation by tightening or relaxing open/close constraints by modifying big M values. In general, it is easy to calculate very loose M values that ensure valid formulations or to calculate very tight and invalid M values that only ensure the existence of some feasible solution (for example by using heuristics that push demand towards source nodes in the context of network connectivity problems). However, in general it is NP-hard to calculate the tightest valid M value.

We will present the results of an initial computational study on algorithms we developed that, starting from very tight or loose Ms, iteratively search the space of the M values In the second part of our presentation, we will discuss searching the dual space of an integer program for better bounds by iteratively relaxing constraints. One of the challenges in this framework is to identify the important constraints. We will present initial computational results of our algorithms on the multi-knapsack problem where we employ various simple search strategies guided by the solutions to linear programing relaxations and exact solutions to previous formulations.

New One-Step Bayes-Optimal Algorithms for Global Optimization: Parallel Computing and Common Random Numbers

Peter Frazier
Cornell University

We consider Bayesian global optimization for use with expensive or time-consuming function evaluations, and develop new algorithms for two specific problem contexts.

In the first problem context, multiple function evaluations can be performed in parallel, and function evaluations are noise-free. Such settings are increasingly common with the current ubiquity of parallel computing. In this setting, with a Gaussian process prior on the objective function, a one-step Bayes-optimal expected-improvement-type algorithm was proposed as a conceptual algorithm by Ginsbourger, Le Riche, and Carraro (2008), but this algorithm was deemed too difficult to actually implement in practice. Using stochastic approximation, we show how this conceptual algorithm can be implemented efficiently, and demonstrate that the resulting practical algorithm provides a speedup over the single-threaded expected improvement algorithm.

In the second problem context, function evaluations are obtained from a stochastic simulation, and we may induce correlations in the output across input points by using a stream of common random numbers. Common random numbers are used broadly within simulation to improve algorithm efficiency, but have been underutilized within Bayesian simulation optimization. We show how the value of information from observing the difference between a pair of correlated samples may be computed analytically, and use this analytic computation to develop a new knowledge-gradient algorithm that takes advantage of common random numbers. When sampling correlations are high, we show that this algorithm is more efficient than algorithms that do not take advantage of common random numbers.

Recent Research on Design Optimization of Wave Propagation in Metamaterials: Fabrication-Robust Design, Binary Optimization with Reduced Basis, Waveguide Optimization of Crystal Fibers, and 3-Dimensional Photonic Crystal Design

Robert Freund, Jaime Peraire
MIT

We report on some of our recent research on design optimization of wave propagation in metamaterials. In fabrication-robust design, we have developed a new methodology for incorporating considerations of implementation/fabrication into optimization models quite generally. It goes without saying that the computed solutions of many optimization problems cannot be implemented directly due to either (i) the deliberate simplification of the model by ignoring integer/binary conditions that would otherwise render the problem intractable, or (ii) human or technological reasons (we cannot be sure our plant or software can implement to the accuracy prescribed in the computed solution). We propose a new alternative paradigm for treating issues of implementation that we call "implementation robustness." This paradigm is applied to the setting of optimizing the fabricable design of photonic crystals with large band-gaps.

Most of the metamaterial designs described in the literature assume infinite periodicity. There is therefore a need for a design methodology capable of computing designs involving two different materials where the underlying design variables correspond to a finite set of pixels in a 2-dimensional mesh, and where the goal is a design with prescribed metamaterial properties. This naturally leads to the consideration of binary optimization models in contrast to classical (continuous) models which generically provide solutions that then need to be "rounded" to binary values. While the potential drawback of binary optimization is that its computational complexity is usually NP-hard and hence theoretically unattractive, we show that binary optimization combined with a reduced basis approach can relatively efficiently produce very good solutions to metamaterial design problems of interest.

