

A Dynamic Mobile Production Capacity and Inventory Control Problem

S. S. Malladi^{a, b}, Alan L. Erera^a, and Chelsea C. White III ^a

^aGeorgia Institute of Technology, USA; ^bTechnical University of Denmark, Denmark

ARTICLE HISTORY

Compiled November 26, 2018

ABSTRACT

We analyze a problem of dynamic logistics planning given uncertain demands for a multi-location production-inventory system with transportable modular production capacity. In such systems, production modules provide capacity, and can be moved from one location to another to produce stock and satisfy demand. We formulate a dynamic programming model for a planning problem that considers production and inventory decisions, and develop suboptimal lookahead and rollout policies that use value function approximations based on geographic decomposition. Mixed-integer programming formulations are provided for several single-period optimization problems that define these policies. These models generalize a formulation for the single-period newsvendor problem, and in some cases the feasible region polyhedra contain only integer extreme points allowing efficient solution computation. A computational study with stationary demand distributions, which should benefit least from mobile capacity, provides an analysis of the effectiveness of these policies and the value that mobile production capacity provides. For problems with 20 production locations, the best suboptimal policies produce on average 13% savings over fixed capacity allocation policies when the costs of module movement, holding, and backordering are accounted for. Greater savings result when the number of locations increases.

KEYWORDS

mobile modular production, joint inventory control and capacity logistics, approximate dynamic programming, rollout heuristic

1. Introduction

Mobile manufacturing capacity is an emerging innovation in the chemical process industry and in the additive manufacturing industry. E-commerce giant Amazon has recently been granted a

patent for mobile additive manufacturing in a make-to-order setting (Geek Wire, 2018). Bayer, a pharmaceutical and agricultural chemicals company, has developed a containerized production unit that operates an intensified continuous batch process for fertilizer production, and has shown that mobile production units require lower setup costs than fixed facilities (F3, 2014). Pfizer is conducting a large-scale collaborative research project on miniaturized modular production technology for oral drugs (Pfizer, 2015). Novartis has developed a refrigerator-sized on-demand pharmacy that can produce common drugs (Novartis, 2016). Supply chain systems that rely on mobile and modular production capacity may have many benefits, including:

- a smaller total capacity investment due capacity sharing (via production module movement);
- a faster response to time-dependent demands for time-sensitive or perishable products by producing locally;
- an alternative strategy for handling uncertain spatial demand variation by relocating capacity inexpensively; and
- an ability to conduct reconnaissance of new markets for products using recoverable production capacity.

To develop a better understanding of supply chain systems that rely on mobile and modular production capacity, this paper explores a dynamic production-inventory planning problem in this emerging context. Suppose that product demands arise in a number of locations, and that each location can host one or more transportable production units referred to hereafter as *modules*. A fleet of production modules is available and is deployed across the system. The module counts at each location will be referred to as the capacity configuration, and this configuration can be altered by moving modules. We introduce the mobile modular production and inventory problem (MMPIP) which seeks to determine an optimal policy for managing capacity configurations, production, and inventory over a finite planning horizon to serve uncertain demands. In addition to the typical trade-off between inventory and shortage costs, there exist trade-offs with the cost of relocating modules.

We define an initial base problem in this study, and for simplicity assume that demand in one location cannot be satisfied by production in another location either directly or via transshipment. This modeling choice was made to focus squarely on capacity movement. There also may be some supply chain systems where production must be local. We focus on a make-to-stock

inventory system, and note that in one of the important applications in chemical manufacturing that holding inventory is typical and that the time required for production is significantly larger than the time taken for demand fulfillment. Finally, we study cases where location-wise demands are stationary and independent. Such settings should provide the least value for mobile production capacity, but are useful to demonstrate the primary modeling ideas.

Successful operation of a production system with transportable, modular capacity depends on other important considerations that will not be specifically addressed here. For example, effective inbound logistics systems for the inputs to production must be available in each location, and must be rapidly scalable with allocated production capacity. Similarly, outbound distribution systems must also accommodate potentially changing production rates. Efficient movement and rapid setup and breakdown of the mobile production modules is also necessary for such systems to be effective.

We organize this paper into the following sections. Section 2 provides a formal problem definition and a formulation using a Markov decision process (MDP), and Section 3 places the research in the context of related literature. Section 4 presents bounds on the optimal cost function. In Section 5, we propose various heuristics for finding sub-optimal model solutions. Section 6 provides a numerical study of computational experiments using the heuristics on three different sets of instances. We summarize our findings and conclude the paper in Section 7.

2. Problem Description

Consider a production-inventory system facing uncertain demands at a set of L locations, operated over a finite time horizon. Each location can produce and store inventory of product to meet its demand over time, with unmet demand backlogged. At any time, a fleet of homogeneous production modules is available and distributed across the system, with some units installed in place and the others moving between locations. Production capacity at each location is limited by the number of production modules present. The objective is to determine module movement, production, and inventory decisions over time to minimize total system costs.

Consider a decision model with a planning horizon T discretized into $T - 1$ consecutive, equal-duration decision periods, $\{1, \dots, T - 1\}$. Each location i must satisfy or backlog demand $D_i(t)$ during epoch t . Let $Y(t)$ be the total number of production modules available in the fleet at time t , and suppose that each module can produce a maximum of G units of product per time period. Finally, let α_{ij} be the time that a module is unable to produce when moved from location i to j , measured in fractional periods. Then, in each period two types of primary decisions are

made: module movement decisions to relocate capacity, and production decisions to use installed modules at each location. When combined with observed demand, production decisions imply changes to inventory positions at locations.

Our models assume the following sequence of events in each period, as depicted in Figure 1. First, module movement decisions are determined and executed, yielding a new capacity configuration. Next, production decisions are determined and executed, yielding a post-replenishment inventory state for each location. Finally, uncertain demands are observed and the inventory state is updated after filling or backlogging demands. Costs are incurred for module movement, production, inventory holding, and demand backlogging.




Figure 1. Sequence of events within a period

In this paper, we define a simple initial problem that we denote the *mobile modular production and inventory problem* (MMPIP). The MMPIP specifically models systems of the type introduced above with the following additional assumptions:

- Location demands are modeled as discrete random variables, and are independent and stationary across locations and time periods;
- Total available capacity is constant throughout the planning horizon, such that $Y(t) = Y$ for all t ;
- Module movement costs are time-invariant, linear, and separable in the number of modules moved between location pairs;
- Module movements require no movement or setup time such that $\alpha_{ij} = 0$ for all pairs of

locations i and j , and thus modules moved during period t are immediately available for production at a new location in that period;

- Production costs are linear in the number of units produced, and per period inventory holding costs and backordering costs are linear in the number of units; and
- Production decisions executed during period t create new items available for immediate use in period t .

In the next subsection, we present a detailed formulation for the MMPIP.

2.1. Formulation

We formulate the MMPIP as a Markov decision process (MDP) for a finite horizon with T periods. Decisions are made in the first $T - 1$ epochs, $\mathcal{T} = \{1, \dots, T - 1\}$, and epoch T models the end state. At every decision epoch $t \in \mathcal{T}$, the state vector is comprised of $2L$ components: the number of modules, $u_i(t)$, before modules are moved and the inventory position, $s_i(t)$, before production at each location i . The action vector at t is comprised of L^2 components: the number of modules to move, $\Delta_{ij}^M(t)$, from location i to location $j \neq i$ and the production quantity, $q_i(t)$, at each location i . Let $D_i(t)$ be the discrete demand random variable for location i in period t , where stationarity implies that the probability mass function for $D_i(t)$ is given by P_i for all time periods t .

Module movement decisions outbound from location i in period t are limited by its inventory of modules, $u_i(t)$. The production decision at each location i in period t is limited by the capacity provided by the post-movement module state, $u_i(t + 1)$. Note that this coupling of production decisions to module movement decisions defines the extension that this model proposes to single-location inventory control models.

Costs are incurred for actions as follows. Module movement costs are linear in the number of modules moved between locations i and j , and given by $K_{ij}^M \Delta_{ij}^M(t)$. Production and inventory costs are separable by location i , and follow the usual form found in base stock models. Production costs are linear, and given by $c_i q_i(t)$. Holding costs accrue for each unit of positive inventory position at the end of period t , $h_i(s_i(t) + q_i(t) - D_i(t))^+$ where $(y)^+ \equiv \max(y, 0)$. Similarly, backorder costs accrue for each unit of negative inventory position, and are given by $b_i(D_i(t) - s_i(t) - q_i(t))^+$. We assume no costs associated with any state in the final horizon period T .

If we let $\xi(t) = (\{u_i(t)\}, \{s_i(t)\})$ represent the complete state variable tuple, we can formulate the MDP as follows:

$$\begin{aligned}
V_t(\xi(t)) = & \min_{\substack{\Delta_{ij}^M(t): \\ \sum_j \Delta_{ij}^M(t) \leq u_i(t)}} \min_{\substack{\forall i q_i(t): \\ q_i(t) \leq Gu_i(t+1)}} \mathbb{E}_D \left[\sum_i \left\{ \sum_j K_{ij}^M \Delta_{ij}^M(t) \right. \right. \\
& + c_i q_i(t) + h_i(s_i(t) + q_i(t) - D_i(t))^+ \\
& \left. \left. + b_i(D_i(t) - s_i(t) - q_i(t))^+ + V_{t+1}(\xi(t+1)) \right\} \right], \quad \forall \xi(t), \forall t \in \mathcal{T} \quad (1)
\end{aligned}$$

$$\begin{aligned}
\text{where } \xi(t) = & (s_1(t), s_2(t), \dots, s_L(t), u_1(t), u_2(t), \dots, u_L(t)), \quad \forall t \in \mathcal{T} \\
s_i(t+1) = & s_i(t) + q_i(t) - D_i(t), \quad \forall i \in \{1, \dots, L\}, \quad \forall t \in \mathcal{T} \\
u_i(t+1) = & u_i(t) - \sum_j \Delta_{ij}^M(t) + \sum_k \Delta_{ki}^M(t), \quad \forall i \in \{1, \dots, L\}, \quad \forall t \in \mathcal{T} \\
\sum_i u_i(1) = & Y
\end{aligned}$$

$$V_T(\xi(T)) = 0 \quad \forall \xi(T). \quad (2)$$

where $V_t(\xi(t))$ is the expected cost-to-go function of MMPIP from decision epoch t to the end of the horizon. Extending this formulation and the solution approaches presented in this paper to include a per-period discount rate is straightforward.

This MDP formulation uses a state space whose cardinality is exponential in the number of locations L . The space of possible movement decisions each period is similarly large, and thus finding exact optimal solutions to this model will not be possible except for the smallest instances. We therefore will develop heuristic solution methods for identifying high-quality suboptimal designs.

3. Related Literature

The emergence of reconfigurable, mobile, decentralized/distributed manufacturing units has generated significant interest in the manufacturing/process industry in recent years (Bi et al., 2008; Clausen et al., 2015; Kessler and Bruell, 2015; Koren, 2010; Lier and Grünewald, 2011; Lier et al., 2013,1; Rogers and Bottaci, 1997; Wörsdörfer and Lier, 2017; Wörsdörfer et al., 2017). Reconfigurable or mobile modular production systems are characterized by *transformability*: scalability, adaptability (modularity, universality, compatibility), and mobility (Lier and Grünewald, 2011; Wörsdörfer et al., 2017). Marcotte et al. and Marcotte and Montreuil (2016) present mathematical models for make-to-stock and make-to-order scenarios of *hyperconnected* mobile production.

The MMPIP model considers both dynamic capacity allocation and also joint inventory management across multiple locations. When demands are deterministic, dynamic capacity allocation can be viewed as a special case of the dynamic facility location problem (DFLP). The DFLP is the problem of determining locations and opening schedules for multiple facilities

(equivalently, units of capacity) over the planning horizon (Jena et al., 2015). Wörsdörfer and Lier (2017) model a special case of the MMPIP, a mobile modular production-inventory problem with deterministic demands, as a DFLP. A DFLP with modular capacities is presented in (Jena et al., 2015), which provides a good review of DFLP literature. Shifting a module in the MMPIP is equivalent to opening a facility at the new location and closing one at the old location in the DFLP. Ghiani et al. (2002) present a capacitated DFLP with multiple facilities in the same site for which Melo et al. (2005) allow transfer of capacity between sites. The mobile facility routing problem with deterministic demands (Halper and Raghavan, 2011) is a DFLP with mobile facilities with duration-based capacity.

