

# Markov Decision Process Design: A Framework for Integrating Strategic and Operational Decisions

Seth Brown\*      Saumya Sinha†      Andrew Schaefer\*

## Abstract

We consider the problem of optimally designing a system for repeated use under uncertainty. We develop a modeling framework that integrates design and operational phases, which are represented by a mixed-integer program and discounted-cost infinite-horizon Markov decision processes, respectively. We seek to simultaneously minimize the design costs and the subsequent expected operational costs. This problem setting arises naturally in several application areas, as we illustrate through examples. We derive a bilevel mixed-integer linear programming formulation for the problem and perform a computational study to demonstrate that realistic instances can be solved numerically.

**Keywords:** Markov decision processes, bilevel optimization, design optimization

## 1 Introduction

We introduce a general modeling framework for the problem of optimally designing a system for repeated use under uncertainty. Consider an agent who seeks to make decisions in two phases, namely, a design phase followed by an operational phase. The design phase corresponds to strategic decisions that have long implementation cycles and are difficult to modify once executed. As such, these are here-and-now static decisions that are made upfront. The operational phase begins once the design decisions have been executed and any uncertainty in their outcomes has been realized. This phase corresponds to day-to-day dynamic decisions that are routinely made in an uncertain environment once the system is in use.

---

\*Computational Applied Mathematics & Operations Research, Rice University, Houston, TX, USA

†Industrial & Systems Engineering, University of Minnesota, Minneapolis, MN, USA (Corresponding author)

Examples of this setting arise naturally in several application areas. For instance, consider a firm that seeks to acquire a new fleet of vehicles. In the design phase, the firm determines the number and specifications of the vehicles to purchase, subject to budgetary constraints. The purchase orders may take years to fulfill and are difficult to revise. Moreover, the actual performance of the delivered vehicles is uncertain and unknown at the time of procurement. The operational phase corresponds to the routine deployment of the fleet after vehicle delivery. As another example, consider a retailer who seeks to open new retail outlets from a set of available choices. The design phase corresponds to choices regarding the number, location, and specifications (e.g., size and capacity) of the new outlets, while the operational phase corresponds to inventory management at the retail locations once they are ready for use.

We highlight two important features of this decision problem:

- (i) interdependence of the two decision phases: As evident from the examples, the two phases of decision-making are naturally interdependent – optimal design decisions must factor in the long-term operational cost of the enterprise, while the parameters of the operational problems are determined by the design phase.
- (ii) two sources of uncertainty: As first noted by Bailey et al. (2006), the decisions are affected by two distinct sources of uncertainty. Design uncertainty refers to uncertainty in the outcome of the design phase (e.g., variability in vehicle performance after purchase), and is not realized until the design phase has been executed. Operational uncertainty refers to the stochastic environment in which operational decisions are made (e.g., uncertain demand in the inventory management problem).

Despite the natural interdependence, standard modeling approaches typically view the two phases as disparate sub-components, due in part to the complexity of their relationship and high dimensionality of the resulting problem Burnak et al. (2019). Several authors have discussed the need to integrate different phases of decision-making Burnak et al. (2019); Grossmann (2005); Avraamidou and Pistikopoulos (2019); Badejo and Ierapetritou (2022), noting that the synergistic interactions across different phases can lead to gains in efficiency and profits. In this work, we present an optimization model that captures the hierarchy of these decisions, and incorporates both types of uncertainty. We model the design phase as a mixed-integer program (MIP), and the operational decisions via infinite-horizon discounted cost Markov decision processes (MDPs). We represent design uncertainty via a finite set of scenarios and operational uncertainty through

probabilistic state transitions in the MDPs. Then, we seek to minimize the combined expected cost of the MIP and the subsequent discounted-cost MDPs that are determined after the first-stage decisions are made. This problem may alternatively be viewed as one where we choose optimal design parameters for MDPs so as to minimize the expected long-run operational cost under a family of uncertain scenarios. To the best of our knowledge, we are the first to consider this problem of optimal MDP design.

