



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

In This Issue

To cite this article:

(2016) In This Issue. Operations Research 64(6):iii-vi. <https://doi.org/10.1287/opre.2016.1568>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2016, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

In This Issue

Strong SOCP Relaxations for Optimal Power Flow Problem

AC optimal power flow (OPF) is the most fundamental optimization problem that needs to be solved every 5–15 minutes in the operation of electric power systems. It can be formulated as a non-convex quadratically constrained quadratic program. Despite of many years' extensive study, solving OPF in real time with a guarantee of solution quality still remains a key challenge. In "Strong SOCP Relaxations for the Optimal Power Flow Problem," Kocuk, Dey, and Sun propose three strong second order cone programming (SOCP) relaxations for AC OPF, which are incomparable to each other and two of which are incomparable to the standard SDP relaxation. Extensive computational experiments show that the solution quality of the SOCP relaxations is extremely close to the SDP relaxation and the solution time is 10 times faster than the SDP relaxation. A real-world 3375-bus system can now be solved by the proposed methods with a provable 0.13% global optimality gap within 157.20 seconds on a modest personal computer. This provides a practical approach to solve AC OPF with extremely good solution quality for the real-time operation of power systems.

Models and Methods for Testing Collective Rationality of Consumption Behavior

When given data corresponding to household consumption, it is by no means a trivial task to test whether particular collective generalizations of GARP (the Generalized Axiom of Revealed Preference) or SARP (the Strong Axiom of Revealed Preference) are satisfied. In "Revealed Preference Tests of Collectively Rational Consumption Behavior: Formulations and Algorithms," Talla Nobibon, Cherchye, Crama, Demuynck, De Rock, and Spieksma provide explanation for the perceived difficulty of this task: it is shown that the corresponding decision questions are NP-complete. Further, they analyze integer programming models for these tests of revealed preference that can be used for data sets of medium size. Finally, they propose a simulated annealing algorithm to test the collective model, which is particularly suited for large data sets.

The Paradox of Unlicensed Spectrum

The surge in demand for wireless services has generated interest in making additional bands of unlicensed (also referred to as open access or commons spectrum) radio spectrum available. This permits access by any device (e.g., for WiFi access) that abides by restrictions, such as a limit on transmit power. Unlicensed spectrum would encourage market expansion and increased competition among providers but it would increase congestion (interference) among consumers of wireless services. In "The Cost of Free Spectrum," Nguyen, Zhou, Berry, Honig, and Vohra consider a congestion pricing game, where agents choose a provider based on the minimum delivered price; the weighted sum of the price of the service and a congestion cost. They show that the social welfare

depends on the amount of additional unlicensed spectrum, and can actually decrease over a significant range of unlicensed bandwidths. With nonuniform weighting, introducing unlicensed spectrum could reduce consumer welfare.

Sharing Costs to Optimize Equilibria

The problem of coordinating independent individuals who compete for the resources of a system is ubiquitous in many important applications, from routing Internet traffic to shipping freight through transportation hubs. Due to the sheer size and the dynamic nature of these networks, a policy maker who is interested in the welfare of the system's users is often restricted to defining local policies, such as packet forwarding priorities or prices for the transportation of freight, which will lead the users of the network to an efficient equilibrium. In "Optimal Cost-Sharing in General Resource Selection Games," Gkatzelis, Kollias, and Roughgarden prove that for a very large spectrum of applications, the policy that corresponds to sharing the induced resource costs using the Shapley value is optimal.

Optimal Selection of Loan Portfolios

Although portfolio optimization has been widely studied for stocks, loans such as mortgages, credit cards, auto loans, or student loans have received very little attention. The optimal selection of loan portfolios is a large-scale nonlinear integer optimization problem and has several unique computational challenges: Loan portfolios can contain thousands to even hundreds of thousands of loans. Furthermore, the optimization must account for the loan-level data features of each individual loan (e.g., FICO score, interest rate, debt-to-income ratio, LTV ratio, etc.) which specify each loan's delinquency and prepayment behavior. In "Large-scale Loan Portfolio Selection," Sirignano, Tsoukalas, and Giesecke address these challenges by developing an approximate optimization approach based upon a weak convergence analysis. This approximate optimization method is typically orders of magnitude faster than nonlinear integer program solvers while also being highly accurate even for moderate-sized portfolios. The paper demonstrates the approach using actual mortgage data.