The feature of managing mobile capacity and controlling inventory under uncertain demands over all the locations is missing in the DFLP literature. In order to manage the mobile modular production-inventory system effectively, it is necessary to consider inventory and capacity management simultaneously. Hence, the MMPIP cannot be treated as a special case of existing DFLP models.

In the context of joint capacity and inventory decision-making, Angelus and Porteus (2002) study a make-to-stock production system with the ability to buy and sell capacity and prove that a target interval policy is optimal in two cases (with and without carryover of inventory). Simultaneously planning inventory actions and capacity change decisions at a single location is also studied in (Bradley and Glynn, 2002; Rajagopalan and Swaminathan, 2001). Qiu and Sharkey (2013) jointly plan the location and inventory of a single facility facing spatially distributed deterministic demand, in the context of managing a collection of ships that serves as a military base at sea. Iyer and Jain (2004) analyze priority-based operating rules for a production inventory system in which two locations with low and high demand variabilities choose to pool their capacities.

We represent the MMPIP using a Markov decision process (MDP) model which for realistically-sized problems is computationally intractable. Hence, our focus is on finding good heuristics that rely on techniques that such as value function approximations (Powell, 2007, Ch. 10) and rollout methods (Goodson et al., 2017). For a rollout heuristic to be effective, the computation of the guidance mechanism must be tractable (Bertsekas et al., 1997; Ryzhov et al., 2012; Secomandi, 2001; Tesauro and Galperin, 1996). In the next sections, we will develop tractable heuristics for the model using these techniques.

4. Bounds on the Optimal Expected Cost Function

In this section, we develop two lower bounds for the optimal expected cost function of the MMPIP, specified by (1) and (2). We denote them as the *perfect information relaxation* (PIR) and the *most flexible system* (LB) lower bounding functions respectively. Related to the most flexible system bound, we also develop an upper bound based on a fixed module configuration.

4.1. Perfect Information Relaxation: Lower Bound (PIR)

The structure of the recursion in (1) ensures that decisions made in time period t do not anticipate the outcomes of $D_i(\tau)$ for $\tau \geq t$. A common approach for developing a lower bound for $V_t(\xi(t))$ is to relax this condition, and to solve a deterministic planning problem for each possible demand trajectory defined by outcomes of $(\{D_i(t)\}, \{D_i(t+1)\}, \dots, \{D_i(T-1)\})$. A lower bound on the optimal expected cost is then given by the average of the resultant optimal objective functions for the demand trajectories considered. The deterministic planning problems are solved by value iteration for each trajectory, solving the deterministic case of (1) and (2) with the expectation removed.

Since the number of demand trajectories may be very large for longer horizon problems with many production locations, this PIR bound can be approximated by computing an average over a Monte Carlo sample of trajectories. We note that bounds of this type (sometimes also called *a posteriori* bounds) are quite common for dynamic programming models, but we have found that they are weak for MMPIP problem instances. We note that PIR bounds for problems in which capacity is fixed at locations in advance also tend to be weak. The weakness of PIR bounds for the MMPIP does lead to a natural conjecture: there may be significant value to move production modules in response to specific demand trajectories, and therefore solving MMPIP problems effectively may lead to significant value beyond systems with fixed installed capacity. Through our computational analysis, we will find later that it is indeed the case.

4.2. Most Flexible System: Lower Bound (LB)

We now present a lower bound for the MMPIP that we have found to be much tighter than the PIR in computational experiments. A mobile modular production system is highly flexible, since the capacity configuration can be altered during the time horizon. Module movement costs, however, mitigate the value of this flexibility. An approach to developing lower bounds, which we now take, is to assume that production capacity can be allocated to locations immediately in any period, and without module movement costs, in response to the current vector inventory state.

Since lower bounds of this type do not depend on the current capacity configuration, we specify a lower bounding function $\tilde{V}_t^{\mathcal{L}}(\{s_i(t)\})$, where

$$\tilde{V}_t^{\mathcal{L}}(\{s_i(t)\}) \leq V_t(\{s_i(t)\}, \{u_i(t)\}) \quad \forall \{u_i(t)\}.$$

We compute $\tilde{V}_t^{\mathcal{L}}$ recursively using a simpler dynamic program, where we assume that the lowest-cost capacity configuration can be used for each possible inventory state, as follows:

$$\begin{aligned} \tilde{V}_t^{\mathcal{L}}(\{s_i(t)\}) &= \min_{\substack{\{u_i(t+1)\} \\ \sum_i u_i(t+1)=Y}} \min_{q_i(t) \leq Gu_i(t+1) \forall i} \mathbb{E}_D \sum_{i \in \mathcal{I}} \left[h_i(s_i(t) + q_i(t) - D_i(t))^+ \right. \\ &\quad \left. + b_i(D_i(t) - s_i(t) - q_i(t))^+ + \tilde{V}_{t+1}^{\mathcal{L}}(\{s_i(t) + q_i(t) - D_i(t)\}) \right], \quad \forall t \in \mathcal{T}, \\ \tilde{V}_T^{\mathcal{L}}(\{s_i(T)\}) &= 0 \forall \{s_i(T)\}. \end{aligned}$$

Note that it remains computationally expensive to compute the bounding function $\tilde{V}_t^{\mathcal{L}}$ in general. It is more difficult than solving L independent inventory management problems. Since the total number of modules defines total production capacity, it is necessary in most instances to decide which locations, given a current inventory state, should be prevented from selecting an optimal unconstrained production value by restricting their capacities. However, it requires significantly less computation than solving the MMPIP problem to optimality. Consider solving both problems by value iteration. The number of capacity configurations possible is $\mathcal{O}((Y+L)^{L-1})$, since the number of ways to allocate Y modules to L locations is given by $\binom{Y+L-1}{L-1}$. Thus, at each epoch, for a given vector inventory state, value iteration for MMPIP requires $\mathcal{O}((Y+L)^{L-1})$ times more effort than solving for LB.

4.3. Fixed Capacity System: Upper Bound (UB)

A conventional production system with stationary capacity at all locations can be viewed as a mobile modular production system where the capacity configuration is fixed for the planning horizon. Given a fixed capacity configuration, an upper bound on the optimal expected cost given an initial inventory state can be determined by solving L independent, constrained inventory management problems. More specifically, let $\tilde{V}_t^F(\{s_i(t)\}, \{u_i\})$ be the optimal cost-to-go function of the multilocation fixed system with capacity configuration $\{u_i\}$ fixed from t until the end of the planning horizon, given initial inventory position state $\{s_i(t)\}$. Let $V_{i,t}^F(s, C)$ be the optimal cost-to-go function of the single location inventory control problem with production capacity C and initial inventory position s at location i at time epoch t . When there are u_i modules at i ,

then $C = Gu_i$. To determine this upper bound, we use the following optimality equations:

$$\begin{aligned}
\tilde{V}_t^F(\{s_i(t)\}, \{u_i\}) &= \sum_{i=1}^L V_{i,t}^F(s_i(t), Gu_i) \\
&= \sum_{i=1}^L \min_{q_i(t) \leq Gu_i} \mathbb{E}_D \left[h_i(s_i(t) + q_i(t) - D_i(t))^+ \right. \\
&\quad \left. + b_i(D_i(t) - s_i(t) - q_i(t))^+ + V_{i,t}^F((s_i(t) + q_i(t) - D_i(t), Gu_i) \right] \\
&\quad \forall \{s_i(t)\}, \forall t \in \mathcal{T} \\
\tilde{V}_T^F(\{s_i(T)\}, \{u_i\}) &= 0 \quad \forall \{s_i(T)\}.
\end{aligned}$$

We note that the computation of $\tilde{V}_t^F(\cdot, \{u_i\})$ can be decoupled across locations. Let \mathcal{Q} be the cardinality of the inventory position state space for each location. Given capacity configuration $\{u_i\}$, for each location at each epoch computing the cost function above requires $\mathcal{O}(\mathcal{Q}Gu_i)$ steps, and thus the total computational effort at each epoch is $\mathcal{O}(\mathcal{Q}GY)$. Again, this effort is significantly smaller than the corresponding $\mathcal{O}(\mathcal{Q}^L(Y+L)^{L-1}GY)$ effort required for computing the optimal value function for MMPIP at each epoch for a given capacity configuration.

If we use this upper bounding approach beginning at the initial time epoch 1, we can also compute the minimum (UB_{\min}) and the maximum (UB_{\max}) possible expected total cost of a fixed system, given an initial inventory state of zero at all locations. Define the capacity configuration that corresponds to UB_{\min} as $u_{\min} = \arg \min_{\{u_i\}} \tilde{V}_1^F(\bar{0}, \{u_i\})$. Determining UB_{\min} and u_{\min} requires comparing $\mathcal{O}((Y+L)^{L-1})$ capacity configurations, or solving a multiple choice knapsack problem (Ibaraki et al., 1978; Sinha and Zoltners, 1979) as presented below:

$$\begin{aligned}
&\min \sum_{i=1}^L \sum_{y=0}^Y V_{i,1}^F(0, Gy) z_{iy} \\
&\quad \text{s.t.} \quad \sum_{i=1}^L \sum_{y=0}^Y y z_{iy} = Y \\
&\quad \sum_{y=0}^Y z_{iy} = 1 \quad \forall i \in \{1, \dots, L\} \\
&\quad z_{iy} \in \{0, 1\} \quad \forall i \in \{1, \dots, L\}, y \in \{0, \dots, Y\}
\end{aligned}$$

The idea of the knapsack formulation is to add one item for each location, where the size of the item is its number of assigned modules and the knapsack size is Y . The cost of adding item iy to the knapsack is the expected cost-to-go at epoch 1 for location i with y modules. In the multiple choice problem, a set of constraints ensures that exactly one item is chosen from each class of items (location). We note that the computational effort to solve this integer program in

practice is much smaller than that required to determine the cost-to-go values for all locations and capacity levels.

We will see later fixed system configurations beginning at some time period can play the role of a base heuristic within rollout approaches for determining good (but suboptimal) dynamic policies for the MMPIP problem. Additionally, fixed system bounds will be used as benchmarks to assess the value addition created by mobile modular production systems and to evaluate the performance of heuristics for the MMPIP.

5. Heuristics

Since the MMPIP is characterized by a state space whose size is exponential in the number of locations, L , and the length of the horizon, T , the model suffers from a curse of dimensionality. Hence, we seek suboptimal policies for the problem constructed using approximate dynamic programming techniques, such as rollout algorithms and decomposition-based approaches, that do not require complete characterization of the optimal expected cost function. Following the terminology presented in Goodson et al. (2017), a one-step rollout algorithm is a value function approximation approach in which the decision at the current epoch is determined by approximating the cost-to-go by the expected cost of implementing a base policy for states beginning in the next epoch. We propose a decomposition-based rollout algorithm, RF, that approximates the cost-to-go function by assuming that the capacity configuration does not change again after decisions made in current epoch, and that an optimal inventory control policy is used for future replenishment decisions given this fixed capacity. Since this rollout can still be computationally expensive, we also present an alternative one-step lookahead policy, LAF, that approximates the expected cost of the optimal RF rollout. A one-step lookahead policy is a suboptimal policy obtained by minimizing the sum of the immediate cost for the current period and an approximation of the cost-to-go function for the remaining horizon (Bertsekas, 2000, Chapter 6). In addition to these two core heuristics, we also develop additional value function approximation policies that can be used for special case problem instances with $L = 2$ locations.