Bailey et al. (2006) consider an *adversarial* version of this problem, where an agent first makes a design decision to minimize costs, and an opponent subsequently makes operational decisions to maximize the long-run costs to induce the greatest disruption to the agent. The contradiction of goals in the two stages allows for the reformulation of the problem as a stochastic MIP with continuous second-stage variables. In contrast, we follow a *cooperative* approach where the same agent makes the decisions in both stages, and the objectives in both stages are aligned so that we simultaneously minimize both the design and operational costs. This alignment renders the stochastic MIP formulation of Bailey et al. (2006) no longer achievable. Instead, we formulate the problem as a bilevel program. Note that this contrast between the cooperative and adversarial approaches is different from the usual distinction between ‘optimistic’ and ‘pessimistic’ formulations in bilevel optimization Dempe et al. (2015). In our framework, the decision-maker in the strategic phase (i.e., the leader problem) is only interested in the optimal values of the follower problems, and is unaffected by the particular choice of the followers’ optimal solutions. As such, questions about optimistic and pessimistic formulations do not arise.

Bilevel optimization has been used in several domains to integrate decisions that were previously modeled independently, such as facility location and routing González et al. (2022); Tordecilla et al. (2023), and design and scheduling of production systems Avraamidou and Pistikopoulos (2019); Leenders et al. (2023). To our knowledge, we are the first to consider the general problem of integrating the design and operational phases where the latter is modeled as an MDP, incorporating both types of uncertainty.

Recent years have witnessed significant advances in the development of solution approaches for bilevel programming problems, which are generally characterized by their (non-)linearity and the presence of discrete and/or continuous variables. Our formulation consists of a linear bilevel program with a mixed-integer leader problem and multiple follower problems with continuous variables. Audet et al. (1997) show that linear bilevel programs can be reformulated as mixed binary linear programs, and Fischetti et al. (2016) describe a general approach that focuses on adapting existing

MIP solvers to tackle bilevel problems. Saharidis and Ierapetritou (2009) present a Benders-like decomposition approach for linear problems, and Xu and Wang (2014) solve the problem with bounded integer leader variables. Anandalingam (1988) discusses a decentralized bilevel problem that has a single leader and multiple followers, which is one way of viewing the stochastic variant of a bilevel programming problem. Tahernejad et al. (2016) describe the implementation of the bilevel solution software MibS, which uses a branch-and-cut approach to solve mixed-integer linear bilevel programs. Several strategies for solving fully general bilevel MIPs have been developed, many of which are discussed by Fischetti et al. (2017) in the context of their software. Saharidis et al. (2013) survey solution approaches for linear bilevel programs. In particular, the authors discuss three categories of methods for mixed-integer problems: reformulation, branch-and-bound/branch-and-cut, and parametric programming. Kleinert et al. (2021) provide a more recent survey of solution algorithms, including a discussion of some nonlinear cases.

In Section 2, we describe the problem setup in detail, and derive a bilevel linear MIP formulation. In Section 3, we provide examples to illustrate the broad applicability of the proposed modeling environment. In Section 4, we present a numerical study to demonstrate the feasibility of this framework for realistically-sized problems using existing optimization software.

## 2 MDP Design: An Integrated Framework for Design and Operations

We consider a two-phase integrated decision framework, consisting of a design phase and an operational phase. In the design phase, an agent seeks to choose optimal design variables, modeled via a MIP, and during the operational phase, she determines optimal decision rules for an infinite-horizon MDP that depends on the initial design decisions. Uncertainty in the operational phase is modeled via probabilistic state transitions in an MDP, while a scenario approach is used to account for uncertainty in design execution between the two phases. Our objective is to minimize the combined expected cost of the MIP representing the design phase and the subsequent discounted-cost MDPs for the operational phase, taking into account both types of uncertainty. Dependence of the MDP on the design decisions is modeled through the MDP cost functions, which are assumed to be affine functions of the leader variables. We note that this dependence may be extended to instances where the MDP state and action spaces depend on the design decisions by means of a big-M approach that retains the linear relationship; see Section 3.2 for an example.