Photonic crystal fiber (PCF) is a new class of optical fiber based on the microstructures and the properties of photonic crystals. Compared to conventional optical fibers, the higher degrees of freedom and versatile structures of PCF enable its guidance of light with lower losses, higher nonlinearities, higher birefriengence, more thermal resistance, and more engineerable dispersions. Because of all these unique characteristics, PCFs have found increasingly novel applications in telecommunications, sensors, high-power lasers, optical imaging probes, and many other areas. We have successfully applied our methodology tools from bandgap optimization (semidefinite optimization, locally fixing orthogonal spaces, "active subspaces") to the optimization of the bandwidth of 3-D single-polarization single-mode photonic crystal fibers (SPSM-PCF).

Large-scale optimal design of 3-D photonic crystals poses many more challenges. Although our algorithm was shown to be able to successfully increase the gap-midgap ratio of a given initial photonic crystal configuration, no significant complete gap (positive gap-midgap ratio) have been obtained in the simple cubic (sc) lattice. For one thing, the existence of a complete 3-D band gap could be very rare. Nevertheless, we plan to build lattice structures with more high-symmetry points in the irreducible Brillouin zone, e.g., the face centered cubic (fcc) lattice, body centered cubic (bcc) lattice, and the diamond lattice.

Speeding Up the Cross Entropy Method for Global Optimization

Yuting Wang, Alfredo Garcia
University of Virginia

We analyze a multi-start implementation

of the CE method for global optimization in which sampled solutions are used as initial solutions or seeds for independent local searches. We provide a formal characterization for the speed of convergence (both worst-case and average) by developing a Markov chain model in which the state space is the set of all locally optimal solutions. The speed of convergence (worst-case) is determined by second largest eigenvalue associated with the transition probability matrix. This eigenvalue has a straightforward interpretation in terms of the "worst" possible state in which the process remains for a relatively large number of iterations. The average performance is characterized in terms of a classification of states into "clusters". The average performance of the single- thread CE methods deteriorates in problems with many clusters with relatively large basins of attraction. These results motivate a new parallel implementation of the method that is guaranteed to speed up convergence by means of an acceptance-rejection test aimed to prevent duplication in search effort.

Alternative Methods for Obtaining Optimization Bounds

John Hooker
Carnegie Mellon University

In this talk I report on two general methods for obtaining bounds for discrete optimization problems. They provide alternatives to bounds obtained from the continuous relaxation of 0-1 models. One method reformulates the problem using finite domain variables, as is frequently done in constraint programming, and analyzes the polyhedral structure of the resulting model. Facet-defining cuts are mapped back into the 0-1 space. We apply this to graph coloring and find that the cuts deliver tighter bounds than known 0-1 cuts in a fraction of the time. The second method obtains bounds from a discrete relaxation based on binary decision diagrams. We apply this to the independent set problem and obtain significantly tighter bounds than CPLEX obtains with cutting planes at the root node, and in a small fraction of the time. This is joint work with David Bergman, Andre Cire, and Willem van Hoeve.

Sequential Stochastic Assignment Problems: Dependency and Uncertainty

Sheldon H. Jacobson
University of Illinois at Urbana-Champaign

Consider the stochastic sequential assignment problem (SSAP), where N workers are allocated to N IID sequentially-arriving tasks so as to maximize the expected total reward. This presentation discusses a variation of SSAP, where the task values in each time period depend on the value of the preceding task. Note that this problem contains as a special case an SSAP in which the dependency between task values is governed by a Markov chain. Also, the tasks arrive with a certain probability in each time period (independent of the other time periods), providing an added dimension of uncertainty. The objective function is to maximize the expected total reward, and an optimal assignment policy is established, with the threshold interval breakpoints random variables. Furthermore, a generalization of this problem is discussed where the total number of arriving tasks is unknown until after the final arrival and is a random variable following an arbitrary distribution with finite support. This problem is further extended to the case where the distribution of the number of tasks has infinite support, and an optimal policy is proposed so as to achieve the maximum expected total reward over an infinite-horizon. The present work deviates from the existing SSAP literature by combining dependency between task values with randomness in the number of tasks, hence opening up a new direction of exploring uncertainty for such problems.

