



Operations Research

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

In This Issue

To cite this article:

(2019) In This Issue. Operations Research 67(1):ii-iv. <https://doi.org/10.1287/opre.2019.1855>

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 © 2019, 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

<https://doi.org/10.1287/opre.2019.1855>

Copyright: © 2019 INFORMS

Optimal Control of Large-Scale Energy Storage

Electricity storage is likely to play a significant role in balancing future energy systems. Often, much of the value of large-scale storage (e.g., pumped storage and hydro) may be captured in price arbitrage. In “Control of Energy Storage with Market Impact: Lagrangian Approach and Horizons,” Cruise, Flatley, Gibbens, and Zachary study the optimal control of storage that is sufficiently large as to impact on the market in which it operates and that seeks to maximize the profit that can be made by buying and selling over time. They develop the associated strong Lagrangian theory, which is of economic significance when storage is embedded in wider systems, and provide forecast and decision horizons for the optimal control, together with an algorithm suitable for control over an indefinite time period.

Promoting Effectiveness and Equity in the Daily Logistic Operations of Food Banks

Increasing amounts of excess food that is perfectly edible but would otherwise be destroyed are transferred every day from food banks, supplied by donors in the supply chain, to welfare agencies that serve the needy. In “The Humanitarian Pickup and Distribution Problem,” Eisenhandler and Tzur study such operations, motivated by the activity of nongovernmental organizations in Israel and in the United States. This setting gives rise to a complex logistic problem, with simultaneous vehicle routing and resource allocation decisions. An innovative objective function, based on the Gini index, is used to balance two colliding goals: maintaining equitable allocations to the different agencies and delivering as much food as possible. This function satisfies desired properties of the allocation and is easy to compute and implement within a mathematical formulation. A heuristic procedure is developed and shown to produce high-quality solutions for real-life as well as randomly generated data sets.

An Impulse Control Approach to Capacity Expansions

In “Sequential Capacity Expansion Options,” Bensoussan and Chevalier-Roignant consider a firm expanding its production capacity in stages under uncertainty. The firm decides on the investment times and the size of the investment lumps. This setting gives rise to an impulse control problem for which we derive a quasi-variational inequality that involves two state variables: a stochastic price process and a controlled capacity process. The authors provide a general verification theorem and construct an optimal two-dimensional (s, S) -type policy for a specific case. The firm “waits and sees” before investing until the perpetuity value of newly installed capacity exceeds the total opportunity cost by a sufficient

margin. The authors’ model generalizes extant models dealing with “the option to expand” capacity and makes predictions of both the optimal investment timing and the optimal scale of production.

Forward Commodity Trading with Private Information

Firms that trade a commodity that has a random price can reduce risk by trading forward. In “Forward Commodity Trading with Private Information,” Anderson and Philpott consider a setting in which a buyer and a seller of a divisible commodity have different private information on the probability distribution of its future price. This situation arises in electricity pool markets where contracts for differences are traded by buyers and sellers of electricity to hedge future price risk. The paper compares several mechanisms for settling the price and quantity of such a contract when buyers and sellers have private information. These include Nash bargaining and supply-function equilibrium models. In the latter setting, a player can deduce the other player’s private information from its offered supply function and use this to improve its own supply function. Examples are given to show that this strategy, when used by both firms, may make both firms worse off in equilibrium.

Learning from the Adversary: An Approach to Interdiction with Incomplete Information

In “Sequential Interdiction with Incomplete Information and Learning,” Borrero, Prokopyev, and Sauré study a general class of interdiction problems when the leader has incomplete information regarding the formulation solved by the follower and both interact repeatedly over time. In the setting, the leader learns by observing the follower’s reactions to the leader’s actions. The authors show that strong notions of optimality are not attainable, but a form of weak optimality is attained by the set of proposed policies. These policies are greedy and robust in the sense that the leader reacts “optimally” in each period based on the information at hand and takes a robust approach to handling missing information. The policies are shown to provide a real-time certificate of optimality and to consistently outperform a reasonable benchmark in a series of numerical experiments

The Newsvendor Meets Machine Learning: New OR Approaches for Big Data

In “The Big Data Newsvendor: Practical Insights from Machine Learning,” Ban and Rudin take an innovative machine-learning approach to a classic problem solved by almost every company, every day, for inventory management. By allowing companies to use large amounts of data to predict

the *correct answers to decisions directly*, they avoid intermediate questions, such as “how many customers will we get tomorrow?” and instead can tell the company how much inventory to stock for these customers. This has implications for almost all other decision-making problems considered in operations research, which has traditionally considered data estimation separately from the decision optimization. Their proposed methods are shown to work both analytically and empirically, with the latter explored in a hospital nurse staffing example in which the best one-step, feature-based newsvendor algorithm (the kernel-weights optimization method) is shown to beat the best-practice benchmark by 24% in the out-of-sample cost at a fraction of the speed.