The design problem is represented by mixed-integer decision variables  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^{n_1} \times \mathbb{Z}^{n_2}$ , and parameterized by the cost vector  $\mathbf{c} \in \mathbb{R}^n$  (where  $n = n_1 + n_2$ ), constraint matrix  $\mathbf{A} \in \mathbb{R}^{m \times n}$ , and constraint right-hand-side  $\mathbf{b} \in \mathbb{R}^m$ . We use boldface letters (e.g.,  $\mathbf{x}$ ) to denote vectors, and subscripted regular typeface letters (e.g.,  $x_i$ ) to denote their components. Matrices are denoted by boldface upper case letters. To model uncertainty in the outcomes of the design phase, we consider a finite set of scenarios  $\mathcal{K}$ , where each scenario  $k \in \mathcal{K}$  occurs with probability  $q_k > 0$ . Thus,  $\mathbf{q} \in [0, 1]^{|\mathcal{K}|}$  and  $\sum_{k \in \mathcal{K}} q_k = 1$ .

For a fixed design decision  $\mathbf{x}$  and scenario  $k \in \mathcal{K}$ , the operational phase is modeled as an infinite-horizon discounted cost MDP with a (possibly scenario-dependent) discount factor  $\lambda_k \in (0, 1)$ , finite state space  $\mathcal{S}^k$ , and finite action space  $\mathcal{A}^k$ . For each state-action pair  $(s, a) \in \mathcal{S}^k \times \mathcal{A}^k$ , the agent incurs an immediate cost  $(\mathbf{f}_{s,a}^k)^\top \mathbf{x} + g_{s,a}^k$  that depends linearly on  $\mathbf{x}$ . Transition probabilities are denoted by  $\mathbf{p}^k(\cdot | s, a)$ .

A (stationary deterministic) decision rule  $\mathbf{d}^k$  is a function that assigns one action to every state in  $\mathcal{S}^k$ . Given a decision rule  $\mathbf{d}^k$ , let  $\mathbf{P}_{\mathbf{d}^k}$  be the associated transition probability matrix and  $\mathbf{h}_{\mathbf{d}^k}(\mathbf{x})$  be the cost vector whose  $s^{\text{th}}$  entry is the cost for state-action pair  $(s, d_s^k)$ . Then, the expected value of policy  $\mathbf{d}^k$  is

$$\mathbf{v}_{\mathbf{d}^k}(\mathbf{x}) = (\mathbf{I} - \lambda_k \mathbf{P}_{\mathbf{d}^k})^{-1} \mathbf{h}_{\mathbf{d}^k}(\mathbf{x}), \quad (1)$$

where  $\mathbf{I}$  denotes the identity matrix. The optimal value function of the MDP, denoted by  $\mathbf{v}^k(\mathbf{x})$ , is obtained by minimizing (1) componentwise over all possible decision-rules  $\mathbf{d}^k$  (or equivalently, given positive weights  $\beta$ , by minimizing  $\beta^\top \mathbf{v}_{\mathbf{d}^k}$  Puterman (2014)).

Standard methods for solving infinite-horizon MDPs include value iteration, policy iteration, and linear programming (LP) (Puterman, 2014). While the first two are more popular for solving stand-alone MDPs, we use the LP approach to formulate the integrated problem as a bilevel program whose second level is an LP. In fact, our approach can be generalized to other forms of operational decision-models that have a suitable linear programming representation, such as finite-horizon total cost MDPs Derman and Klein (1965) and finite-horizon risk-sensitive MDPs Kumar et al. (2015).

The LP method uses Bellman's equations to derive linear constraints on the optimal values of states, so that  $\mathbf{v}^k(\mathbf{x})$  is the optimal solution of  $(\text{LP}(\mathbf{x}, k))$ .