5.1. Myopic Policy (MP)

In a myopic policy, we ignore the cost-to-go in the next epoch, and thus $V_{t+1}(\{s_i(t+1)\}, \{u_i(t+1)\}) \approx 0$ for all system states. The myopic action in the current epoch can be found by solving an integer linear program (IP) that extends the classic formulation for the discrete demand distribution newsvendor model. Let the set of demand outcomes from stationary distribution P_i at every location i be $\{d_i^k\}$, where for notational convenience k indexes the outcomes in \mathcal{K}_i . Given

the current state $(\{s_i\}, \{u_i\})$ of module allocation and inventory positions, the IP formulation (3) is given by:

$$\begin{aligned}
\text{MMPIP-MP. } \min \sum_{i=1}^L & \left[\sum_{j=1}^L K_{ij}^M \Delta_{ij}^M + \sum_{k=1}^M p_i^k (h_i r_i^k + b_i o_i^k) \right] \\
& r_i^k \geq s_i + q_i - d_i^k, \quad \forall k \in \mathcal{K}_i, i \in \{1, \dots, L\} \\
& o_i^k \geq d_i^k - s_i - q_i, \quad \forall k \in \mathcal{K}_i, i \in \{1, \dots, L\} \\
& 0 \leq q_i \leq G \left(u_i - \sum_j \Delta_{ij}^M + \sum_l \Delta_{li}^M \right) \quad \forall i \in \{1, \dots, L\} \\
& \sum_j \Delta_{ij}^M \leq u_i \quad \forall i \in \{1, \dots, L\} \\
& q_i, \Delta_{ij}^M \in \mathcal{Z}^+ \quad \forall i, j; \quad r_i^k, o_i^k \in \mathcal{Z}^+ \quad \forall k, \forall i.
\end{aligned} \tag{3}$$

The decision variables are the module movements $\{\Delta_{ij}^M\}$, the replenishment quantities $\{q_i\}$, and the positive and negative parts of post-decision, post-information inventory position $\{r_i^k\}$ and $\{o_i^k\}$ respectively. MMPIP-MP minimizes the immediate cost of module movement and inventory holding or backordering. The first two constraints characterize a newsvendor problem. The third constraint ensures that the post-module movement production capacity is not exceeded at each location. The last constraint prevents removing more modules than those available at any location.

The formulation for the single location discrete demand distribution newsvendor model is proved to be a linear program (Chen et al., 2014). We now present analogous structural results on the integer program MMPIP-MP to enable increased computational efficiency.

Theorem 5.1. *For integer values of G , d_i^k , u_i , and s_i for all $k \in \mathcal{K}_i$ and $i \in \{1, \dots, L\}$,*

- (a) *the integrality constraints on $\{o_i^k\}$, $\{r_i^k\}$, and $\{q_i\}$ for all $k \in \mathcal{K}_i$, $i \in \{1, \dots, L\}$ in the integer program MMPIP-MP are redundant.*
- (b) *when capacity per module $G = 1$, the linear programming relaxation of the integer program MMPIP-MP has an integral optimal solution.*

Proof of Theorem 5.1 is provided in Appendix A.1. Theorem 5.1(a) implies that it is sufficient to impose integrality constraints on the variables $\{\Delta_{ij}^M\}$ only. Theorem 5.1(b) presents a condition, namely $G = 1$, under which, all the integrality constraints are redundant and thus the MMPIP-MP can be solved by its linear programming relaxation. This implies fast compute

times when $G = 1$ even for relatively large values of L and T .

5.2. Rollout of Fixed Future (RF)

In this rollout heuristic, the base heuristic assumes that beginning in the next period modules will be fixed at their current locations until the end of the horizon. Thus, this approach approximates the flexible capacity production system with one in which flexibility is only available for the current period. We propose integer linear program (4) to determine the optimal one-step decisions for this rollout, given the current state $(\{s_i\}, \{u_i\})$ of module allocation and inventory positions. This problem is an extension of the integer program for the myopic single-period action selection problem.

In (4), the objective is to minimize movement cost, expected inventory cost at all locations, and the expected optimal fixed future cost. Note that since the capacity state is fixed beginning in the next period, it is possible to decompose the optimal cost-to-go function by location, and to only require the local production capacity and inventory state as inputs to the precomputed functions $V_{i,t+1}^F(s, C)$. Thus, in addition to the $\{q_i\}$ and $\{\{\Delta_{ij}^M\}\}$ decision variables used in the myopic integer program, binary variables $z_i(\Delta^M, q)$ are specified that take value 1 if Δ^M modules are transferred to location i , and then used in the current period to produce q items (note that $q \leq G(u_i + \Delta^M)$ and that Δ^M may be negative). The first two constraints again are used to compute single-period underage or overage units. The next three constraints ensure that only one set of module movement and production decisions is made for each location, and that the module movement variables result in the selected capacity state for each location i . The formulation is provided here:

$$\begin{aligned}
\text{MMPIP-RF. } \min & \sum_{i=1}^L \left[\sum_{j=1}^L K_{ij}^M \Delta_{ij}^M + \sum_{k=1}^M p_i^k \{h_i r_i^k + b_i o_i^k\} \right. \\
& + \left. \sum_{i=1}^L \sum_{\Delta^M=-u_i}^{Y-u_i} \sum_{x=0}^{G(u_i+\Delta^M)} z_i(\Delta^M, q) \sum_k p_i^k V_{i,t+1}^F \left(s_i + q - d_i^k, G(u_i + \Delta^M) \right) \right] \\
& r_i^k \geq s_i + q_i - d_i^k, \quad \forall k \in \mathcal{K}_i, i \in \{1, \dots, L\} \\
& o_i^k \geq d_i^k - s_i - q_i, \quad \forall k \in \mathcal{K}_i, i \in \{1, \dots, L\} \\
& \sum_{\Delta^M=-u_i}^{Y-u_i} \sum_{q=0}^{G(u_i+\Delta^M)} z_i(\Delta^M, q) = 1 \quad \forall i \in \{1, \dots, L\} \\
& - \sum_j \Delta_{ij}^M + \sum_l \Delta_{li}^M = \sum_{\Delta^M=-u_i}^{Y-u_i} \sum_{q=0}^{G(u_i+\Delta^M)} \Delta^M z_i(\Delta^M, q) \quad \forall i \in \{1, \dots, L\} \\
& q_i = \sum_{\Delta^M=-u_i}^{Y-u_i} \sum_{q=0}^{G(u_i+\Delta^M)} q z_i(\Delta^M, q) \quad \forall i \in \{1, \dots, L\} \\
& z_i(\Delta^M, q) \in \{0, 1\}, \quad \forall q \in \{0, \dots, G(u_i + \Delta^M)\}, \\
& \Delta^M \in \{-u_i, \dots, Y - u_i\}, \quad i \in \{1, \dots, L\} \\
& q_i, \Delta_{ij}^M \in \mathcal{Z}^+, \quad \forall i, j \in \{1, \dots, L\}; \quad r_i^k, o_i^k \in \mathcal{Z}^+ \quad \forall k \in \mathcal{K}_i, \forall i \in \{1, \dots, L\}. \tag{4}
\end{aligned}$$

It should be clear that the integer program (4) includes a large number of binary variables for larger values of L , Y , and G ; the variable count grows at $\mathcal{O}(GY^2L)$. Furthermore, before the formulation can be used, it is necessary to compute the function lookup tables $V_{i,t+1}^F(s, C)$ for each possible module change $\Delta_i^M \in \{-u_i, \dots, Y - u_i\}$ and inventory state s at all locations using the approach described earlier. We also note that (4) is a potentially useful model when $V_{i,t+1}^F(s, C)$ is an approximate value function, decoupled by location, developed using any alternative approach and not necessarily limited to the case where these functions represent the optimal cost-to-go of the single location fixed capacity problem.

5.3. Lookahead with Approximate Fixed Future (LAF)

The complete rollout heuristic RF can be computationally expensive since the integer program becomes difficult in practice for larger problems. We therefore now develop an approximation of RF that is more computationally tractable. To do so, we approximate the single location cost function by the average of two piecewise linear and convex functions. Doing so bypasses the computationally expensive cost function lookup table modeling required in the RF heuristic. In this method, at every epoch decisions are made under the assumption that the cost-to-go

function is approximated by the following expression:

$$V_{t+1}(\{s_i(t+1)\}, \{u_i(t+1)\}) \approx (\widetilde{V}_{t+1}^F(\{\bar{s}_i(t+1)\}, \{u_i(t)\}) + \widetilde{V}_{t+1}^F(\{s_i(t)\}, \{u_i(t+1)\}))/2,$$

where $\bar{s}_i(t+1) = s_i(t) + q_i(t) - \lceil E[D_i(t)] \rceil$ and $\lceil a \rceil$ rounds a to the nearest integer.

We can model this cost approximation using the optimal cost-to-go function of the fixed system by leveraging the structural properties which establish that cost-to-go function of a capacitated single location inventory control problem is convex in inventory position for a fixed capacity level and convex in capacity level for a fixed inventory position (Aviv and Federgruen, 1997). This result implies that $V_{i,t}^F(s_i, G_i)$ is piecewise linear (due to discrete inventory state space) and convex in s_i for a fixed G_i . Thus, it can be represented as $\max\{\gamma_j^i s_i + \widehat{\gamma}_j^i : (\gamma_j^i, \widehat{\gamma}_j^i) \in \Gamma_t^i(u_i)\}$, $\forall i \in \{1, \dots, L\}$. Similarly, for a fixed s_i , the function $V_{i,t}^F(s_i, G_i)$ is piecewise linear and convex in u_i (since $G_i = Gu_i$) that can be expressed as $\max\{\theta_j^i u_i + \widehat{\theta}_j^i : (\theta_j^i, \widehat{\theta}_j^i) \in \Theta_t^i(s_i)\}$, $\forall i \in \{1, \dots, L\}$.

This approximation is again implemented by modifying the integer program (3), as follows:

$$\begin{aligned} \text{MMPIP-LAF.} \quad & \min \sum_{i=1}^L \left[\sum_{j=1}^L K_{ij}^M \Delta_{ij}^M + \sum_{k=1}^M p_i^k \{h_i r_i^k + b_i o_i^k\} + (\zeta_i + \eta_i)/2 \right] \\ & \zeta_i \geq \gamma_j^i (s_i + q_i - \lceil E[D_i(t)] \rceil) + \widehat{\gamma}_j^i, \quad \forall (\gamma_j^i, \widehat{\gamma}_j^i) \in \Gamma_{t+1}^i(u_i), \quad i \in \{1, \dots, L\} \\ & \eta_i \geq \theta_j^i y_i + \widehat{\theta}_j^i, \quad \forall (\theta_j^i, \widehat{\theta}_j^i) \in \Theta_{t+1}^i(s_i), \quad i \in \{1, \dots, L\} \\ & r_i^k \geq s_i + q_i - d_i^k, \quad \forall k \in \mathcal{K}_i, \quad i \in \{1, \dots, L\} \\ & o_i^k \geq d_i^k - s_i - q_i, \quad \forall k \in \mathcal{K}_i, \quad i \in \{1, \dots, L\} \\ & 0 \leq q_i \leq Gy_i \quad \forall i \in \{1, \dots, L\} \\ & u_i - \sum_j \Delta_{ij}^M + \sum_l \Delta_{li}^M = y_i \quad \forall i \in \{1, \dots, L\} \\ & q_i, y_i, \Delta_{ij}^M \in \mathcal{Z}^+, \quad \forall i, j; r_i^k, o_i^k \in \mathcal{Z}^+ \quad \forall k \in \mathcal{K}_i, \quad \forall i \in \{1, \dots, L\}; \\ & \eta_i, \zeta_i \in \mathbb{R} \quad \forall i \in \{1, \dots, L\}. \end{aligned} \tag{5}$$

In addition to the decision variables described in the implementation of MP, ζ_i and η_i represent the single location future costs at i expressed as a function of next period's inventory when capacity is held at the initial level of Gu_i , and as a function of new capacity when inventory is held at the initial level of s_i respectively. The number of integer variables required to represent the future cost-to-go is significantly lower by $\mathcal{O}(GY^2L)$ in MIP (5) when compared with IP (4).

This reduces the computational effort required to solve this MIP dramatically. Furthermore, when module capacity $G = 1$, this MIP reduces to a linear program.

Theorem 5.2. *For integer values of G , d_i^k , u_i , and s_i for all $k \in \mathcal{K}_i$ and $i \in \{1, \dots, L\}$, when capacity per module $G = 1$, the linear programming relaxation of the mixed integer program MMPIP-LAF has an optimal solution where decision variables q , y , w , r , and o are integral.*

Proof of Theorem 5.2 is presented in Appendix A.2.