Data-Driven Online And Real-Time Combinatorial Optimization

Patrick Jaillet
MIT

Concentrating on autonomous spatial exploration and information harvesting problems, the research under this grant look at the formulation and algorithmic design/analysis of online optimization problems with (i) incomplete and uncertain input streams, (ii) time-sensitive objectives, and (iii) short time requirements and/or capacity constraints for some decisions. After briefly motivating the problems and the methodologies, we will present recent results on online quota traveling salesman problems and online resource allocation problems, including worst-case competitive ratio bounds for deterministic and/or randomized online algorithms for these problems, as well as some results under stochastic assumptions for the input of the instances.

Probabilistic Analysis of Multidimensional Assignment Problems

Pavlo Krokhmal
University of Iowa

We consider a class of discrete optimization problems where the underlying combinatorial structure is based on hypergraph matchings, which generalize the well-known problems on bipartite graph matchings, such as the Linear Assignment Problems, Quadratic Assignment Problems, etc., and are also known as multidimensional assignment problems (MAPs). Properties of large-scale randomized instances of MAPs are investigated under assumption that their assignment costs are iid random variables. In particular, we consider linear and quadratic problems with sum and bottleneck objectives. For a broad class of probability distributions, we demonstrate strong convergence properties of optimal values of random MAPs as problem size increases. The analysis allows for identifying a subset of the feasible region containing high-quality solutions of randomized MAPs. We also investigate the average-case behavior of Linear Sum MAP in the case when the assumption regarding independence of the assignment costs is relaxed, and a correlation structure is present in the array of assignment costs. In particular, we consider the case of decomposable assignment costs, when the cost of a hyperedge in d-partite d-uniform hypergraph is given by the cost of a clique in the corresponding d-partite graph. Bounds on the expected optimal cost of LSMAP are constructed, and average-case performance of several heuristics is discussed.

Instant Approximate 1-center on Roads

Piyush Kumar
Florida State University

In this talk, I will present an algorithm to compute 1-center solutions on road networks, an important problem in GIS. Using Euclidean embeddings, and reduction to fast nearest neighbor search, we devise an approximation algorithm for this problem. Our initial experiments on real world data sets indicate fast computation of constant factor approximate solutions for query sets much larger than previously computable using exact techniques. I will also show extensions to k-clustering problems.

An Optimization Framework for Air Force Logistics Models

Retsef Levi, Tom Magnanti
MIT

We consider a class of models that capture central issue in maintenance of fighter aircrafts with low-observable (LO) capabilities (i.e., F-22). These aircrafts pose unique challenges from a maintenance scheduling standpoint. Over time, an aircrafts LO capability deteriorates in a random manner as a result of everyday flying operations. In order to ensure that each aircraft retains its LO capability and remains fully mission capable (FMC), maintenance squadrons must constantly work to repair the exterior skin of the these aircraft. When considering an entire fleet of aircraft, maintenance squadrons must make difficult decisions regarding when to maintain each aircraft and to what extent, given that there is a limited amount of maintenance resources. In addition, one has to decide which aircrafts to use to comply with the daily sorties requirements. Currently, the Air Force does not have a published policy as to how LO maintenance should be planned or scheduled.

We show how to model the problem as a variant of the restless multi-arm bandit. This leads to simple index policies, which can be shown (through computational experiments) to significantly increase the overall FMC rate of the respective fleet. The computational experiments are based on real maintenance data of the AF. The index policies are computed off a linear programming relaxation of the problem. We then discuss several interesting related theoretical questions.

Joint work with Phil Chu, Vivek Farias, John Kessler, Yaron Shaposhnik and Eric Zarybnisky.

Some Recent Research in Simulation Optimization

Steve Marcus, Michael Fu, Jiaqiao Hu
University of Maryland

We report on some of our current research in simulation optimization, considering both the discrete and continuous cases. Specifically, we present preliminary results for newly proposed approaches using stochastic gradient estimation techniques for augmenting traditional regression and stochastic kriging for fitting response surfaces and metamodels estimated from simulations. For statistical ranking and selection among a fixed set of alternatives, we present a new algorithm for optimally allocating a computational/simulation budget when there is unknown correlation in the performance estimation of the alternatives.