$$\begin{aligned} & \max_{\mathbf{v} \in \mathbb{R}^{|\mathcal{S}^k|}} \mathbf{1}^\top \mathbf{v} \\ & \text{s.t. } v_s \leq (\mathbf{f}_{s,a}^k)^\top \mathbf{x} + g_{s,a}^k + \lambda_k \sum_{s' \in \mathcal{S}^k} p^k(s' | s, a) v_{s'}, \quad \forall a \in \mathcal{A}^k, s \in \mathcal{S}^k, \end{aligned} \quad (\text{LP}(\mathbf{x}, k))$$

where  $\mathbf{1}$  is the vector of all ones. Let  $\alpha_s^k$  be the probability that the MDP in scenario  $k$  initially occupies state  $s \in \mathcal{S}^k$ . Then, given the value function  $\mathbf{v}^k(\mathbf{x})$  for each scenario  $k$ , the problem of finding the optimal design parameters is given by

$$\begin{aligned} \min \quad & \mathbf{c}^\top \mathbf{x} + \sum_{k \in \mathcal{K}} q^k ((\boldsymbol{\alpha}^k)^\top \mathbf{v}^k(\mathbf{x})) \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \in \mathcal{X}. \end{aligned} \tag{MIP}$$

The objective function of (MIP) is the sum of design costs and the expected *optimal* costs from the MDPs. Unlike the problem with adversarial recourse in Bailey et al. (2006), the function  $\mathbf{v}^k(\mathbf{x})$  is not convex; instead, it is piecewise linear concave Bard (1991). To see this, note from Equation (1) that  $\mathbf{v}^k(\mathbf{x})$  is the minimum of a family of linear functions of  $\mathbf{x}$  and therefore piecewise linear and concave Boyd and Vandenberghe (2004). In fact, it continues to be concave if the immediate costs of the MDP are concave in the design variables  $\mathbf{x}$ . If  $\mathcal{X}$  is a small finite set, then a brute-force approach to solving (MIP) consists of iterating over all feasible solutions and solving  $|\mathcal{K}|$  MDPs to optimality for each objective function evaluation. However, this approach is not feasible if  $\mathcal{X}$  is large or if the design phase includes continuous variables.

## 2.1 Bilevel Programming Formulation

To derive tractable reformulations of (MIP), we replace the functions  $\mathbf{v}^k(\mathbf{x})$  with auxiliary variables  $\mathbf{v}^k$ , and explicitly constrain these variables to equal the optimal values of the corresponding MDPs. Thus, we model (MIP) as a (mixed-integer) linear bilevel program with a single leader corresponding to the choice of the design variable  $\mathbf{x}$ , and  $|\mathcal{K}|$  follower problems, each comprised of the LP for the associated MDP. As such, the decision variables for the  $k^{\text{th}}$  follower problem are  $v_s^k$ , the value of state  $s$  in scenario  $k$ , for all states  $s \in \mathcal{S}^k$ . Therefore, a bilevel programming formulation for (MIP) is given by

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{v}^k} \quad & \mathbf{c}^\top \mathbf{x} + \sum_{k \in \mathcal{K}} q^k ((\boldsymbol{\alpha}^k)^\top \mathbf{v}^k) \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \in \mathcal{X}, \\ & \mathbf{v}^k \in \operatorname{argmax}_{\mathbf{v} \in \mathbb{R}^{|\mathcal{S}^k|}} \left\{ \mathbf{1}^\top \mathbf{v} : v_s \leq (\mathbf{f}_{s,a}^k)^\top \mathbf{x} + g_{s,a}^k + \lambda_k \sum_{s' \in \mathcal{S}^k} p^k(s'|s, a) v_{s'}, \forall a \in \mathcal{A}^k, s \in \mathcal{S}^k \right\}, \quad \forall k \in \mathcal{K}. \end{aligned} \tag{2}$$

Problem (2) is a linear bilevel program with a mixed-integer leader problem and  $|\mathcal{K}|$  independent continuous follower problems—an equivalent formulation with a single follower is obtained by moving the expectation from the leader’s objective to the follower’s objective and using a single follower variable  $\mathbf{v} = \sum_{k \in \mathcal{K}} q^k ((\boldsymbol{\alpha}^k)^\top \mathbf{v}^k)$ . Because the MDPs have unique optimal value functions Puterman (2014), the question of optimistic versus pessimistic formulations of the bilevel problem (Colson et al., 2005) do not arise even though the followers’ optimal policies may not be unique. Linear bilevel optimization is strongly NP-hard Hansen et al. (1992), but recent years have seen significant advances in the development of tractable solution methods for mixed-integer linear bilevel programs Kleinert et al. (2021). As such, (2) can be solved using existing bilevel solvers.