5.4. $L = 2$ Heuristics

The suboptimal policies proposed in the prior sections rely on approximating the future value function with a form that decomposes by location. In this section, we explore suboptimal policies and solution heuristics that no longer have this feature to understand how much incremental value may be gained. For computational tractability, we focus on small problem instances with $L = 2$ production locations.

5.4.1. Lookahead with Fixed or Purchasable Most Flexible Future (LFP)

The cost function computation of the lower bound LB presented in Section 4.2 is coupled across locations and has severe computational drawbacks for larger values of L . For fewer locations, however, it may serve as a useful basis for a heuristic. Since the lower bound assumes that modules can be moved without cost, we call this case the most flexible future. In this section, we investigate a value function approximation that blends this lower bound with the fixed future upper bound a suboptimal lookahead policy.

We denote the blending heuristic LFP, denoting a lookahead with a fixed future capacity or a “purchasable” most flexible future capacity. The LFP assumes that the decision-maker chooses the best action at the current decision epoch by assuming that she has a choice between (i) keeping capacity fixed for the remainder of the planning horizon and (ii) making a one-time payment now for an unlimited number of future module movements. That is, the cost-to-go function in Eq. 1 is approximated as follows:

$$V_{t+1}(\{s_i(t+1)\}, \{u_i(t+1)\}) \approx \min \left\{ \tilde{V}_{t+1}^F(\{s_i(t+1)\}, \{u_i(t+1)\}), \kappa^F(t) + \tilde{V}_{t+1}^{\mathcal{L}}(\{s_i(t+1)\}) \right\},$$

where \tilde{V}_{t+1}^F as usual can be decomposed into $\sum_{i=1}^L V_{i,t+1}^F(s_i(t+1), u_i(t+1))$.

To use this approach, it is necessary to define a mobility purchase cost $\kappa^F(t)$. We use the following definition, which works well in practice:

$$\kappa^F(t) = \frac{1}{2} \bar{K}^M (T - t) P(\{D_1 = d_1^{\max}\} \cup \{D_2 = d_2^{\max}\}), \quad \forall t \in \mathcal{T}.$$

This value is an approximation of the expected cost of preventing a single period stockout by moving modules; a decision maker should be willing to pay at least this cost in return for an unlimited number of module movements in future periods. Let \overline{K}^M be the average (directed) cost of moving a module between the locations. The likelihood of a stockout is approximated by the probability of the maximum demand occurring at either location. Finally, the expression assumes that roughly half of the time a potential stock-out will be addressed by the movement of a production modules. Again, it is possible to use different cost expressions here but this value created good results in our computational study.

We note that if the mobility purchase cost $\kappa^F(t)$ is very large, the LFP policy is equivalent to **RF**. On the other hand, if mobility purchase cost $\kappa^F(t) = 0 \forall t \in \mathcal{T}$, we denote the resulting suboptimal policy as *lookahead with most flexible future* (**RLB**). Under the RLB policy, decisions are made each period assuming that in the next period there will be no additional cost of module movements.

5.4.2. Lookahead with Iteratively Updated Cost-to-go (LIU)

The policy LIU is defined by a simulation-based optimization method. The method seeks an approximate characterization of the expected cost function in the form of a lookup table covering the complete discrete state space of MMP. Once again, since such a method becomes intractable for larger values of L , we develop and test this heuristic policy only for systems with $L = 2$.

In each pass of the algorithm presented in Appendix B of the Appendix, the initial estimate of the cost-to-go function is set to the cost-to-go function of the fixed capacity system. Demand outcomes are generated for all periods from the given stationary demand distributions by Monte-Carlo simulation. The state at $t = 1$, $\xi(1)$, has zero inventory at all locations and the best fixed module configuration u^{\min} . At every epoch t at the current state in the sample trajectory, a new iterate is created by approximating the cost-to-go of the current state. The new iterate is determined by blending the current approximation of the current state cost-to-go with a new estimate created by finding the expected cost of an optimal single-period action when using the current cost-to-go approximation at epoch $t + 1$. This approximation approach is repeated and cost-to-go functions updated along N sample trajectories. The heuristic policy induced by the final estimate of the cost-to-go function is referred to as LIU.

Let α_n be the blending coefficient used in the n^{th} iteration of the approach (for trajectory n). The condition $\sum_{n=1}^{\infty} \alpha_n \rightarrow \infty$ prevents premature stalling of the algorithm (Powell, 2007,

Ch. 11). Additionally, the condition $\sum_{n=1}^{\infty} \alpha_n^2 < \infty$ ensures fast convergence of the iterates by limiting their variance. We test three blending coefficients, namely, a constant (0.5), $1/(n+1)$, and $1/(\text{number of visits to a particular state})$, where n is the trajectory counter. We note that although the constant coefficient violates the second condition, it performs better in practice than the other two candidates. Thus, we present computational results only for the constant blending coefficient.

6. Computational Study

We now present a computational study of these ideas. The study is designed both to assess the quality of the suboptimal policies developed in this paper for the MMPIP and also to build an initial understanding of the value of production capacity mobility in supply chain systems. This study focuses only on instances where location demand processes are characterized by stationary parameters, and where demand outcomes are independent across locations.

The study uses three sets of instances:


- (1) **Set 1** is a set of 1280 instances with $L = 2$ locations with sufficiently small values of other parameters that allow the computation of the exact optimal solution cost (OPT) by value iteration.
- (2) **Set 2** is a set of 720 $L = 2$ instances with larger values of the module fleet size Y , planning horizon T , and capacity per module G than Set 1.
- (3) **Set 3** is a set of 540 instances with $L \geq 2$.

The detailed designs of these three sets of instances are presented in Appendix C of the Appendix. To evaluate the suboptimal policies determined by our heuristics, we compute the average total cost of implementing each policy on 50 random demand sample paths. Sample average cost is compared to the optimal total expected cost for each instance in Set 1. For Sets 2 and 3, sample average cost is compared to the expected optimal total cost of the best fixed capacity system configuration given by UB_{\min} .

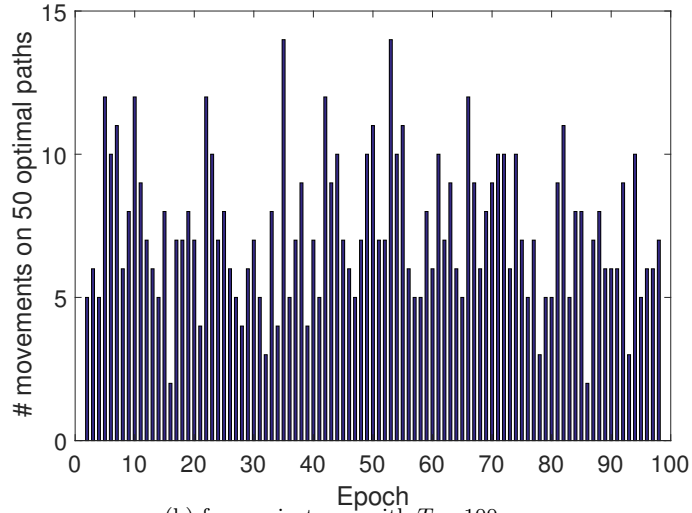
6.1. Results for Set 1

We first determine optimal solutions for each of the instances in Set 1, all of which have only two production locations and short planning horizons of no greater than 15 decision epochs. We note that sometimes optimal policies for stationary policies are unlikely to prescribe any movement of production modules during the horizon. However, in roughly half of the instances (595 of the 1280), the optimal policy leads to at least one module movement in at least one of the 50 sample

trajectories.



(a) total for all $T = 10$ and $T = 15$ instances



(b) for one instance with $T = 100$

Figure 2. Module movements induced by optimal policies for Instance Set 1

Figure 2a presents histograms that depict the total number of module movements induced by optimal policies for all instances and trajectories considered with planning horizon lengths $T = 10$ and $T = 15$ during specific decision epochs. Note that movements are observed in both early and late decision epochs, but that end effects tend to decrease the number of movements at both the beginning and the end of the planning horizon. This is especially true in the initial decision epoch when modules are already positioned optimally for a fixed capacity system. We additionally solved a single instance with a longer planning horizon of $T = 100$ decision epochs, and depict the sum of module movements per epoch over the 50 sample trajectories in Figure 2b. We note that the figure demonstrates that performance metrics for problems with longer planning horizons are less impacted by end effects (initial capacity configuration and inventory levels, and zero value of end-of-horizon inventory); module movements are somewhat evenly dispersed across the epochs in the decision horizon.

We label each instance of Set 1 as (AB, CD) where $A, B \in \{L, H\}$ respectively indicate

the levels of expected demand at locations 1 and 2 and $C, D \in \{L, H\}$ indicate the levels of coefficient of variation (CV) at locations 1 and 2 respectively. For either location, the level of expected demand is considered high (H) if it is greater than or equal to $0.3GY$ and low (L) otherwise (recall that total system production capacity is GY). Likewise, coefficient of variation is high (H) if it is greater than or equal to 0.5 and low (L) otherwise. In Figure E1 of Appendix E of the Appendix, we observe that instances with high expected demand at both locations exhibit the most module movements. In such problems, total demand in a period may frequently exceed capacity and create an incentive to relocate modules to decrease backorder costs. Instances with high variation at both locations and high expected demand at either location follow next. Instances with opposite expected demand and variation levels at the two locations exhibit a moderate number of module movements. The lowest number of movements is observed when both locations have low expected demand and/or low coefficients of variation. As expected, the relative magnitudes of expected demands and coefficients of variation play a significant role in the number of module movements observed in an instance. Table 1 summarizes the usefulness

Table 1. Usefulness of mobile modular production for Instance Set 1

	low ED at both	low, high EDs	high ED at both	Overall
high CV at both	Little	Great	Greatest	Great
low, high CVs	Little	Moderate	Greatest	Moderate
low CV at both	Least	Little	Greatest	Moderate
Overall	Little	Moderate	Greatest	

of mobile modular production determined by the percentage of instances exhibiting any module movement in each subset as “Least” for $[0, 10)\%$, “Little” for $[10, 20)\%$, “Moderate” for $[20, 45)\%$, “Great” for $[45, 70)\%$, and “Greatest” for $[70, 100]\%$. We observe that there is greater benefit in instances with high coefficients of variation (CV) and expected demands (ED). Another key takeaway is that high expected demand induces more module movement than high coefficient of variation; as total expected demand approaches total system capacity, module movements become more valuable for correcting poor inventory position. The subset of instances of Set 1 which have high expected demand at both locations (and hence are expected to show a higher tendency of module movement than the other instances) will be referred to as the HH instances of Set 1.

In Table 2, we summarize the quality of the lower and upper bounds given varying total number of modules Y , length of horizon T , backorder unit cost b , holding unit cost h , and movement unit cost \bar{K}^M for Set 1. PIR is a consistently weak lower bound for all T , and grows

Table 2. Variation of bound quality compared to optimal objective with varying T , Y , b , h , and \bar{K}^M for Instance Set 1

T	PIR	LB	UB _{max}	UB _{min}	Y	PIR	LB	UB _{max}	UB _{min}
5	0.224	0.946	4.8	1.021	3	0.520	0.871	4.2	1.119
7	0.228	0.938	6.5	1.033	5	0.221	0.923	7.4	1.059
10	0.221	0.930	9.2	1.062	7	0.105	0.961	9.6	1.020
15	0.199	0.931	13.2	1.086	10	0.026	0.990	12.4	1.004

(a) T (b) Y

b	PIR	LB	UB _{max}	UB _{min}	h	PIR	LB	UB _{max}	UB _{min}
1	0.218	0.938	6.4	1.036	0.5	0.234	0.927	10.5	1.060
2	0.218	0.934	10.5	1.065	1	0.201	0.945	6.4	1.041

(c) b (d) h

\bar{K}^M	PIR	LB	UB _{max}	UB _{min}
0.5	0.181	0.960	8.6	1.079
5	0.255	0.912	8.2	1.021
Overall	0.218	0.936	8.4	1.050

(e) \bar{K}^M

weaker for larger values of Y . It is not particularly surprising that the PIR bounds are so weak in this setting, since perfect information allows a planner to avoid all cost sources by eliminating inventory holding, backordering, and module movements when possible with optimally-selected production quantities. This is especially true when Y grows larger; in such instances, the large number of modules allows a substantial fraction of capacity to remain fixed in place and module movements under perfect information become even more rare. We note that the PIR bounds, although bad, are significantly better for HH instances (Table E1 in Section E of the Appendix). For these instances, the optimal policy under uncertainty generates decisions closer to those that would result under a perfect information assumption. Since these instances also result in more module movement, a natural hypothesis is that problem instances where good solutions require many module movements will benefit less overall from complete information than instances where module movements are rare even under uncertainty.

