



## Operations Research

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

### In This Issue

To cite this article:

(2020) In This Issue. Operations Research 68(6):ii-iv. <https://doi.org/10.1287/opre.2020.2063>

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](mailto: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 © 2020, 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.2020.2063>

Copyright: © 2020 INFORMS

### Simple Bayesian Algorithms for Best-Arm Identification

In “Simple Bayesian Algorithms for Best-Arm Identification,” Russo considers the optimal adaptive allocation of measurement effort for identifying the best among a finite set of options or designs. An experimenter sequentially chooses designs to measure and observes noisy signals of their quality with the goal of confidently identifying the best design after a small number of measurements. Just as the multiarmed bandit problem crystallizes the tradeoff between exploration and exploitation, this “pure exploration” variant crystallizes the challenge of rapidly gathering information before committing to a final decision. The author proposes several simple Bayesian algorithms for allocating measurement effort and, by characterizing fundamental asymptotic limits on the performance of any algorithm, formalizes a sense in which these seemingly naive algorithms are the best possible.

### Reducing Delay in Retrial Queues by Simultaneously Differentiating Service and Retrial Rates

Customer retrials commonly occur in many service systems, such as healthcare, call centers, mobile networks, computer systems, and inventory systems. However, because of their complex nature, retrial queues are often more difficult to analyze than queues without retrials. In “Reducing Delay in Retrial Queues by Simultaneously Differentiating Service and Retrial Rates,” J. Wang, Z. Wang, and Liu develop a service grade differentiation policy for queueing models with customer retrials. They show that the average waiting time can be reduced through strategically allocating the rates of service and retrial times without needing additional service capacity. Counter to the intuition that higher service variability usually yields a larger delay, the authors show that the benefits of this simultaneous service-and-retrial differentiation (SSRD) policy outweigh the impact of the increased service variability. To validate the effectiveness of the new SSRD policy, the authors provide (i) conditions under which SSRD is more beneficial, (ii) closed-form expressions of the optimal policy, (iii) asymptotic reduction of customer delays when the system is in heavy traffic, and (iv) insightful observations/discussions and numerical results.

### Optimal Reflection Control

The so-called reflection control is easy to implement and widely applied in many applications such as inventory management and financial systems. To apply reflection control in a production-inventory system, for example, production

will stop when the finished-goods inventory reaches a certain level. What is the best level for this control? In what sense is it optimal? In “Technical Note—On the Optimality of Reflection Control,” Yang, Yao, and Ye have established the optimality of reflection control under an exponential holding cost in three settings—namely, a Brownian motion model, a single-server system, and a birth–death queue model. The study provides a thorough understanding of the control and extends significantly its domain of applications.

### Sampling the Future in Monte Carlo Tree Search

In “Optimistic Monte Carlo Tree Search with Sampled Information Relaxation Dual Bounds,” Jiang, Al-Kanj, and Powell propose an extension to Monte Carlo tree search that uses the idea of “sampling the future” to produce noisy upper bounds on nodes in the decision tree. These upper bounds can help guide the tree expansion process and produce decision trees that are deeper rather than wider, in effect concentrating computation toward more useful parts of the state space. The algorithm’s effectiveness is illustrated in a ride-sharing setting, where a driver/vehicle needs to make dynamic decisions regarding trip acceptance and relocations.

### The Value of (Imperfect) Information and Memory in Dynamic Resource Allocation

In “Information and Memory in Dynamic Resource Allocation,” Xu and Zhong propose a general framework, dubbed stochastic processing under imperfect information (SPII), to study the impact of information constraints and memories on dynamic resource allocation. The framework involves a stochastic processing network (SPN) scheduling problem in which the scheduler may access the system state only through a noisy channel, and resource allocation decisions must be carried out through the interaction between an encoding policy (that observes the state) and an allocation policy (that chooses the allocation). Applications in the management of large-scale data centers and human-in-the-loop service systems are among our chief motivations. The authors quantify the degree to which information constraints reduce the size of the capacity region in general SPNs and how such reduction depends on the amount of memories available to the encoding and allocation policies. Using a novel metric, capacity factor, their main theorem characterizes the reduction in capacity region (under “optimal” policies) for all nondegenerate channels and across almost all combinations of memory sizes. Notably, the theorem demonstrates, in substantial generality, that (1) the presence of a noisy channel always reduces capacity, (2) more memory for the allocation policy always improves capacity, and (3) more memory for the encoding policy has little to no effect on

capacity. Finally, the authors positive (achievability) results are established through constructive, implementable policies.