A first step in designing a VU-algorithm for nonconvex minimzation

Robert Mifflin
Washington State University

This talk lays the ground work for designing a future VU-type minimization algorithm to run on locally Lipschitz functions for which only one Clarke generalized gradient is known at a point. This entails development of a bundle method sub-algorithm that has provable convergence to stationary points for semismooth functions and can make adequate estimates of "V-subspace" bases in the presence of nonconvexity. Ordinary bundle methods generate consecutive "null steps" from a "bundle center" until a "serious step" point is found, which then becomes the next center. A VU-algorithm is similar except that its serious descent point is "very serious" (called serious in [1]) which means it defines a good "V-step" and also generates a good "U step" to add to its very serious point in order to define the next center.

For an objective function of one variable the desired VU-algoritm exists, but it does not extend directly to functions of n variables. For convex functions about 20 years worth of proximal point and VU theory had to be developed before a rapidly convergent method for the multivariable case could be defined. For the nonconvex case there are some ideas associated with the single variable algorithm which can be adapted for developing an n variable method. One is to use second derivative estimates from differencing generalized gradients to give better "V-model" approximation functions when negative curvature is detected. Another is to employ a certain form of a safeguard to guarantee desired convergence even when the negative curvature estimates are not good enough for proving stationarity in the limit.

The new bundle algorithm described here aims at keeping as many as possible properties of an ordinary bundle method for convex minimization. This leads to basing convergence proofs on first showing that a sequence of V-model proximal point subproblem objective values converges to zero. Associated with this is the desire to preserve the concepts of null and serious steps in the sense that, under reasonable assumptions, there is convergence of a bundle algorithm generated sequence to a stationary point for the problem objective function if there are either (1) an infinite number of serious step iterations or (2) a finite number of serious steps followed by an infinite number of consecutive null step iterations.

1. R. Mifflin & C. Sagastizabal, "A VU-algorithm for convex minimization", *Math. Prog.* 104, 2005.

Particle Filters and Assignment Problems for Target Tracking

Eduardo Pasiliao
Air Force Research Laboratory

This talk will present several research topics in target tracking in which the PI has been involved. We present an application of particle filters to target state estimation in an urban environment. The objective is to improve the estimation efficiency by developing an accurate and comprehensive representation of the target task-space and by incorporating measurements from the readily available soft information sensors. In addition, the sensor observation uncertainties are fused into the particle weight update law to inject robustness with respect to the classification type I and type II errors.

Results and ongoing efforts in graph-theoretic analysis of local search techniques, linear assignment interdiction problems via bilevel programming, computational studies of randomized multidimensional assignment problems, and implementation of assignment problem solutions on high performance computers will also be presented.

Competitive Resource Allocation, QPCCs: A progress report

Jong-Shi Pang
University of Illinois at Urbana-Champaign

John Mitchell
Rensselaer Polytechnic Institute

This talk is an annual progress report of our research on competitive resource allocation problems and the class of quadratic programs with complementarity constraints (QPCCs). On the former, we have investigated a class of VI-constrained hemivariational inequalities and Nash equilibrium problems with minmax objectives, both having applications to power control in ad-hoc networks. In addition, we have analyzed the maximum sum-rate attainable for Nash solutions in distributed Gaussian interference channels with unbounded budgets. We have obtained some new results for the class of QPCCs, linking this class of optimization problems with the class of quadratically constrained quadratic programs and the class of copositive programs, and have extended a logical Benders decomposition method for the global resolution of a QPCC. In addition, we have been inspired for a competitive biofuel supply chain application that leads to a discretely-constrained QPCC whose global solution is a challenging task under investigation. Other applications of these two broad mathematical paradigms are also part of the research program.

We have carried out the above research jointly with two colleagues Francisco Facchinei and Gesualdo Scutari and my student Alberth Alvarado on the first topic, and a co-PI John Mitchell, his student Lijie Bai, a colleague Yanfeng Ouyang and his student Bella Bai, and my student Yuching Lee on the second topic.