The lower bound derived by assuming the most flexible module relocation system (LB) is a much tighter bound. We note that the performance of both LB and UB_{min} deteriorates slightly with increasing T . This behavior is expected, and is likely due primarily to end-of-horizon effects. Both of these measures improve with a larger module fleet size Y . There are two reasons: first, correct base capacity can be installed with more precise control at each location and second, the refined production capacity control enabled by the larger fleet provides added benefit. It is also not unexpected that the policy created for LB moves modules more than a policy that incurs a cost for these moves; this tends to increase the optimality gap at higher movement and backorder

rates \bar{K}^M , b , and lower holding rate h (Table 2). On the other hand, when b is low or \bar{K}^M is high, optimal policies do not move modules very frequently and thus the upper bound given by UB_{\min} grows tighter. The cost of the worst fixed configuration, UB_{\max} , provides a benchmark for the quality of more intelligent configurations and dynamic policies. The weak performance of these upper bounds also show that if stationary demand forecasts were highly inaccurate, very poor fixed system configurations could result with costs 5 to 13 times those from smarter configurations.

Table E2 in Section E of the Appendix presents results that demonstrate the value of mobile production capacity for Instance Set 1 by comparing costs with mobile modules to costs incurred by operating a system with capacity fixed in place. Optimal costs are separated into components for moving modules, backordering demand, and holding inventory, and compared for the fixed capacity policies and optimal policies. Considering only the HH instances, a significant reduction (18% on average and 30% for $T = 15$) in the cost due to backordering is observed when production capacity is mobile.

Figure 3 compares the performance of various heuristics on set 1 instances, and specifically the HH instances. In each of the sub-figures of Figure 3, the ratios of the heuristic costs for policies MP, RF, RLB, LFP, and LIU to OPT are presented. The performance of MP improves with increase in Y and declines for larger horizon lengths. RLB performs better than MP at all levels of Y and T with an overall average of 1.02 times OPT. RF outperforms RLB leading to an overall average gap of 1% over OPT. LFP's performance is always similar to or better than that of LIU. For the $L = 2$ problem, LFP leads to the lowest gap above OPT of 0.3% on average and it improves with increase in Y . Although it performs quite well, LFP requires significantly more computational effort than MP, RLB, or RF.

When only HH instances are considered, the performance of the proposed heuristics remains good although the gaps are slightly higher. Since HH instances require more module movements in optimal solutions, this performance degradation is not unexpected. We now discuss the module movements prescribed by these heuristics comparatively. All instances with module movements prescribed by MP and RLB have module movements prescribed by the optimal policy as well. However, MP and RLB are more conservative than the optimal policy and exhibit movement in only 66% and 74% of the instances where the optimal policy prescribes movements. This result is intuitive as MP does not account for the sustained future benefit after moving modules and RLB would not be eager to shift modules in the present due to its assumption of no costs for future

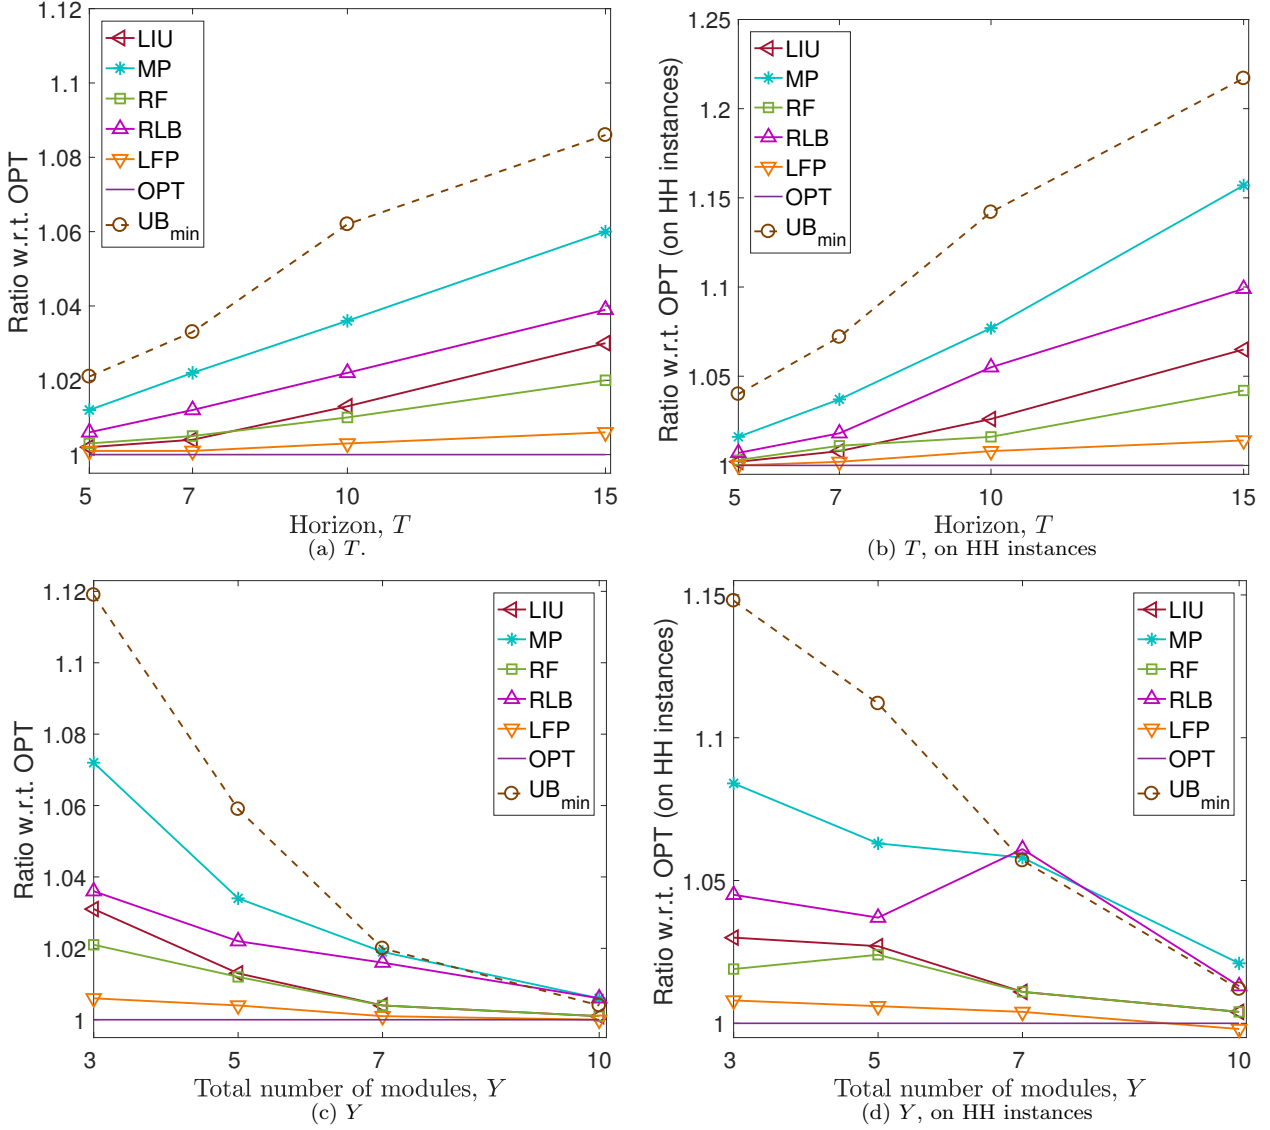


Figure 3. Performance of heuristics compared to optimal policies with varying T and Y for Instance Set 1

module movements. RF and LFP show movement in most (more than 97%) of the instances that show movement by the optimal policy, indicating good policy performance of these heuristics. All the heuristics yield better results for low b and high h (see Table E3b of the Appendix). For high \bar{K}^M , LIU, RF, and LFP perform well but MP and RLB are closer to optimality when \bar{K}^M is low, as expected (see Table E3b of the Appendix).

Focusing on the HH subset of Set 1, we note that bounds LB and UB_{\min} are worse on HH instances than on the overall set, again emphasizing that these instances require module movements to achieve low costs. We observe that the optimality gap of each heuristic on HH instances is higher than the overall average although the magnitude of the gap itself is still within 8%. Figure 3 shows that LFP outperforms all the heuristics and is closely followed by RF.

6.2. Results for Set 2

This set consists of 720 two location instances with relatively higher values of horizon length T and module fleet size Y , along with more spread in backorder cost b , module movement cost K , and production capacity per module G . We compare the performance of the heuristics MP, RLB, RF, LFP, and LIU and the lower bound LB with UB_{\min} .

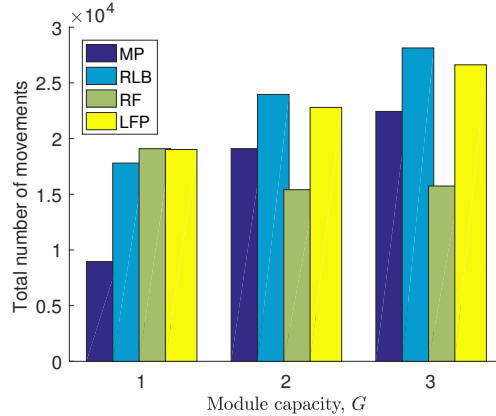


Figure 4. Module movement totals given varying G for Instance Set 2

Figure 4 totals the number of module movements prescribed by the different heuristic policies at different module capacity levels across the instances. MP was the most conservative policy for all values of G since its goal is to minimize immediate costs only. RF restricted movements at higher capacity levels since it assumes that modules cannot be moved again in the future. RLB moved modules more frequently compared to MP and RF since it assumes that the cost of recourse from poor capacity configurations is zero. It is not surprising to note that LFP exhibits a movement count that falls between RF and RLB, since LFP blends both lower and upper bounds in its future cost estimate.

Although we do not present these results in a figure, we note that UB_{\max} increases with increasing T , G , and b but decreases with an increase in Y , as expected. The average ratio between UB_{\max} and UB_{\min} was 13.9 for $G = 1$ but increases 20.1 for $G = 3$, again motivating that poor initial capacity configurations can be very expensive. Figure 5 summarizes overall heuristic performance for instances in Set 2 when Y , G , \bar{K}^M , and b are varied, while Figure 6a summarizes performance with varying T . LB is tighter for higher G and Y , and lower b , \bar{K}^M , and T , which is expected as mobile modularity is not fully utilized in these situations due to lower incentives for movement. We note that RF and LFP again performed much closer to optimality and significantly better than MP, LIU, and RLB.