### **There's No Free Lunch in Bilevel Optimization**

Linear bilevel problems are often reformulated as single-level problems by using the KKT optimality conditions of the lower level. The resulting KKT complementarity conditions are usually linearized by the classical big-M approach. However, this approach requires to bound the lower-level dual variables to obtain a correct single-level reformulation of the original bilevel problem. If the big-M value is not large enough, this yields to a reformulation admitting wrong solutions w.r.t. the original problem. In "Technical Note—There's No Free Lunch: On the Hardness of Choosing a Correct Big-M in Bilevel Optimization," Kleinert, Labbé, Plein, and Schmidt show that verifying that a big-M does not lead to cutting off any feasible vertex of the lower-level dual polyhedron cannot be done in polynomial time unless  $P = NP$ . Moreover, we prove that verifying that a big-M does not lead to cutting off any optimal point of the lower-level dual problem is as hard as solving the original bilevel problem.

### **Diffusion and Seeding in Random Networks**

Motivated by viral marketing in social networks, "Diffusion in Random Networks: Impact of Degree Distribution," by Manshadi, Misra, and Rodilitz studies the diffusion process of a new product on a network where each agent is connected to a random subset of others. The firm chooses a fixed number of agents to seed, knowing only the degree of each agent, and incurs a fixed cost per contact. Under such a setting, the authors exactly characterize the entire diffusion trajectory in the limit of network size without resorting to approximation methods. Using their limit results, the authors uncover a trade-off between cost-efficiency and fast growth. Further, the authors study the impact of degree distribution on optimal seeding strategies. Somewhat surprisingly, they show that to minimize cost, it is optimal to seed low-degree agents. Even if the objective is to minimize time, for certain regimes, the optimal seeding strategy is a mixture of low- and high-degree agents.

### **Asynchronous Schemes for Stochastic and Misspecified Potential Games and Nonconvex Optimization**

In "Asynchronous Schemes for Stochastic and Misspecified Potential Games and Nonconvex Optimization," Lei and Shanbhag consider a class of convex stochastic Nash games, possibly corrupted by parametric misspecification and characterized by a possibly nonconvex potential function. The authors present an asynchronous inexact proximal best-response (BR) scheme in which, at any step, a randomly selected player computes an inexact BR step (via stochastic approximation) and other players keep their strategies invariant. Misspecification is addressed by a simultaneous learning process reliant on an increasing batch size of sampled gradients. Almost-sure convergence guarantees are provided to the set of Nash equilibria, and such claims can be extended to delay-afflicted regimes, generalized potential games (with coupled strategy sets), and

weighted potential games. In fact, equilibria of this potential game are stationary points of the potential function and asynchronous inexact BR schemes are, in essence, randomized block-coordinate schemes for a subclass of stochastic nonconvex optimization problems.

### **Tractable Approximate Linear Programs for Discrete Optimization via Networks and Graph Completions**

Several prescriptive tasks in business and engineering as well as prediction in machine learning entail the solution of challenging discrete optimization problems. In "Network-Based Approximate Linear Programming for Discrete Optimization," Nadarajah and Cire recast the typical optimization formulation of these problems as high-dimensional dynamic programs and approach their approximation via linear programming. They develop tractable approximate linear programs with supporting theory by bringing together tools from state-space aggregations, networks, and perfect graphs (i.e., graph completions). The authors embed these models in a simple branch-and-bound scheme to solve applications in marketing analytics and the maintenance of energy or city-owned assets. They find that the resulting technique substantially outperforms a state-of-the-art commercial solver as well as aggregation-heuristics in terms of both solution quality and time. Their results motivate further consideration of networks and graph theory in approximate linear programming for solving deterministic and stochastic discrete optimization problems.

### **Resource Allocation and Pricing in the Absence of a Demand Forecast**

Motivated by collaborations with airline carriers and hotel chains, Ma and Simchi-Levi analyze in "Algorithms for Online Matching, Assortment, and Pricing with Tight Weight-Dependent Competitive Ratios" a general online resource-allocation problem using competitive analysis, where the sequence of demands is completely unknown. They allow for resources to be sold at multiple feasible prices, and establish tight competitive ratios that are dependent on the given feasible prices, generalizing the well-known online matching and Adwords settings. The algorithms derived are simple and intuitive, based on a class of universal value functions that integrate the resource selection and pricing decisions. Finally, the authors run simulations on publicly available data, and consider how a hotel could have managed its room resources through controlling the room packages offered on its website. The authors find that revenues are maximized when their algorithms (which do not assume anything about the sequence of visitors) are applied as a hybrid with existing algorithms that attempt to forecast and learn the sequence of visitors.

### **Managing Operational Risk**

Financial services firms are subject to various types of risks. In particular, operational risk is difficult to assess and can be devastating, although it is often perceived by a firm's management as being more controllable than the cost of managing other types of risks. Understanding the management problems associated with operational risk is crucial

to the performance of the firm. In “Operational Risk Management: A Stochastic Control Framework with Preventive and Corrective Controls,” Xu, Zhu, and Pinedo introduce a general modeling framework for operational risk management for financial firms. They propose two types of controls and characterize the optimal control policies. The authors apply their model to a data set from a commercial bank, and through a proper investment strategy, one can achieve a significant performance improvement.