Optimal Learning and Approximate Dynamic Programming for High-Dimensional Stochastic Search and Control

Warren B. Powell
Princeton University

I will describe ongoing research into the use of advanced machine learning algorithms based on Dirichlet process models in the context of stochastic search and approximate dynamic programming, research experiences with the knowledge gradient using Gaussian process priors, and the incorporation of optimal learning methods in a broad class of problems, including those with a physical state as well as a new project involving sequential experimental design for research at the interface of nano-bio technologies. We are also starting to work on problems with heavy-tailed stochastic processes, where I will describe a new method based on quantile optimization using asymmetric signum functions. New work is focusing on sequential control problems where the objective function involves a broad class of risk measures to account for extreme events.

Irregular Polyomino Tiling via Integer Programming

Oleg Prokopyev
University of Pittsburgh

A polyomino is a generalization of the domino and is created by connecting a certain number of unit squares along an edge. The problem of tiling a region with a given set of polyominoes is known to be computationally hard. In this talk we focus on tiling a rectangle with a given finite set of polyominoes from the same family. One recent application of this problem is in the design of phased array antennas, where irregular polyomino tilings are used for reducing quantization sidelobes, i.e., beams with directions other than the main beam direction of the antenna. We formulate the irregular polyomino tiling problem as a nonlinear exact set covering model, where the irregularity requirement is incorporated into the objective function via the information theoretic entropy concept. We describe both exact and heuristic approaches developed for solving this problem. Preliminary computational results are encouraging.

Expanding the Reach of Nonlinear Optimization

Michael C. Ferris, Stephen M. Robinson
University of Wisconsin

In this presentation we will review completed and ongoing work under the above grant. Areas of current work include the following, of which we will present a subset consistent with time constraints in the review.

- Using integer linear programming to infer the pathways through which host factors modulate replication of a virus in a host cell. Input is a set of viral phenotypes observed in single-host-gene mutants and a background network consisting of a variety of host intracellular interactions. The output is an ensemble of subnetworks that provides a consistent explanation for the measured phenotypes, predicts which unassayed host factors modulate the virus, and predicts which host factors are the most direct interfaces with the virus.

- A bi-level optimization model to produce policies that meet new requirements for demand response compensation in organized wholesale energy markets by incorporating the demand response decision in a traditional economic dispatch algorithm,. The model proves effective running on industrial scale data and yields insight into dispatch operations under the new compensation rule.

- A new modeling language extension to facilitate the solution of hierarchical models with information-sharing interfaces among model components.

- Shaping dose volume histograms in radiotherapy treatment using stochastic programming techniques (joint work with Andrzej Ruszczyński and Darinka Dentcheva).

- Solving multi-period stochastic equilibrium problems with contracts allowed (joint work with Roger J.-B. Wets).

- Mathematical models for structure of equilibrium problems that enhance both problem analysis and effective numerical solution.

Optimal Budget Allocation for Sample Average Approximation

Johannes Royset
Naval Postgraduate School

Terry Rockafellar
University of Washington

The sample average approximation approach to solving stochastic programs induces a sampling error, caused by replacing an expectation by a sample average, as well as an optimization error due to approximating the solution of the resulting sample average problem. We obtain estimators of an optimal solution and the optimal value of the original stochastic program after executing a finite number of iterations of an optimization algorithm applied to the sample average problem. We examine the convergence rate of the estimators as the computing budget tends to infinity, and characterize the allocation policies that maximize the convergence rate in the case of sublinear, linear, and superlinear convergence regimes for the optimization algorithm.

Recovering Sparse Hessian Structure and Using Random Models in Derivative-Free Optimization

Katya Scheinberg
Lehigh University

We will show how using results from compressed sensing we can construct accurate second order models of black box functions by using relatively few samples. We will apply these ideas within a convergent trust-region based derivative free optimization method.

Dynamic Stochastic Mixed-Integer Programming Algorithms

Suvrajeet Sen
Ohio State University

Many hybrid systems involve discrete and continuous decisions simultaneously within a continually evolving dynamic system. When coupled with the fact that such environments can be stochastic, such systems give rise to dynamic stochastic mixed-integer programming problems. We are studying new decomposition algorithms in which stagewise models are approximated using valid inequalities that are parameterized by the state of the system. This allows us to build approximations in a sequential manner, and prove asymptotic results when the integer variables are restricted to be binary.