The best heuristic LFP delivered an average cost reduction of about 4% compared to

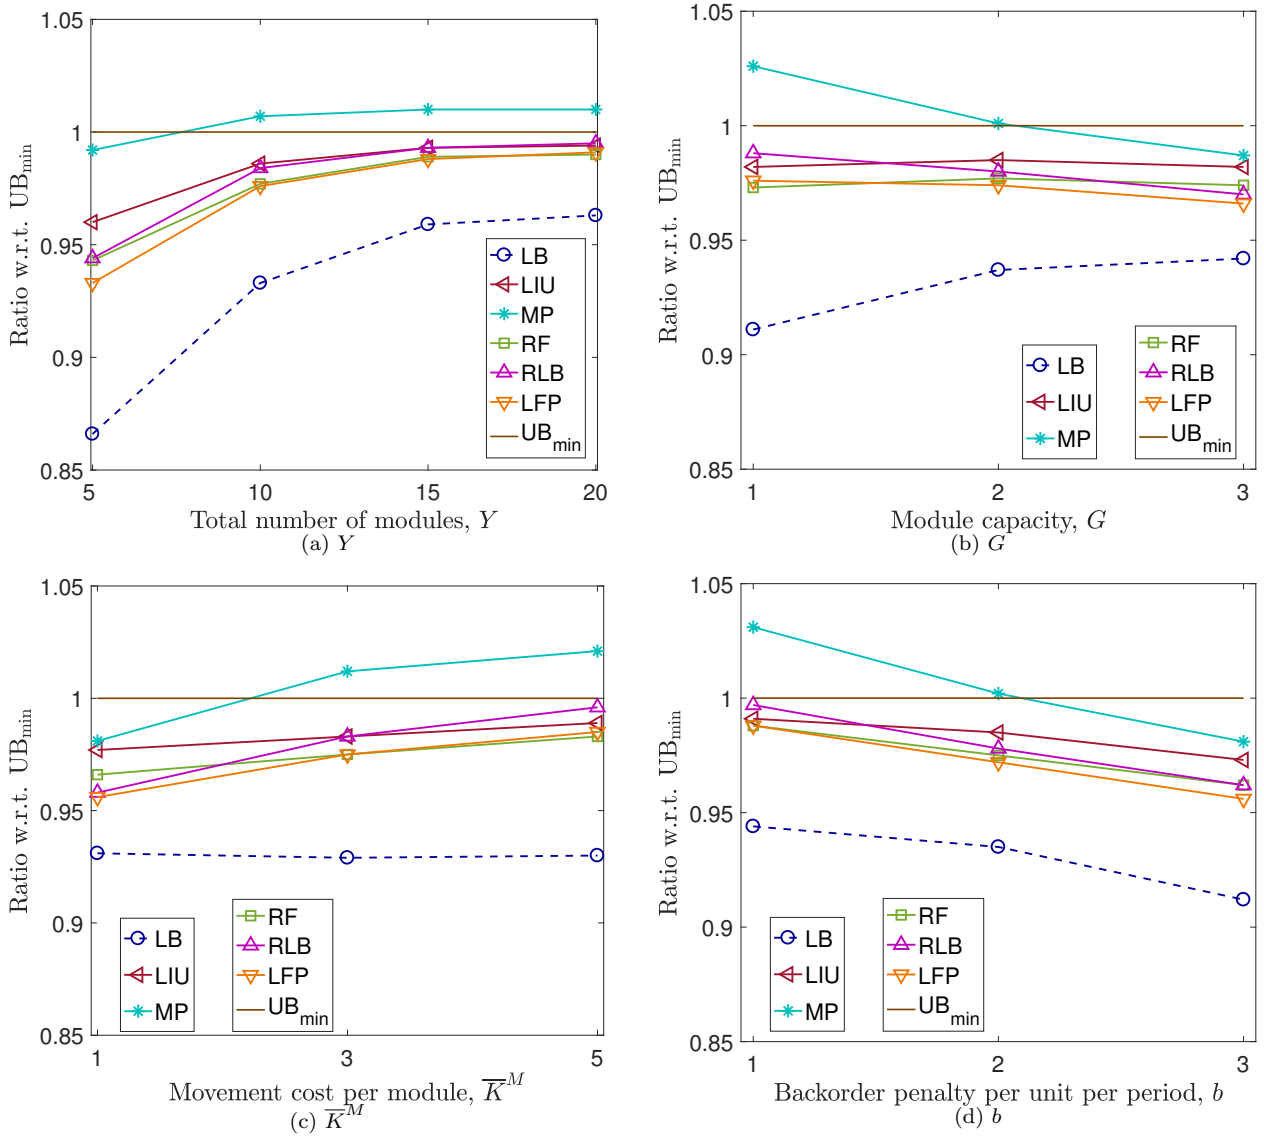


Figure 5. Overall performance of heuristics compared to UB_{\min} with varying Y , G , \bar{K}^M , and b for Instance Set 2

the fixed configuration UB_{\min} on Set 2. Focusing on the HH instances of Set 2 (see Figure 6b), we note a 5% average advantage over fixed systems for LFP, which is close to RF(4.8%), and higher than RLB (4%), LIU (3%), and MP (-0.4%). LFP yielded 9% savings over UB_{\min} when the horizon length is 20. These results provide additional confirmation that the proposed heuristics are effective for problems with $L = 2$ production locations even when underlying demand distributions are stationary.

6.3. Results for Set 3

This set consists of 540 $L \geq 2$ instances generated by the procedure described in Section C of the Appendix. We study the performance of the suboptimal policies generated by three heuristics: the the myopic policy (MP), the lookahead with approximate fixed future (LAF), and the rollout




Figure 6. Heuristic performance compared to UB_{\min} with varying T for Instance Set 2

of fixed future (RF). The initial inventory state of the system is zero inventory at all locations. Two different initial capacity states are compared, where the first assumes that modules are located in the configuration u_{\min} that minimizes the expected total cost of a fixed capacity system and the second assumes modules are allocated to locations in a configuration u_{simple} that depends on the mean and variance of demand at each location (see Section D of the Appendix for details).

Table 3. Average computation time in seconds for $T = 15$ instances with varying L for Instance Set 3

L	UB_{\min} (per instance)	MP (per trajectory)	LAF	RF
2	0.08	0.05	0.08	0.21
3	0.23	0.07	0.13	0.33
5	0.87	0.10	0.21	0.74
10	7.30	0.23	0.48	3.6
20	96.0	0.77	1.30	33.7

The computations for the study on Set 3 are performed on a single server of an Intel Xeon Processor E5-2670 workstation. We first examine the computation times required for finding the best upper bound UB_{\min} (one-time and offline) and for determining the heuristic policies MP, LAF, and RF (Table 3) on $T = 15$ instances. We note that LAF is significantly faster than RF and is comparable to MP in the order of magnitude of seconds required. Figure 7 shows that as L increases the effort required to implement RF increases dramatically and that LAF is on average 35 times faster than RF for $L = 20$. Additionally, we note that using LAF takes less than twice the effort required for implementing the naive MP policy, even when L is high. Thus, LAF presents notable computational advantages over RF. Figure 8 demonstrates the performance

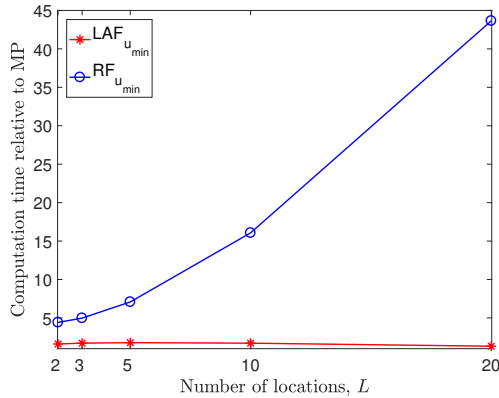


Figure 7. Average computational effort relative to MP with varying L for Instance Set 3

of heuristic policies for Instance Set 3 with varying L , T , \bar{K}^M , and b . We note that LAF and RF perform almost identically, suggesting that the approximation of fixed future proposed in LAF achieves almost identical savings as RF with a remarkable improvement in computational efficiency. LAF and RF outperform UB_{\min} by 4 to 9% and the savings increase with L on average. Also, smaller savings are obtained for higher Y for any given L , since fewer modules are moved between locations due to greater availability of capacity at each location. MP results in average costs that are sometimes higher than UB_{\min} also, indicating that naive heuristic policies may not outperform systems with fixed capacity installations. Figure 8a demonstrates that both RF and LAF are able to create more value when the number of locations L grows, but that the marginal improvement decreases with L . Table E4 in Appendix E of the Appendix depicts the varying quality of the policies found by RF and LAF with varying L and Y . Figure 8c confirms the intuition that with increasing \bar{K}^M , LAF and RF are not able to generate as much benefit over a fixed capacity configuration. MP appears to be more sensitive to higher values of \bar{K}^M , as moves made for immediate cost benefits may induce additional movements to the escape the poor state that was reached. As the backorder cost b increases, there is greater use of mobility of production capacity (Figure 8d). In Figure 8b, we again see that longer horizons yield more accumulated benefit for mobile production. Motivated by this observation, we present the performance of the heuristics with varying L for the longest horizon, $T = 15$, in Figure 9b. We observe a significant average value addition ranging between 9 to 16%.

We finally compare the performance of both fixed capacity systems and mobile modular production systems operated with our heuristics when the initial capacity configuration is optimized versus when it is set heuristically. Configuration u_{simple} is determined with a simple approach that is based on each location's expected demand and variability (see Figure 9), while

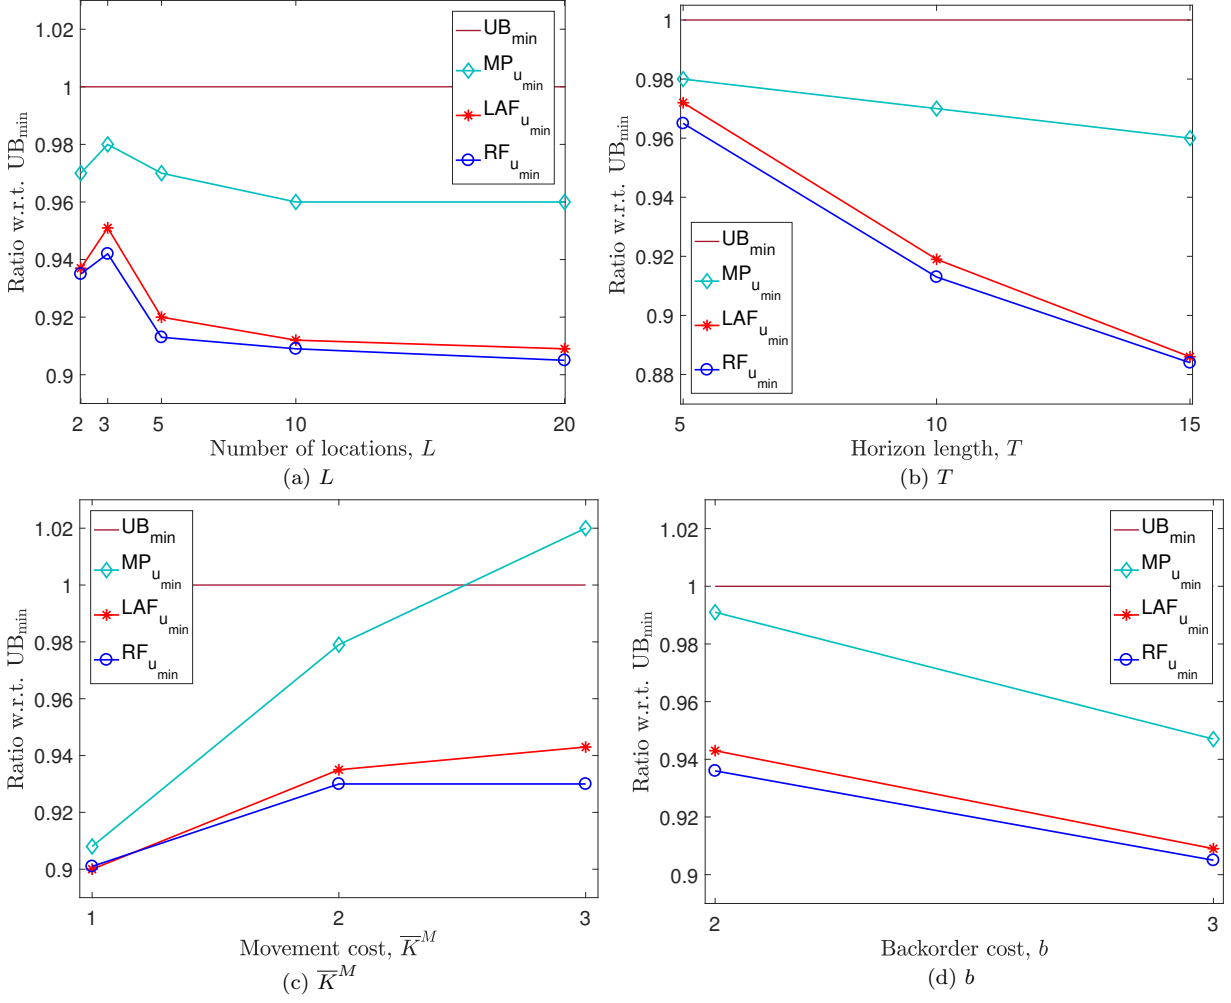


Figure 8. Performance of heuristics with varying L , T , \bar{K}^M , and b when capacity initialized to u_{\min} for Instance Set 3