## 2.2 MIP Formulation

To obviate the need for specialized software or solution approaches for bilevel programs, we may also reformulate (2) as a MIP that uses the dual of the follower problem to obtain a single-level MIP. The dual for  $(\text{LP}(\mathbf{x}, k))$  is given by

$$\begin{aligned}
& \min \sum_{s \in \mathcal{S}^k} \sum_{a \in \mathcal{A}^k} \gamma_{s,a}^k \cdot ((\mathbf{f}_{s,a}^k)^\top \mathbf{x} + g_{s,a}^k) \\
& \text{s.t.} \quad \sum_{a \in \mathcal{A}^k} \left[ \gamma_{s,a}^k - \lambda \sum_{s' \in \mathcal{S}^k} \gamma_{s',a}^k p^k(s | s', a) \right] = 1 \quad \forall s \in \mathcal{S}^k, \\
& \quad \gamma_{s,a}^k \geq 0 \quad \forall s \in \mathcal{S}^k, a \in \mathcal{A}^k.
\end{aligned} \tag{3}$$

Then, a nonlinear MIP reformulation of (2) is obtained by substituting the dual problems (3) for the follower problems, and including complementary slackness conditions as (nonlinear) constraints Fortuny-Amat and McCarl (1981). Moreover, a big-M approach may then be used to derive a linear

MIP reformulation as follows Fortuny-Amat and McCarl (1981):

$$\begin{aligned}
& \min \mathbf{c}^\top \mathbf{x} + \sum_{k \in \mathcal{K}} q^k \sum_{s \in \mathcal{S}^k} \alpha_s^k v_s^k \\
& \text{s.t. } \mathbf{Ax} = \mathbf{b}, \\
& \mathbf{x} \in \mathcal{X}, \\
& v_s^k \leq (\mathbf{f}_{s,a}^k)^\top \mathbf{x} + g_{s,a}^k + \lambda_k \sum_{j \in \mathcal{S}} p^k(j|s,a) v_j^k \quad \forall s \in \mathcal{S}^k, a \in \mathcal{A}^k, k \in \mathcal{K}, \\
& \sum_{a \in \mathcal{A}^k} [\gamma_{s,a,k} - \lambda_k \sum_{j \in \mathcal{S}^k} \gamma_{a,j,k} p^k(s|j,a)] = q^k \alpha_s^k \quad \forall s \in \mathcal{S}^k, k \in \mathcal{K}, \\
& \gamma_{s,a,k} \leq M_k \delta_{s,a,k} \quad \forall s \in \mathcal{S}^k, a \in \mathcal{A}^k, k \in \mathcal{K}, \\
& (\mathbf{f}_{s,a}^k)^\top \mathbf{x} + g_{s,a}^k - v_s^k + \lambda_k \sum_{j \in \mathcal{S}} p^k(j|s,a) v_j^k \leq M'_k (1 - \delta_{s,a,k}) \quad \forall s \in \mathcal{S}^k, a \in \mathcal{A}^k, k \in \mathcal{K}, \\
& \gamma_{s,a,k} \geq 0, \delta_{s,a,k} \in \{0, 1\} \quad \forall s \in \mathcal{S}^k, a \in \mathcal{A}^k, k \in \mathcal{K}.
\end{aligned} \tag{4}$$

The values of  $M_k$  and  $M'_k$  can be chosen as  $M_{s,k}$ , the maximum number of (discounted) times state  $s$  could be reached in scenario  $k$ , and  $M'_{s,k}$ , the most costly possible outcome starting in state  $s$  in scenario  $k$  for any feasible  $\mathbf{x}$ , respectively. The latter is found by solving an LP for each state-scenario pair. More simply, we may choose  $M_k = 1/(1 - \lambda_k)$  and  $M'_k = t^k/(1 - \lambda_k)$ , where  $t^k$  is the maximum cost for any state-action pair in scenario  $k$ .