Bounded Error Approximation Algorithms for Risk-Based Intrusion Response

K. Subramani
West Virginia University

Risk analysis has been utilized to manage the security of systems for several decades. However, its use has been limited to offline risk computation and manual response. In contrast, we use the risk computation to drive changes in an operating system's security configuration. This allows risk management to occur in real time and reduces the window of exposure to attack. We posit that it is possible to protect a system by reducing its functionality temporarily when under siege. Our goal is to minimize the tension between security and usability by trading them dynamically. Instead of statically configuring a system, we aim to monitor the risk level and use it to drive the trade-off between security and utility. The advantage of this approach is that it provides users with the maximum possible functionality for any predefined level of risk tolerance.

In this talk, we frame Risk management as an exercise in managing the constraints on edge and vertex weights of a tripartite graph, with the partitions corresponding to the threats, vulnerabilities, and assets in the system. If a threat requires a specific permission and affects a particular asset, an edge is added between the threat and the permission that mediates access to the vulnerable resource, and another between the permission and the asset. The presence of a path from a threat, through a permission check, to an asset contributes an element of risk. Since the set of threats cannot be altered by the response apparatus, we can merge the first partition, which contains the threats, into the second by scaling each permission's weight (which represents its probability of being granted) with the sum of the threat likelihoods that have incident edges on the permission. We note that the semantics of risk management require that the risk must be reduced below the threshold of tolerance each time it is found to exceed it.

Risk can be reduced by denying access to a resource that contains a vulnerability or activating data protection measures. This is modeled as the removal of edges representing risk in the aforementioned graph. Depending on whether the risk estimates are integers or reals, whether the vulnerabilities and consequences are independent or conditional, whether the application workload is known in advance or not, whether the workload is stable or changing rapidly, and the semantics of response selection, there are different underlying graph problems. We seek to analyze the complexity of and devise approximation algorithms for these and related problems that form the algorithmic underpinnings of optimal risk management.

Mixed Integer Second Order Conic Optimization (MISOCO): Disjunctive Conic Cuts

Tamás Terlaky
Lehigh University

Second Order Conic Optimization (SOCO) has been the subject of intense study in the past two decades. Interior Point Methods (IPMs) provide polynomial time algorithms in theory, and powerful software tools in computational practice.

The use of integer variables naturally occur in SOCO problems, thus the need for dedicated MISOCO algorithms and software is evident. In this project we design design efficiently computable disjunctive conic cuts for MISOCO problems. The novel disjunctive conic cuts may be used to design Branch-and-Cut algorithms for MISOCO. Finally, some illustrative, preliminary computational results as presented.

Joint work with: Julio Gez and Ted Ralphs (Lehigh U.), Pietro Belotti1 (Clemson U.) Imre Plik, (SAS)

New Developments in Uncertainty: Linking Risk Management, Reliability, Statistics and Stochastic Optimization

Stan Uryasev
University of Florida

Terry Rockafellar
University of Washington

Decisions often have to be made in the face of uncertainty. A structure must be designed to withstand forces for which there is only statistical knowledge, at best. A system for responding to threats like oil spills, disease outbreaks, or terrorist acts must be put in place without knowing in advance the timing, location and magnitude of the incidents that will occur. The decisions in such situations must nonetheless confront constraints about the uncertain outcomes that flow from them, and they need to be open to comparisons which enable some kind of optimization to take place. In engineering, the reliability of a design is especially important and has generated a large literature. Serious difficulties have emerged, however, not only in assessing reliability and how it can be controlled by the means available, but even in the workability of the models that have been adopted. The central idea of "probability of failure," for instance, is inherently troublesome for its mathematical behavior when utilized in optimization, and this brings into question the soundness of the "design solutions" that may be calculated in terms of it.