Border Patrols

In recent months the problem of how to infiltrate a border, and the opposing problem of how to defend it against infiltration, have received much attention in the news. In "Patrolling a Border" by Papadaki, Alpern, Lidbetter, and Morton, this problem is modelled as game between the Patroller and the Infiltrator. The border is modelled as a line graph with n nodes. The nodes might be the points of potential infiltration or more generally places where the network might be attacked. The overall framework in the analysis is the general model of "Patrolling Games" introduced earlier in this journal by Alpern, Morton and Papadaki to analyse arbitrary

networks. It turns out that the optimal strategies for both players are highly sensitive to the number of nodes of the line graph and the time required to infiltrate the network (or equivalently, to attack a node). The problem with simply patrolling back and forth from end to end is that this strategy leaves the end nodes particularly vulnerable to attack.

Managing Perishable Inventories When They are Depleted Last-In–First-Out (LIFO)

Perishable inventories in supermarkets such as milk, juice, packaged meat, etc. have a short life span and are mostly sold on a LIFO basis. How should a supermarket replenish its perishable inventories and clear them when they are approaching their expiration dates? In “Managing Perishable Inventories in Retailing: Replenishment, Clearance Sales, and Segregation,” Li, Yu, and Wu show that in both optimal policy structure and approximation, the age group of inventory with a remaining lifetime of one period plays a critical role. In particular, clearance sales may occur when its level is too high or too low. The phenomenon that a clearance sale happens when the inventory is low is driven by the need to segregate the newest inventory from the oldest inventory and is unique to the LIFO issuing rule. The optimal policy requires a full inventory record of every age group and its computation is challenging. We consider two myopic heuristics that require only partial information. The first requires only the information about the total inventory and the second requires the information about the total inventory as well as the information about the inventory with a one-period remaining lifetime. Our numerical studies show that the second outperforms the first significantly and its performance is consistently very close to that of the optimal policy.

Repairable Stocking and Expediting in a Fluctuating Demand Environment: Optimal Policy and Heuristics

Demand intensity for repairable spare parts fluctuates over time due to maintenance programming and operating conditions. These fluctuations can be buffered by stocking spare parts and by expediting the repair of a failed spare part when necessary. In “Repairable Stocking and Expediting in a Fluctuating Demand Environment: Optimal Policy and Heuristics,” Arts, Basten and van Houtum model demand fluctuations with a Markov Modulated Poisson process and characterize optimal expediting policies. They also study the asymptotics of this system for increasingly slow demand fluctuations and show a decomposition result into a weighted average of stationary systems. These asymptotics are leveraged to show that ignoring demand fluctuations can be detrimental when ignored. These asymptotics also lead to computationally efficient heuristics with average optimality gaps below 0.33% in a large numerical experiment. Comparison with naive heuristics shows that there is great value in leveraging knowledge about demand fluctuations when making repairable expediting and stocking decisions.

Simple Policies for Inventory Systems with Dual Delivery Modes and Fixed Ordering Costs

Efficient inventory management when multiple delivery modes are available requires managers make ordering decisions that carefully balance costs that arise from ordering, shipping and receiving as well as those from inventory holding and shortage. In “Continuous-review (R, nQ) Policies for Inventory Systems with Dual Delivery Modes,” Zhou and Yang propose a class of

simple policies called single-index (R, nQ) policies to manage inventory systems with dual delivery modes. Ordering from either mode incurs a fixed cost and the expedited mode provides a shorter lead time than the regular mode. The paper provides an exact procedure to compute the expected long-run average cost and proposes two simple heuristics for computing near-optimal policy parameters. The single-index (R, nQ) policy is shown to perform very well via numerical comparisons with the more complicated dual-index (R, nQ) policies and the policy computed via dynamic programming.