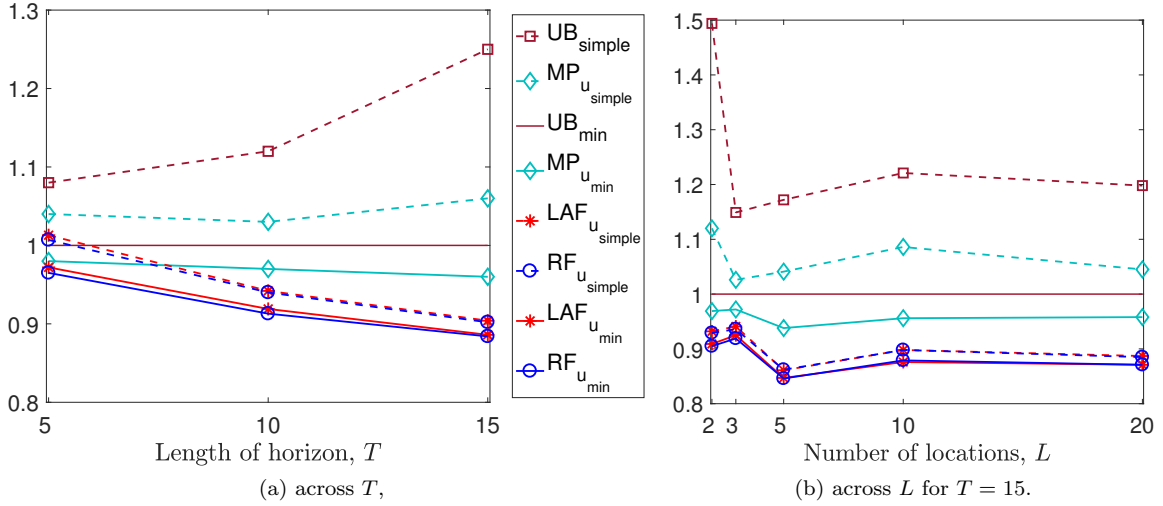


Figure 9. Performance of heuristics with varying T and L when initialized to u_{\min} and u_{simple} for Instance Set 3

u_{\min} is the optimal fixed configuration that leads to lowest cost when capacity is fixed. In Figure 9a, we note that the gap between the implementations of RF with two different initial con-

figurations reduces as the horizon grows longer. We may infer that the difference in heuristic performance is due to the amortization of a one time movement cost incurred to switch from subsequent poor states induced by u_{simple} to better states potentially induced by u_{min} . The fixed system configured to u_{simple} costs about 25% more than UB_{min} when the length of the horizon T is 15 (Figure 9b). However, setting the initial state to u_{simple} instead of u_{min} leads to only about 2% decrease in the gaps of LAF and RF at $T = 15$. This observation establishes that the mobile modular system is indeed very resilient compared to fixed systems since it is able to retrieve most of the savings over fixed systems even with suboptimal initial configurations (about 25-30% savings over the fixed system configured to u_{simple}). These observations indicate that LAF and RF perform robustly irrespective of the initial capacity configuration.

7. Conclusion

In this paper, we introduced the multilocation mobile modular production and inventory problem and presented efficient and effective techniques for determining effective operating policies based on approximate dynamic programming. The heuristic policies generate significant value addition over fixed capacity systems, providing evidence of the efficiency, practicality, and resiliency of mobile modular production systems. Our heuristics perform very close to the optimal solution for $L = 2$ and significantly better than fixed capacity systems for $L \geq 2$. Through our computational study, we demonstrate the cost-effectiveness and resilience of mobile modular systems in comparison to fixed systems. The heuristics RF and LFP mimic the optimal policy very closely for $L = 2$ and the heuristic LAF achieves excellent results and remarkable computational efficiency for $L \geq 2$. Future work will focus on systems with non-stationary demand with parameters that must be learned over time. Additionally, future models may also allow inventory transshipment between production locations as an alternative to module movements as well as module movement lead times.

References

- Angelus, A. and Porteus, E. L. (2002). Simultaneous capacity and production management of short-life-cycle, produce-to-stock goods under stochastic demand. *Management Science*, 48(3):399–413.
- Aviv, Y. and Federgruen, A. (1997). Stochastic inventory models with limited production capacity and periodically varying parameters. *Probability in the Engineering and Informational Sciences*, 11:107–135.
- Bertsekas, D. P. (2000). *Dynamic Programming and Optimal Control*. Athena Scientific, 3rd edition.

- Bertsekas, D. P., Tsitsiklis, J. N., and Wu, C. (1997). Rollout algorithms for combinatorial optimization. *Journal of Heuristics*, 3:245–262.
- Bi, Z. M., Lang, S. Y. T., Shen, W., and Wang, L. (2008). Reconfigurable manufacturing systems: the state of the art. *International Journal of Production Research*, 46(4):967–992.
- Bradley, J. R. and Glynn, P. W. (2002). Managing capacity and inventory jointly in manufacturing systems. *Management Science*, 48(2):273–288.
- Chen, W., Dawande, M., and Janakiraman, G. (2014). Integrality in stochastic inventory models. *Production and Operations Management*, 23(9):1646–1663.
- Clausen, U., Langkau, S., Goedicke, I., and Lier, S. (2015). Location and network planning for modular container plants in the process industry. *Chemie Ingenieur Technik*, 87(10):1365–1372.
- F3 (2014). Flexible, fast and future production processes. Technical report, Bayer Technology Services GMBH.
- Geek Wire (2018). Amazon finally wins a patent for 3-D printing on demand, for pickup or delivery. [Online; accessed Nov 25, 2018].
- Ghiani, G., Guerriero, F., and Musmanno, R. (2002). The capacitated plant location problem with multiple facilities in the same site. *Computers & Operations Research*, 29(13):1903–1912.
- Goodson, J. C., Thomas, B. W., and Ohlmann, J. W. (2017). A rollout algorithm framework for heuristic solutions to finite-horizon stochastic dynamic programs. *European Journal of Operational Research*, 258(1):216 – 229.
- Halper, R. and Raghavan, S. (2011). The mobile facility routing problem. *Transportation Science*, 45(3):413–434.
- Ibaraki, T., Hasegawa, T., Teranaka, K., and Iwase, J. (1978). The multiple-choice knapsack problem. *Journal of the Operations Research Society of Japan*, 21(1):59–95.
- Iyer, A. V. and Jain, A. (2004). Modeling the impact of merging capacity in production-inventory systems. *Management Science*, 50(8):1082–1094.
- Jena, S. D., Cordeau, J.-F., and Gendron, B. (2015). Dynamic facility location with generalized modular capacities. *Transportation Science*, 49(3):489–499.
- Kessler, S. and Bruell, L. (2015). New production concepts for the process industry require modular logistics solutions. *ChemBioEng Reviews*, 2(5):351–355.
- Koren, Y. (2010). Reconfigurable manufacturing systems. In *The Global Manufacturing Revolution: Product-Process-Business Integration and Reconfigurable Systems*. John Wiley & Sons, Inc.
- Lier, S. and Grünewald, M. (2011). Net present value analysis of modular chemical production plants. *Chemical Engineering & Technology*, 34(5):809–816.
- Lier, S., Wörsdörfer, D., and Gesing, J. (2013). Business models and product service systems for transformable, modular plants in the chemical process industry. In *Product-Service Integration for Sustainable Solutions: Proceedings of the 5th CIRP International Conference on Industrial Product-Service Systems*, pages 227–238. Springer Berlin Heidelberg.

- Lier, S., Wörsdörfer, D., and Grünewald, M. (2016). Transformable production concepts: Flexible, mobile, decentralized, modular, fast. *ChemBioEng Reviews*, 3(1).
- Marcotte, S. and Montreuil, B. (2016). Introducing the concept of hyperconnected mobile production. *Progress in Material Handling Research*.
- Marcotte, S., Montreuil, B., and Coelho, L. C. (2015). Modeling of physical internet enabled hyperconnected modular production. In *Proceedings of 2nd Physical Internet Conference, 2015*.
- Melo, M., Nickel, S., and Saldanha da Gama, F. (2005). Dynamic multi-commodity capacitated facility location: a mathematical modeling framework for strategic supply chain planning. *Computers & Operations Research*, 33(1):181–208.
- Nemhauser, G. L. and Wolsey, L. A. (1988). *Integer and Combinatorial Optimization*. Wiley-Interscience, New York, NY, USA.
- Novartis (2016). Pharmacy on Demand. Accessed: 2016-05-18.
- Pfizer (2015). Pfizer announces collaboration with GSK on next-generation design of portable, continuous, miniature and modular (PCMM) oral solid dose development and manufacturing units. Accessed: 2016-05-18.
- Powell, W. B. (2007). *Approximate Dynamic Programming: Solving the Curses of Dimensionality (Wiley Series in Probability and Statistics)*. John Wiley & Sons, Inc., Hoboken, New Jersey, second edition.
- Qiu, J. and Sharkey, T. C. (2013). Integrated dynamic single-facility location and inventory planning problems. *IIE Transactions*, 45(8):883–895.
- Rajagopalan, S. and Swaminathan, J. M. (2001). A coordinated production planning model with capacity expansion and inventory management. *Management Science*, 47(1):1562–1580.
- Rogers, G. G. and Bottaci, L. (1997). Modular production systems: a new manufacturing paradigm. *Journal of Intelligent Manufacturing*, 8(2):147–156.
- Ryzhov, I. O., Powell, W. B., and I, F. P. (2012). The knowledge gradient algorithm for a general class of online learning problems. *Operations Research*, 60(1):180–195.
- Secomandi, N. (2001). A rollout policy for the vehicle routing problem with stochastic demands. *Operations Research*, 49(5):796–802.
- Sinha, P. and Zoltners, A. A. (1979). The multiple-choice knapsack problem. *Operations Research*, 27(3):503–515.
- Tesauro, G. and Galperin, G. R. (1996). On-line policy improvement using monte carlo search. *Advances in Neural Information Processing Systems*, 9:1068–1074.
- Wörsdörfer, D. and Lier, S. (2017). Optimized modular production networks in the process industry. In *Operations Research Proceedings 2015*. Springer International Publishing.
- Wörsdörfer, D., Lier, S., and Crasselt, N. (2017). Real options-based evaluation model for transformable plant designs in the process industry. *Journal of Manufacturing Systems*, 42:29 – 43.

Appendix A. Proofs of Theorems

We present the proofs of Theorem 5.1 and Theorem 5.2 in Section A.1 and Section A.2 respectively.

A.1. Proof of Theorem 5.1

Proof. (a) Since the newsvendor constraint matrices are totally unimodular (TU) already, their rows can be separated into partitions suitably to satisfy (Nemhauser and Wolsey, 1988, Theorem 2.7, III.1). Total unimodularity is retained in the presence of location-wise capacity constraints when $G = 1$. Hence, the constraint matrix of MMPIP-MP is TU when $G = 1$. Since the right hand side is integral, the polyhedron of the IP presented here for the implementation of the myopic policy MP is integral.