### Maximum-Entropy Sampling

Maximum-entropy sampling is a difficult nonlinear discrete optimization problem that arises in spatial statistics, for example, in the design of weather-monitoring networks. An exact algorithm for maximum-entropy sampling was first described in 1995, and subsequent papers have devised a variety of methods that obtain bounds and, in some cases, exact solutions. In “Efficient Solution of Maximum-Entropy Sampling Problems,” Anstreicher describes a new bound for the maximum-entropy sampling problem that is superior to all previously known bounds and is also efficiently computable. A branch-and-bound algorithm that incorporates the new bound solves challenging benchmark instances to optimality for the first time.

### Input Efficiency Measures: A Generalized, Encompassing Formulation

In “Input Efficiency Measures: A Generalized, Encompassing Formulation” Briec, Cavaignac, and Kerstens develop a generalized, encompassing formulation unifying four traditional input efficiency measures: radial, Färe-Lovell, asymmetric Färe, and multiplicative Färe-Lovell. This is basically motivated by the fact that observations on production need not be situated near the efficient subset, but could also be positioned close to the isoquant or even the boundary of the technology. This new generalized Färe-Lovell input efficiency measure shares its axiomatic properties with the original Färe-Lovell input efficiency measure. In addition, the authors can derive new dual interpretations for this generalized Färe-Lovell input efficiency measure. Finally, they derive mathematical programming formulations, with a special focus on cases where linear programming applies.

### Large-Scale Parallel Bayesian Optimization

Bayesian optimization is a machine-learning-based method for optimizing time-consuming-to-evaluate black-box objective functions without derivatives. It is most widely used for optimizing deep learning models, tuning systems through A/B testing, aircraft designing, and developing new materials and drugs. Bayesian optimization relies on optimizing a fast-to-evaluate function, called the acquisition function, to decide where to evaluate the time-consuming objective. This is easy for single-threaded objective function evaluations but is much harder when we want to evaluate the time-consuming objective in parallel. In “Parallel Bayesian Global Optimization of Expensive Functions,” Wang, Clark, Liu, and Frazier provide a more efficient approach to optimizing the parallel version of the most widely used acquisition function for Bayesian optimization, expected improvement. This enables practical large-scale Bayesian optimization. Numerical

experiments demonstrate this new approach scales to at least 128 parallel function evaluations. This new approach was a core component of the widely used open-source Bayesian optimization codebase, the Metrics Optimization Engine (MOE), and was also implemented within Facebook’s Bayesian optimization codebase, BoTorch.

### Why Is Maximum Clique Often Easy in Practice?

“Why Is Maximum Clique Often Easy in Practice?,” by Walteros and Buchanan focuses on providing a rigorous, worst-case explanation for why the maximum clique problem is apparently so easy in naturally occurring graphs, despite its intractability in the worst case. To this day, there exist unsolved benchmark instances with just 1,000 vertices. In contrast, *real-life* instances appear to be much easier. Indeed, relatively simple algorithms can solve instances with millions of vertices in just a few seconds. However, these algorithms lack any worst-case guarantees on their running time. Exactly *why* these real-life instances are so easy has been somewhat of a puzzle. The authors show that maximum clique can be solved in time polynomial in the size of the graph, but exponential in a graph invariant that we denote by  $g$ . Typically,  $g$  is very small for real-life graphs, being 0, 1 or 2 in over half of the instances coming from commonly used testbeds. In these and other cases where  $g$  can be thought of as a small constant, their algorithm runs in time  $O(m^{1.5})$ .

### A Simulation-Based Approach for Calibrating Stochastic Models

Fitting a stochastic model to output data provides a mechanism to calibrate the model without the availability of direct input data and demonstrably improves output prediction accuracy. This idea is analogous to supervised learning built from black-box models but with causal stochastic structure. In “Maximum Likelihood Estimation by Monte Carlo Simulation: Toward Data-Driven Stochastic Modeling,” Peng, Fu, Heidergott, and Lam propose a simulation-based optimization approach that can efficiently fit various forms of causal stochastic models to the output data via maximum likelihood estimation. In a sense, this work brings “light” to the black box.

### Decision Diagrams Fuel Decomposition Techniques for Notoriously Hard Optimization Problems

In “On the Consistent Path Problem,” Lozano, Bergman, and Smith study a novel decomposition scheme, utilizing decision diagrams for modeling elements of a problem where typical linear relaxations fail to provide sufficiently tight bounds. Given a collection of decision diagrams, each representing a portion of the problem, together with linear inequalities modeling other portions of the problem, how can one efficiently optimize over such a representation? The authors model the problem as a consistent path problem, where a path in each diagram has to be identified, all of which agree on the value assignments to variables. They establish complexity results and propose a branch-and-cut framework for solving the decomposition. Through application to binary cubic optimization and a variant of the market split problem, the authors show that the decomposition approach provides significant improvement gains over standard linear models.