Supply Chain Coordination with Multiple Shipments

In many supply chains involving a supplier-retailer relationship, the retailer has to place an order prior to the selling season while the order can be delivered in multiple batches during the season. In “Supply Chain Coordination with Multiple Shipments: the Optimal Inventory Subsidizing Contracts,” Chen, Lee, and Moizadeh study the retailer’s optimal order quantity and timing for each shipment in a centralized supply chain, as well as coordination and incentive alignment in a decentralized one. They show that with multi-shipment, incentive conflicts arise from not only the double marginalization effect, but also ineffective inventory cost allocation between the parties. Hence, certain inventory subsidizing scheme is key for coordination: it can be implemented either by monetary subsidies for the retailer’s carrying inventory or by allowing the retailer to pay only a portion of the full payment upon receipts of the shipments, namely a delayed payment scheme.

How to Optimally Manage an Assembly System with Expediting and Advance Demand Information

Advance demand information allows a company to shift its production from build-to-stock to build-to-order, and thus deal with demand variability with lower inventory requirements. Lower inventory levels, however, leave the supply chain vulnerable to high realizations of demand. One strategy to mitigate exposure to high demand realizations is to allow expediting of inventory. With both advance demand information and expediting of stock, supply chains with assembly structure become imbued with the severe curse of dimensionality. In “Knowledge You Can Act On: Optimal Policies for Assembly Systems with Expediting and Advance Demand Information,” Angelus and Özer develop a new approach for solving complex assembly systems. They identify the structure of the optimal policy for such an assembly problem that reduces its state space significantly. Advance demand information and expediting of stock are found to be complementary under short demand information horizons and substitutes under longer information horizons.

How Much Can Delivery Services Help Reduce Carbon Emissions?

What happens to the carbon footprint in a geographical region when people make use of last-mile delivery services? How much efficiency is gained when one performs several errands in a single outing, as opposed to performing them individually? In “Household-level Economies of Scale in Transportation,” Carlsson, Behroozi, Devulapalli, and Meng show how to use a continuous approximation of the “generalized travelling salesman problem” to answer these questions.

Exact Routing of Electric Commercial Vehicles

The effective employment of new technologies like, e.g., electric commercial vehicles (ECVs) for commercial transport, is key to a widespread adoption of such technologies. The electric vehicle-routing problem with time windows (EVRPTW) models the delivery operation of a homogeneous fleet of capacitated ECVs to a set of customers with given delivery time windows. In “Exact Algorithms for Electric Vehicle-Routing Problems with Time Windows,” Desaulniers, Errico, Irnich, and Schneider describe branch-price-and-cut-algorithms for four variants of the EVRPTW, which are defined by the possible combinations of two attributes: (i) the number of recharges per route (single or multiple), and (ii) the type of the recharge operation (full recharge or partial recharge). Numerical studies show that the algorithms using bidirectional labeling for generating feasible vehicle routes are superior to the ones with monodirectional labeling, and that multiple and partial recharges both help to reduce routing costs.

Approximating the Distribution of Project Completion Time with Random Activity Durations

One of the fundamental problems in project management is to estimate the project completion time when the activity durations are random. This is a special case of the more general problem of estimating the distribution of the solution value to a mixed binary linear optimization problem, where the objective coefficients are random. In “Least Squares Approximation to the Distribution of Project Completion Times with Gaussian Uncertainty,” using the classical Stein’s Identity in probability theory, Zheng, Natarajan, and Teo show that the optimal least squares linear and quadratic estimators of the random optimal value can be solved in closed form, and are related to the Persistency Problem studied by Bertsimas et al. (2006). We apply our methods to the problems in project management and statistical timing analysis, and show that the new approach provides more accurate estimates of the distributions compared to existing methods.

Can Speed-Ups for Stochastic Linear Programming (SLP) Outpace Moore’s Law?

About 50 years ago, Gordon Moore, the co-founder of Intel predicted that the density of transistors in integrated circuits would continue to double every 18 months or so. This remarkable trend has spawned inexpensive and powerful computing platforms. In “Mitigating Uncertainty via Compromise Decisions in Two-stage Stochastic Linear Programming: Variance Reduction” by Sen and Liu, the algorithmic progress discussed in the paper reports speed-ups (after accounting for hardware changes) for SLP problems by a factor of over 128, a rate that exceeds Moore’s Law for a decade. Interestingly, this speed-up is also accompanied by reduced variance decisions when compared with the highly cited sample average approximation. The key here lies in the scalable integration of statistical methodology (e.g. bootstrapping) with optimization algorithms (e.g. stochastic bundle methods). This innovation promotes a simultaneous view of predictive and prescriptive models, and will spawn a truly integrative approach to analytics.