(b) The objective function of each of the L newsvendor problems is given by

$$\sum_{k \in \mathcal{K}} p_i^k \left\{ h_i (s_i + q_i - d_i^k)^+ + b_i (d_i^k - s_i - q_i)^+ \right\}.$$

We note that $h_i P(D_i \leq d_i^m) - b_i P(D_i > d_i^m)$ is the slope of the newsvendor objective function in the interval $d_i^m \leq q_i + s_i \leq d_i^{m+1}$. As the function is convex, its slope is decreasing from left to right and increasing from right to left. The current problem involves re-allocation of capacity followed by replenishment. In the one period problem, a module movement from location i to location j occurs if

- i. the available capacity is greater than the unconstrained optimal order quantity at the sending location i , or
- ii. if the positive gain per unit at the receiving location j , $h_j P(D_j \leq d_j^m) - b_j P(D_j > d_j^m)$, is higher than the (absolute value of) loss at the sending location i , $h_i P(D_i \leq d_i^m) - b_i P(D_i > d_i^m)$. The movement is feasible only if $h_i P(D_i \leq d_i^m) - b_i P(D_i > d_i^m) + h_j P(D_j \leq d_j^m) - b_j P(D_j > d_j^m) > K_{ij}^M / G$. The movement will be beneficial until this inequality is reversed when either $s_j + q_j$ or $s_i + q_i$ reach a slope change point, namely, a demand outcome (which is integral for this problem setup). Since the starting inventory (s_i or s_j) is an integer and the order-up-to level at one of the locations is integral, it follows that the amount of effective potential inventory increment/decrement at the locations is integral and hence the other order-up-to level is also integral. This mechanism is in action for all module shifts induced by the cost structure and hence, irrespective of the integrality of $\{\{\Delta_{ij}^M\}\}$, for all i , replenishment quantity q_i will be

take integer values even without integrality constraints. □

A.2. Proof of Theorem 5.2

Proof. MMPIP-LAF is different from MMPIP-MP in the objective function that now contains ζ_i and η_i terms additionally. The constraints now include the description of the ζ_i and η_i terms as the maximum over piecewise linear facets of sets of convex curves. As these inventory and capacity curves have with integer slope transition points, the minimization of $\sum_{i=1}^L (\zeta_i + \eta_i)/2$ ensures that the y_i 's and q_i 's will still be integers along with the constraints of MMPIP-MP. Having integer values for y_i ensures that Δ_{ij}^M 's are integers as y_i 's are placeholder variables. Integer values of q_i 's guarantees that r_i^k and o_i^k are integers as $r_i^k + o_i^k = s_i + q_i - d_i^k$ and d_i^k 's are integers. Hence, MMPIP-LAF can be solved as a linear program. □

Appendix B. Algorithm for LIU

In this section, we present Algorithm 1 for implementing the heuristic LIU. Algorithm 1 is the

Algorithm 1: Approximate Value Iteration over finite horizon: LIU

Step 1. Initialization:

- a. Initialize $\bar{V}_t^0(\xi(t))$ for all states $\xi(t)$.
- b. Choose an initial state $\xi(1)$.
- c. Set $n = 1$.

Step 2. Choose a sample path d^n .

Step 3. For $t = 1, \dots, T - 1$ do:

- a. Solve $\hat{v}_t^n = \min_{a \in A(\xi^n(t))} \mathbb{E}_D \left(C_t(\xi^n(t), a, D) + \bar{V}_{t+1}^{n-1}(f(\xi^n(t), a, D)) \right)$ and let a_t^n be a minimizer.
- b. Blending: $\bar{V}_t^n(\xi^n(t)) = \alpha_n \hat{v}_t^n + (1 - \alpha_n) \bar{V}_t^{n-1}(\xi^n(t))$.
- c. Greedy trajectory following: $\xi^n(t+1) = f(\xi^n(t), a_t^n, d^n(t))$.

Step 4. Let $n = n + 1$. If $n < N$, go to step 2.

value function approximation algorithm used to implement LIU.

Appendix C. Instance Sets for Computational Study

In the section, we describe the procedures to generate the three instance sets used in our computational study.

C.1. Set 1

This set contain $L = 2$ instances. We fix $G = 1$ and $c = 0$. For each combination of the parameters $Y \in \{3, 5, 7, 10\}$, $T \in \{5, 7, 10, 15\}$, $b \in \{1, 2\}$, $h \in \{0.5, 1\}$, and $\bar{K}^M \in \{0.5, 5\}$, we generate 10 demand distributions in the following fashion. For each of the two stations:

- set the expected total demand to αGY , where $\alpha = 0.75$,
- randomly obtain the fraction of expected demand at station 1 as $\beta \in [0.2, 0.8]$, setting the expected demands to $\mu_1 = \beta\mu$; $\mu_2 = (1 - \beta)\mu$,
- randomly generate three distinct samples each from $\{0, \dots, \lceil \gamma\mu_1 \rceil\}$ and $\{0, \dots, \lceil \gamma\mu_2 \rceil\}$, where $\gamma = 1.1$, and sort them to obtain the vector of demand outcomes, and
- randomly generate three numbers from $\{1, 2, 3, 4, 5, 6\}$ and obtain probabilities from them by dividing each by their sum.

Thus, we generate a total of $4 \times 4 \times 2 \times 2 \times 2 \times 10 = 1280$ instances.

C.2. Set 2

This set contains $L = 2$ instances. We fix $h = 1$ and $c = 0$. For each combination of the parameters $Y \in \{5, 10, 15, 20\}$, $T \in \{5, 10, 15, 20\}$, $G \in \{1, 2, 3\}$, $b \in \{1, 2, 3\}$, and $\overline{K}^M \in \{1, 3, 5\}$, we generate 5 demand distributions in the following fashion. For each of the two stations:

- set expected total demand to αGY , where $\alpha = 0.8$,
- randomly obtain the fraction of expected demand at station 1: $\beta \in [0.2, 0.8]$, setting the expected demands to $\mu_1 = \beta\mu$; $\mu_2 = (1 - \beta)\mu$,
- randomly generate three distinct samples each from $\{0, \dots, \lceil \gamma\mu_1 \rceil\}$ and $\{0, \dots, \lceil \gamma\mu_2 \rceil\}$, where $\gamma = 1.5$, and sort them to obtain the vector of demand outcomes, and
- randomly generate three numbers from $\{1, 2, 3, 4, 5, 6\}$ and obtain probabilities from them by dividing each by their sum.

Thus, we generate a total of $4 \times 4 \times 3 \times 3 \times 3 \times 5 = 2160$ instances.

C.3. Set 3

We fix the following parameters: $G = 1$, $c = 0$, and $h = 1$. For each combination of the parameters $L \in \{2, 3, 5, 10, 20\}$, $Y \in \{\lceil 1.9L \rceil, \lceil 2.4L \rceil\}$, $T \in \{5, 10, 15\}$, $b \in \{2, 3\}$, and $\overline{K}^M \in \{1, 2, 3\}$, we generate 3 demand distributions in the following fashion. For each of the L locations:

- randomly generate three distinct samples each from $\{0, 1, 2, 3\}$ and sort them to obtain the vector of demand outcomes and
- randomly generate three numbers and obtain probabilities from them by dividing each by their sum.

Thus, we generate a total of $5 \times 2 \times 3 \times 2 \times 3 \times 3 = 540$ instances.

Appendix D. A Simple Strategy for Module Allocation u_{simple}

The module allocation strategy u_{simple} is given below:

- (1) Assign $\{\lceil \rho_i Y \rceil\}$ to each location i where, $\rho_i = \frac{\sigma_i}{\sum_{j=1}^L \sigma_j} \times \beta + \frac{\mu_i}{\sum_{j=1}^L \mu_j} \times (1 - \beta)$, $\beta = \frac{\sum_{j=1}^L z_\alpha \sigma_j}{\sum_{j=1}^L (z_\alpha \sigma_j + \mu_j)}$, and $z_\alpha = 1.64$ for a 95% service level.
- (2) If $\sum_{i=1}^L \lceil \rho_i Y \rceil < Y$, then allocate the remaining modules in the decreasing order of coefficient of variation, σ_i / μ_i , of locations.

Appendix E. Additional Figures and Tables in Computational Study

We present additional figures and tables that complement the narrative of the paper.




Figure E1. Effect of variability of demand on movement tendency for Instance Set 1

Table E1. Variation of bounds compared to OPT for HH instances with varying T , Y , and \bar{K}^M for Instance Set 1

T	PIR	LB	UB _{max}	UB _{min}	Y	PIR	LB	UB _{max}	UB _{min}
5	0.498	0.900	3.9	1.040	3	0.626	0.865	3.9	1.148
7	0.482	0.893	5.9	1.072	5	0.490	0.859	7.1	1.112
10	0.521	0.866	6.5	1.142	7	0.246	0.917	11.3	1.057
15	0.499	0.855	9.8	1.217	10	0.061	0.975	16.0	1.012

(a) T

(b) Y

b	PIR	LB	UB _{max}	UB _{min}	h	PIR	LB	UB _{max}	UB _{min}
1	0.500	0.883	5.0	1.086	0.5	0.525	0.861	7.7	1.132
2	0.500	0.874	8.1	1.148	1	0.472	0.898	5.3	1.101

(c) b

(d) h

\bar{K}^M	PIR	LB	UB _{max}	UB _{min}
0.5	0.437	0.928	6.8	1.180
5	0.560	0.831	6.3	1.056
Overall	0.500	0.879	6.6	1.117

(e) \bar{K}^M

Table E2. Value addition of optimal mobile modular production system over fixed system with varying Y , T , \bar{K}^M , h , and b for Instance Set 1

	All			HH Instances		
Y	ΔK	$-\Delta B$	$-\Delta H$	ΔK	$-\Delta B$	$-\Delta H$
3	0.049	0.182	-0.011	0.055	0.223	-0.019
5	0.025	0.093	-0.010	0.044	0.170	-0.019
7	0.012	0.033	-0.001	0.026	0.097	-0.011
10	0.003	0.009	-0.002	0.008	0.033	-0.004
T	ΔK	$-\Delta B$	$-\Delta H$	ΔK	$-\Delta B$	$-\Delta H$
5	0.012	0.038	-0.005	0.025	0.075	-0.011
7	0.020	0.061	-0.004	0.035	0.123	-0.014
10	0.027	0.098	-0.010	0.057	0.221	-0.027
15	0.029	0.121	-0.004	0.065	0.304	-0.015
\bar{K}^M	ΔK	$-\Delta B$	$-\Delta H$	ΔK	$-\Delta B$	$-\Delta H$
0.5	0.026	0.112	-0.005	0.046	0.248	-0.021
5	0.019	0.047	-0.007	0.044	0.113	-0.012
h	ΔK	$-\Delta B$	$-\Delta H$	ΔK	$-\Delta B$	$-\Delta H$
0.5	0.025	0.093	-0.007	0.048	0.199	-0.017
1	0.019	0.065	-0.005	0.041	0.158	-0.017
b	ΔK	$-\Delta B$	$-\Delta H$	ΔK	$-\Delta B$	$-\Delta H$
1	0.020	0.061	-0.005	0.041	0.139	-0.012
2	0.024	0.097	-0.006	0.049	0.220	-0.021
Overall	0.022	0.079	-0.006	0.045	0.179	-0.017

Table E3. Performance of heuristics relative to OPT with varying b , h , and \bar{K}^M

		LIU	MP	RF	RLB	LFP
b	1	1.020	1.066	1.014	1.032	1.005
	2	1.030	1.076	1.022	1.056	1.007
h	0.5	1.028	1.068	1.019	1.049	1.007
	1	1.022	1.075	1.017	1.039	1.005
\bar{K}^M	0.5	1.037	1.052	1.027	1.009	1.004
	5	1.013	1.090	1.009	1.078	1.008
Overall		1.025	1.071	1.018	1.044	1.006

(a) for HH instances of Instance Set 1,

		LIU	MP	RF	RLB	LFP
b	1	1.009	1.031	1.006	1.016	1.002
	2	1.016	1.034	1.013	1.023	1.003
h	0.5	1.014	1.032	1.010	1.022	1.003
	1	1.011	1.033	1.009	1.018	1.003
\bar{K}^M	0.5	1.020	1.023	1.015	1.005	1.003
	5	1.005	1.042	1.004	1.034	1.003
Overall		1.013	1.033	1.010	1.020	1.003

(b) for Instance Set 1.

Table E4. Performance of heuristics compared to UB_{\min} with varying L and Y when initialized to u_{\min} for Instance Set 3

L	2		3		5		10		20		Overall
Y	4	5	6	8	10	12	19	24	38	48	
MP	0.96	0.97	0.96	1	0.94	1	0.93	1	0.93	1	0.97
LAF	0.914	0.96	0.913	0.99	0.866	0.974	0.851	0.973	0.834	0.985	0.93
RF	0.908	0.961	0.897	0.986	0.86	0.965	0.848	0.97	0.829	0.981	0.92