Proceedings of the 2016 Winter Simulation Conference
T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, eds.

SIMULATION OPTIMIZATION WITH SENSITIVITY INFORMATION: AN APPLICATION TO ONLINE-RETAIL INVENTORY REPLENISHMENT

Annie I. Chen

Dept. of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
77 Massachusetts Ave
Cambridge, MA 02139, USA

Stephen C. Graves

Sloan School of Management
Massachusetts Institute of Technology
77 Massachusetts Ave
Cambridge, MA 02139, USA

ABSTRACT

We study a simulation optimization approach for the online-retail inventory replenishment problem, where the goal is to minimize the total operational cost among a class of parametrized replenishment policies. We model the problem as an infinite-horizon average-cost dynamic program with discrete states and controls, which can be difficult to solve due to the curse of dimensionality. We propose a simulation optimization method that searches for low-cost policies using cost sensitivity information with respect to policy parameters. This information is analogous to the gradient in continuous-space descent methods, and we show that it can be collected as a by-product of the simulation process. Numerical experiments with realistic inventory networks show that our method is able to identify low-cost solutions efficiently, requiring significantly fewer iterations and less computation time than a simple random search method that only samples the local neighborhood.

1 INTRODUCTION

A typical online-retail inventory system consists of a set of fulfillment centers, spread across different geographic locations, in which items are held in storage and shipped to customers on demand. Inventory at every fulfillment center is reviewed periodically by a centralized decision maker, and replenishment orders are placed concurrently, although their lead times may differ. Our goal is to find a replenishment policy that minimizes the total operational cost.

There are three major components of online-retail operational costs: lost sales, inventory holding, and outbound shipping costs. While lost sales and holding costs have been widely studied, the cost of shipping items to customers can often be difficult to characterize analytically, as it depends on the choice of fulfillment center from which to ship. A common heuristic in practice is the *myopic fulfillment policy*, where demand is satisfied using the fulfillment center with available inventory that can meet the customer's service request at the lowest shipping cost. As a consequence, when the closest (i.e., cheapest) fulfillment center experiences a stockout, demand is routed to the next closest fulfillment center, a situation known as *spillover*. If there is no inventory in any fulfillment center, demand is lost; but as long as there is inventory somewhere in the system, demand will be fulfilled, even if the spillover requires shipping from a distant fulfillment center at a high cost. This is different from traditional brick-and-mortar warehouses, in which a local stockout directly leads to lost sales rather than spillovers.

Inventory replenishment in online retail has been studied recently by (Acimovic and Graves 2016), who also provide a detailed review on background contexts and related literature. Our model differs from (Acimovic and Graves 2016) in three ways: First, we optimize the total operational cost, rather than minimizing the shipping cost subject to a given system safety stock. Secondly, we allow for non-identical deterministic lead times. Thirdly, we account for integrality explicitly in the model, which is an important consideration for low-demand items that often constitute a significant proportion of the catalog offered by online retailers. To the best of our knowledge, our work is the first systematic investigation of simulation optimization applied to online-retail inventory replenishment.

2 DYNAMIC PROGRAM FORMULATION AND APPROXIMATION

We model the online-retail inventory replenishment problem for a single item as an infinite-horizon average-cost dynamic program. The state is the vector of fulfillment center on-hand inventory levels at the review epoch, and the control is the vector of replenishment order quantities (with simplifying assumption that longest lead time is less than the review period.) Under the myopic fulfillment policy, we provide an exact characterization of the transition probabilities and expected stage cost. This allows us to solve for the optimal replenishment policy, using standard approaches such as relative value iteration and policy iteration. However, due to the complexity caused by lead time and spillover, it is computationally infeasible to scale up these exact solution methods to real-world problem sizes.

This leads us to two approximation approaches. First, we approximate in policy space by restricting our attention to certain classes of parametrized policies that are simple to implement in practice. In addition to the widely-used base-stock policy, we also consider alternatives such as a constant-basestock hybrid policy, as well as a projected base-stock policy similar to that proposed by (Acimovic and Graves 2016).

Secondly, we approximate in value space by using simulation to evaluate candidate polices. Moreover, as described below, we also extract information from the simulation process that is useful for improving upon the current candidate policy, thus enabling us to perform optimization with simulation.

3 SIMULATION OPTIMIZATION WITH SENSITIVITY INFORMATION

We design a discrete-event simulation process where, as a natural by-product, we can also collect information on cost sensitivity with respect to changes in input parameters. Specifically, as we simulate the inventory replenishment and demand arrival events, we keep an account of the marginal cost of order quantities, i.e., how much cost we would have saved if we had previously ordered an additional unit of inventory at some fulfillment center. At the end of each simulation run, we take the average marginal cost as an approximate measure of cost sensitivity with respect to the order quantities, which we then translate into sensitivity with respect to the current replenishment order policy parameters, using perturbation analysis techniques.

Much like gradients in descent methods for continuous variables, sensitivity information can be used to guide our search process for the best policy in a discrete parameter space. Under the framework of simulation optimization via random search, such as that described in (Andradóttir 2006), our method can be viewed as a random search method where the search is focused in the direction of negative sensitivity, i.e., where policies with lower costs are more likely to be found.

Numerical experiment results with realistic synthetic fulfillment networks suggest that, with the help of sensitivity information, our proposed method is more efficient than a simple random search method that only searches candidates in the local neighborhood. The improved efficiency is due to the fact that sensitivity information enables a more focused search (fewer candidates to sample per iteration) in a good direction (where the cost is likely to improve) along which we can move more quickly (allowing for jumps beyond the local neighborhood).

REFERENCES

Jason Acimovic and Stephen C. Graves 2016. "Mitigating spillover in online retailing via replenishment". Available at SSRN: http://ssrn.com/abstract=2459097.

Andradóttir, S. 2006. "An overview of simulation optimization via random search". *Handbooks in operations research and management science* 13:617–631.