Managing Waiting Lines of Selfish Agents

How to minimize the expected costs for managing a waiting line of customers when the information about their necessary service

time and urgency is only known to the customers themselves? This question is answered in the paper “Optimal Mechanism Design for a Sequencing Problem with Two-Dimensional Types” by Hoeksma and Uetz. The paper asks for a Bayes-Nash incentive compatible mechanism that minimizes expected costs for sequencing job-agents when their weights *and* processing times are private information. As their main result, they show how to solve this problem efficiently by linear programming techniques. The paper adds to the literature on optimal mechanism design, which generally asks to solve optimization problems in the face of private data and actions of the participating agents. The seminal publication of R. B. Myerson on “Optimal Auction Design” (1981) marks an important milestone in that area, and the question if and how optimal mechanisms can be designed for problems in which agents have multidimensional types has spurred a lot of attention since then.

Regularization Methods in Optimization with Stochastic Orders

Stochastic orders formalize preferences among random outcomes and are widely used in statistics and economics for comparison of random quantities. The inherent uncertainty in many decision-making problems creates the need for mathematical models of risk that allow the control of the associated risk. Stochastic models using stochastic-order relations as constraints constitute an effective approach to risk-averse optimization, allowing for precise shaping of the probability distributions of interest. In the paper “Augmented Lagrangian Methods for Solving Optimization Problems with Stochastic-Order Constraints,” by Dentcheva, Martinz, and Wolfhagen, optimization problems with stochastic-order constraints on vector-valued random outcomes are studied and their efficient numerical solution is addressed. Several numerical methods based on augmented Lagrangian framework are developed and their convergence is analyzed. The proposed methods construct a sequence of finite-dimensional risk-neutral approximations to address the infinite-dimensional nature of the stochastic-order constraints. Special attention is paid to the case of a single un-certain outcome, in which an augmented Lagrangian method using the quantile representation of the stochastic dominance relation is presented.

Finding Large 2-Hop Clusters

The growing interest in network analysis and mining of complex networks has reignited the study of clustering problems in many disciplines. An extensive body of literature exists, in fields as diverse as physics, biology, sociology, finance, mathematics, and engineering. In “On the 2-Club Polytope of Graphs,” Mahdavi Pajouh, Balasundaram, and Hicks focus on a model that formalizes the notion of a “friend-of-a-friend cluster” known as a 2-club. Introduced in the field of social network analysis, it is based on the premise that two strangers could be included in a cohesive social subgroup if they have a common friend in that cohesive subgroup. Finding large 2-clubs turns out to be a challenging computational problem. The authors provide theoretical results that unify several previous results yielding new facet-inducing inequalities that play an important role in algorithms for this problem. The computational characteristics of the newly discovered inequalities are studied theoretically and empirically by the authors.

Optimal Self-Scheduling Bidding Strategies for Independent Power Producers

With the increasing penetration of renewable energy into the power grid system, the volatility of real-time electricity prices increases significantly. This brings challenges for independent power producers to provide optimal bidding strategies. In “Strong Formulations for the Multistage Stochastic Self-Scheduling Unit Commitment,” Pan and Guan explore optimal self-scheduling strategies to participate in the real-time market considering price volatility, with the objective of maximizing the total expected profit. A multistage stochastic scenario tree is studied to capture the price uncertainties, and accordingly the derived multistage stochastic self-scheduling unit commitment problem is transformed as a deterministic equivalent mixed-integer linear programming formulation. Strong formulations are derived and the corresponding discovered strong valid inequalities can provide the convex hull descriptions for the two-period case and a special class of the three-period cases. Numerical experiments show the promising results of the derived strong valid inequalities by incorporating them in a branch-and-cut framework.