Supply-Then-Price Competition Yields Cournot Equilibria in Differentiated Product Markets Under General Demand and Spill Conditions

Supply-then-price competition is common in settings in which capacity commitments precede pricing decisions. In “On the Relationship Between Quantity Precommitment and Cournot Games,” Farahat, Huh, and Li provide a characterization of when and why equilibria of such games clear the market at the Cournot outcomes. The paper extends the seminal work of Kreps and Scheinkman (1983) to differentiated product settings with an arbitrary number of firms and general demand and spill functions. Farahat et al. identify two new fundamental properties that drive the result: independence from irrelevant supply and spill inertia. These properties are satisfied by common demand and spill models. The authors’ result provides a shortcut approach to analyzing supply-then-price games and highlights the importance of demand specification, compared with spill specification, in determining the outcome of such games.

Determining Work Packages for Projects

In organizing a project’s tasks into manageable work packages (i.e., forming a work breakdown structure), trade-offs arise. Defining smaller work packages increases project complexity and workload, and reduces economies of scale, whereas defining larger work packages reduces concurrent processing and adversely affects cash flow. In “Work Package Sizing and Project Performance,” Li and Hall study this trade-off by developing and analyzing an optimization model for work package formation. The model minimizes total project cost, subject to a deadline constraint on project makespan. From a study of this model, the authors demonstrate the value of deliberately varying work package sizes within a project, in contrast with typical project management practice. This research enables more precise planning of work packages to improve performance, documents the value of integrating the planning of work packages and schedules, and provides insights that guide resource allocation decisions.

Routing Thousands of Vehicles in Real Time: Scaling Integer Optimization

With the emergence of ride-sharing companies that offer transportation on demand at a large scale and the increasing availability of corresponding demand data sets, new challenges arise to develop routing optimization algorithms

that can solve massive problems in real time. In “Online Vehicle Routing: The Edge of Optimization in Large-Scale Applications,” Bertsimas, Jaillet, and Martin present a novel and generalizable backbone algorithm that uses integer optimization to find high-quality solutions to large routing optimization problems. The algorithm, together with the real-time routing optimization software framework developed and shared by the authors, can dispatch thousands of taxis serving more than 25,000 customers per hour. An extensive study with historical simulations of Yellow Cabs in New York City is used both to show that the algorithm improves on the performance of existing heuristics and to provide insights on the optimization opportunities of a large ride-sharing fleet.

Discrete Convexity in Operations Management

Discrete convexity, which extends submodularity to integer vectors, has been used in the economics and management literature to characterize the behavior of optimal policies. One of its variants, called L^1 -convexity, has enabled recent advances in various operations management systems. In a paper published by *Operations Research* in 2005, an assemble-to-order inventory system was shown to have the L^1 -convexity property, which was used to motivate an efficient algorithm. In a technical note, “Error Noted in ‘Order-Based Cost Optimization in Assemble-to-Order Systems’ by Lu and Song (2005)” by Bolandnazar, Huh, McCormick, and Murota, the authors show that the proof in that paper is incorrect and L^1 -convexity may not hold. Despite this error, the authors credit Lu and Song for introducing this useful concept to the operations management community.

System Optimum Dynamic Traffic Assignment in General Networks

Most current system optimum dynamic traffic assignment (SO-DTA) models do not contain first-in-first-out (FIFO) constraints and are limited to single-destination network applications. In “Link-based System Optimum Dynamic Traffic Assignment Problems in General Networks,” Long and Szeto develop SO-DTA models either with or without FIFO constraints for general network applications using the concept of the link transmission model (LTM). The proposed SO-DTA problems without FIFO constraints are formulated as linear programs and can address the vehicle holding problem by adding a penalty term to the objective function. The SO-DTA problems with FIFO constraints are formulated as nonconvex mixed-integer programs and are solved by branch-and-bound algorithms. Two methods are developed to identify FIFO violations in feasible flow patterns and design a branching scheme for the proposed branch and bound algorithms. The authors show that the proposed LTM-based SO-DTA models are more computationally efficient than the cell transmission model (CTM)-based counterparts and maintain the same level of accuracy in terms of guaranteeing the same total system travel time in general networks.

Oriented Measures of Profit Efficiency

As the most popular measures of technical efficiency, Farrell measures have not been incorporated explicitly in the decomposition of profit efficiency. By synthesizing existing measures of profit efficiency, the paper “A Unifying

Framework for Farrell Profit Efficiency Measurement,” by Färe, He, Li, and Zelenyuk establishes a general oriented model that includes many existing profit efficiency measures as special cases. Furthermore, the oriented measures generated from this model have consistent meaningful interpretation. This is also perhaps the first time in the literature that the Farrell technical efficiency and allocative efficiency are included as multiplicative components of profit efficiency. A new component is introduced to help decision makers distinguish input changes from output changes in improving profit efficiency.