Measures of risk can be used to extract from a random variable that stands for a hazard cost, or loss, a single quantity which can substitute for it in risk management and optimization. Risk in this sense can be derived from penalty expressions of regret about the mix of potential outcomes of such a random variable by means of tradeoffs between an up-front level of cost and the uncertain residual. The concept of regret mirrors, for losses, that of utility for gains relative to a benchmark. Statistical estimation is inevitably a partner with risk management in handling hazard variables, which may be known only through a data base, but a much deeper connection has come to light with statistical theory itself, in particular regression. Very general measures of error can associate with any hazard variable a "statistic" along with a "deviation" which quantifes the variable's nonconstancy. Measures of deviation, on the other hand, are known to be paired closely with measures of risk exhibiting aversity. A direct correspondence can furthermore be identified between measures of error and measures of regret. The fundamental quadrangle of risk puts all of this together in a unifed scheme.

Efficient Algorithms for Large-Scale Convex Optimization

Stephen Vavasis
University of Waterloo

Recent results by our group and others have shown that convex optimization is able to solve NP-hard problems in data mining provided that the input instances are constructed in a certain way. The resulting convex problems can be solved in polynomial time. However, due to the immense size of the original data mining instances, the resulting convex programming problems may have millions or even billions of variables rendering them intractable for state-of-the-art interior point methods. We will present recent progress from our group on efficient algorithms for large-scale convex optimization including conjugate gradient, proximal-point algorithms and facial reduction.

This talk represents joint work with Y.-L. Cheung, X. V. Doan, S. Karimi, N. Krislock, K.-C. Toh, H. Wolkowicz.

Exploiting Elementary Landscape analysis to improve search on NP-Hard Problems

Darrell Whitley, Adele Howe
Colorado State University

Much of the work on Elementary Landscapes has focused on what can be learned about the structure of the search landscape. There has been little work showing how this structure can be exploited so as to improve combinatorial search algorithms. In this talk, we present two new results. 1) We have been able to exploit the decomposability of the Traveling Salesman Problem to improve on the state of the art algorithms. The improvement is most pronounced in clustered instances, which have proven harder for the state of the art algorithms. 2) The decomposability of MAX-kSAT and NK-Landscapes has supported our development of constant $O(1)$ average time complexity next ascent and approximate steepest ascent search methods. This means we can select the best improving move for local search without examining neighbors, and we can analytically determine that we have arrived at a local optimum without examining neighbors.

Game Theoretic Search and Detection

Kevin Wood
Naval Postgraduate School

An *evader* wishes to travel between vertex s and vertex t in a network $G = (V, E)$. To detect this evader, an interdictor controls m_r *searchers* of each type $r \in R$. The evader is detected with probability p_{er} if he traverses edge $e \in E$ when a searcher of type r is present. The interdictor wants to position his searchers, with at most one searcher per edge, so as to maximize probability of detection (P_D). Under reasonable assumptions, this situation describes US forces attempting to detect an enemy that is infiltrating a countrys border region using known roads and tracks; searchers correspond to ground teams, UAVs, electronic detectors, etc. Key contributions of this research are the accurate modeling of multiple searcher types and multiple detection opportunities.

Given a scenario of ongoing infiltration, we define a two-person zero-sum game between evader and interdictor: the evader wishes to minimize P_D by choosing an $s-t$ path according to a probabilistic strategy, and the interdictor wishes to maximize P_D with a probabilistic strategy that assigns searchers to network edges. Assuming Poisson detections, a polynomially sized linear program optimizes P_D in an apparently relaxed formulation: variables assign a fraction of search effort for each searcher type to each edge, while constraints require that at most one searcher be assigned to each edge, on average. Following Bechhofer and von Neumann, we show that this relaxed solution is implementable as an exact solution to the original game through set of at most $|E|^2$ pure search strategies, each having at most one searcher per edge. A polynomial-time solution for the overall model therefore results.

When detections are not Poisson and the number of inspectors is large compared to the size of G, the model described above inappropriately accumulates multiple detections. We devise a finite row-and-column generation algorithm with integer-programming subproblems to solve this more complex problem. Polynomial solvability arises for interesting special case, but not in general.

The methods are demonstrated on a realistic, unclassified model covering a border area in Afghanistan. This is joint work with Alan Washburn and Lee Ewing at the Naval Postgraduate School.