Designing Selection-of-The-Best Procedures Without Indifference Zone

Selection-of-the-best problems are classic in the simulation society. To select the unique best alternative, many existing frequentist procedures are designed based on the indifference-zone (IZ) formulation. Their performance hinges on the relationship between the specified IZ parameter and the true mean differences that is unknown to users. In “Indifference-Zone-Free Selection of the Best,” Fan, Hong and Nelson remove the requirement of IZ parameter and design a new class of sequential IZ-free procedures based on the Law of the Iterated Logarithm. They further add a stopping criterion to these procedures to convert them into IZ procedures. Extensive numerical study shows that the designed procedures outperform the existing procedures.

Two Fundamental Simulation Optimization Methods are Asymptotically Equivalent

The expected improvement (EI) class of methods is a fundamental and well-established methodology for efficient information collection in stochastic and simulation-based optimization. In particular, EI has been a standard workhorse algorithm in global optimization for nearly 20 years. However, despite its repeatedly demonstrated practical potential, EI has thus far not been amenable to any of the standard mathematical techniques for convergence rate analysis. The paper “On the Convergence Rates of Expected Improvement Methods,” by Ryzhov, presents the first rate analysis for EI-type methods in the widely-studied setting of ranking and selection with normally distributed noise. Rate characterizations are obtained for several variants of EI. When the variance of the simulation output is known, one variant is shown to sample suboptimal alternatives at the same relative rates as another prominent algorithmic family known as optimal computing budget allocation (OCBA), thus providing the first bridge between these two research communities.

Base-Stock List-Price Policy is Optimal When Customers are Loss-Averse

As firms dynamically adjust prices, customers form reference price from past prices to judge against the selling price. Customers

are often loss-averse, and thus more sensitive to price increase (loss) than price decrease (gain) relative to the reference price. Incorporating loss-averse behavior into integrated pricing and inventory models is important yet poses significant analytical challenges due to the lack of concavity of the single period expected profit in price and reference price. In “Dynamic Stochastic Inventory Management with Reference Price Effects,” Chen, Hu, Shum and Zhang study a single product integrated pricing and inventory model in which demand is a linear function of the current selling price plus a concave piecewise linear function of the difference of the reference price and the selling price. Despite of the lack of concavity of the single period expected profit, they prove that a base-stock list-price policy is optimal, which is enabled by a novel transformation technique.

A New Method for Solving Large Scale Markov Decision Processes

Markov decision processes (MDPs) have historically been a major area of focus in operations research and have been used to model problems arising in many different areas such as revenue management, finance and healthcare. However, practical MDPs are in general extremely difficult to solve exactly, due to the state space being extremely large in most practical problems. In “Decomposable Markov Decision Processes: A Fluid Optimization Approach,” Bertsimas and Mišić identify a broad class of MDPs, called decomposable MDPs, where the system can be viewed as a collection of subsystems, and propose a novel linear optimization formulation for this class of problems. They show theoretically that this model is stronger than three state-of-the-art formulations: the approximate linear optimization model, the classical Lagrangian relaxation and a new alternate Lagrangian relaxation. In numerical experiments with both regular and restless multiarmed bandit problems, they show that a heuristic derived from their formulation can significantly outperform alternative proposals.

Constant-Order Policy Nearly Optimal for Lost Sales Inventory Models with Small-to-Moderate Lead Times

Lost sales inventory models, which arise in many practical settings, become notoriously challenging to optimize in the presence of positive lead times, since the state-space blows up and dynamic programming techniques become intractable. Recently, Goldberg and co-authors laid the foundations for a new approach to solving these models, by proving that as the lead time grows large, a simple constant-order policy is asymptotically optimal. However, their results required the lead time to be impractically large, and deriving tighter bounds was left as a significant open problem. In “Optimality Gap of Constant-Order Policies Decays Exponentially in the Lead Time for Lost Sales Models,” Goldberg and Xin take a large step towards resolving this problem, by proving that the optimality gap of the constant-order policy actually converges exponentially fast to zero. The authors also derive simple and explicit bounds for the optimality gap by combining ideas from convexity, queueing, and the theory of random walks, and demonstrate good numerical performance.