Efficient Learning for Large Resource Allocation Problems

Approximate dynamic programming has been applied to solve large-scale resource allocation problems in many domains, including transportation, energy, and healthcare. The main algorithmic challenge in these applications is to accurately estimate the value of a resource in a certain state (for example, a battery that is half full, or a truck en route to a certain destination). This “value function” must be learned as part of the decision-making process, giving rise to the “exploration/exploitation” challenge: we may, and often should, choose to experiment with seemingly suboptimal actions because they have high potential to be better than we believe. In “Bayesian Exploration for Approximate Dynamic Programming,” Ryzhov, Mes, Powell, and van den Berg present a principled framework for modeling uncertainty about the value function and measuring the potential of a state to improve a decision-making policy. This method is scalable and performs well in experiments.

Why Do Clustering Algorithms Perform So Well in Practice?

In “Good Clusterings Have Large Volume,” Borgwardt and Happach study the strong contrast between the favorable performance of clustering algorithms in practice and their bad theoretical worst cases. They explain this contrast using polyhedral theory. Classical clustering algorithms can be interpreted as the search for vertices in the so-called bounded-shape partition polytopes, following a “random” direction. The vertices correspond to clusterings with extraordinary separation properties, and the quality of this separation depends on the volume of the normal cone of the vertex—the larger the better. This gives rise to a new quality criterion for clusterings and explains why good clusterings are also the most likely to be found. The authors further characterize the edges of the polytopes, which allows for an explicit description of the normal cone, and in turn the computation of a best separation between the clusters of a given clustering.

Robustifying Stochastic Simulation to Capture Input Model Risks

Decision making using simulation analysis faces the risk of being adversely affected by insufficiently calibrated input models. This risk is especially significant when only limited historical observations of the real-world process are available, or when unexpected risk factors are suspected and require elicitation of expert opinion. In “Robust Analysis in Stochastic Simulation: Computation and Performance Guarantees,” Ghosh and Lam propose a new methodology that solves

robust optimization formulations over generic simulation models to compute bounds on the worst possible performance values. Precise performance guarantees are derived for the proposed methods, and these are further illustrated with key examples from stochastic service and queueing systems.

Gaussian Markov Random Fields for Discrete Optimization via Simulation: Framework and Algorithms

This paper lays the foundation for employing Gaussian Markov random fields (GMRFs) for discrete decision-variable optimization via simulation; that is, optimizing the performance of a simulated system. Gaussian processes have gained popularity for inferential optimization, which iteratively updates a model of the simulated solutions and selects the next solution to simulate by relying on statistical inference from that model. Salemi, Song, Nelson, and Staum show that, for a discrete problem, GMRFs, a type of Gaussian process defined on a graph, provides better inference on the remaining optimality gap than the typical choice of continuous Gaussian process and thereby enables the algorithm to search efficiently and stop correctly when the remaining optimality gap is below a pre-defined threshold. We also introduce the concept of multi-resolution GMRFs for large-scale problems, with which GMRFs of different resolutions interact to efficiently focus the search on promising regions of solutions.

Little’s Law for Periodic Queueing Systems

In “Periodic Little’s Law,” Whitt and Zhang extend the fundamental conservation law established by J. D. C. Little (1961) to settings in which both the arrival rates and waiting-time distributions are periodic. Both sample-path and steady-state stochastic versions are established in discrete time. These results help explain the authors’ data analysis of the patient flow of an emergency department. Just as for the time-varying Little’s law established by Bertsimas and Mourtzinou (1997), the time-varying average number in system depends on the full waiting-time distribution.

Novel Delay Scaling and Load Balancing for Parallel Server Systems

The exact study of control policies for queueing systems, such as prioritizing queues or load balancing among servers, is often notoriously intractable. The usual path to study these questions is an asymptotic analysis where either the system load is increased to its capacity or the size of the queueing system is increased (called many-servers regime) at the same time as its load is increased. In the paper “Load Balancing in the Nondegenerate Slowdown Regime,” Gupta and Walton argue that the traditional asymptotic regimes are insufficient for studying the properties of the most classical load-balancing policy, join the shortest queue (JSQ), where arriving jobs join the server with the fewest number of jobs. By studying load-balancing policies in the novel nondegenerate slowdown asymptotic regime, the authors conclude that JSQ achieves performance close to the ideal centralized policy. Further, the work proposes a simpler load-balancing policy, called idle one first, which achieves the same performance as JSQ but at a much smaller communication overhead between the servers and the load balancer.