The MIP reformulation allows us to solve the integrated problem (2) using readily available off-the-shelf MIP solvers. Moreover, as shown in Section 3.2, this approach can provide additional modeling flexibility in cases where the state and action spaces for the follower depend on the design decision. However, Kleinert and Schmidt (2023) show that this approach is not always computationally advantageous, in part due to the challenge in obtaining suitable big- $M$  parameters. Even with the choice of parameters described above, our preliminary computational tests indicated poorer computational performance for the MIP approach than for directly solving the bilevel model using specialized bilevel optimization software. Therefore, we focus on the bilevel model in our numerical experiments in Section 4.

### 3 Applications

We now illustrate that this problem setting arises in several application areas with interdependent design and operational phases.



### 3.1 Reliability

In designing a system for long-term operation, minimizing downtime and maintenance costs is an important consideration which must be balanced against the need to lower initial costs. As such, there is a tradeoff between, for example, selecting expensive high-quality components to build the system, or using low-cost components that lead to increased system downtime and maintenance requirements during regular operation. An integrated model would make optimal design choices to balance this tradeoff.

Tao et al. (1995) consider this problem and propose two solution approaches. The first aims to iteratively solve the MDP for a range of design decisions, which cannot be guaranteed to recover the true optimal unless the number of design choices is finite. The second reformulates the problem as a single nonlinear non-convex problem without structure that would lend itself to tractability. We discuss this problem in the context of our framework. Specifically, we use MDPs to model system reliability and extend the standard approach to additionally consider the problem of purchasing the best components from a set of options under budgetary constraints.

Let  $T$  be a system dependent on some subset of  $N$  components arranged in series, and let  $B$  be a boolean expression for whether  $T$  is operational based on the current present and working components. This can easily be expanded to a network of agents, each of which has its own components and may or may not be functioning at a given time, with an arbitrary degree of dependence on the other agents to determine the transition probabilities; for the purposes of this discussion, we use a single-system model.

We model the design problem via mixed-integer decisions  $\mathbf{x} \in \mathbb{R}^{|N|}$  for the amount of each component to purchase (where a continuous decision variable may indicate a volume or mass of some component) at costs  $\mathbf{c}$ , such that system operation is possible and the purchase cost does not exceed a budget  $b$ . Let  $\mathcal{X}$  be the set of component choices that allow for system operation. The uncertainty in design outcomes represents the variability in the quality of the components that are eventually delivered, as well as external sources of uncertainty, such as demand for the service provided and the labor market for performing repairs.

In the subsequent operational stage, we define status variables  $\mathbf{y} \in \mathbb{R}^{|N|}$  to indicate the state of each component selected in the first stage, and the system state  $s$  is the product of these status variables. Transition probabilities  $p_i^k(y_i | s, a)$  for each component  $i$  depend on both the current state and the design uncertainty. In any epoch, an action  $a$  corresponds a to subset of the components to

repair, and the immediate cost  $(\mathbf{f}_{s,a}^k)^\top \mathbf{x} + g_{s,a}^k$  includes both the cost of repair and system operation.

We assume that the system is initially in the fully functioning state denoted by  $s_0$ . Then the bilevel problem takes the form:

$$\begin{aligned} \min \quad & \mathbf{c}^\top \mathbf{x} + \sum_{k \in \mathcal{K}} q^k v_{s_0}(\mathbf{x}, k) \\ \text{s.t.} \quad & \mathbf{c}^\top \mathbf{x} \leq b, \\ & \mathbf{x} \in \mathcal{X}, \\ & \mathbf{v}(\mathbf{x}, k) \in \operatorname{argmax} (\operatorname{LP}(\mathbf{x}, k)). \end{aligned}$$

### 3.2 Inventory management

Consider an extended version of the inventory management problem, where the first phase corresponds to business startup location decisions in a particular market, and the second phase is a classic inventory management problem with Markovian demand (see White (1993) for a survey). We assume that the costs of ordering and holding inventory are location-dependent. The discount factor accounts for rental or other location dependent costs that are fixed for each location decision and are incurred at each time step.

For the design phase, suppose the agent needs to select  $r$  locations out of  $R$  available choices. Let  $\mathbf{x} \in \mathbb{B}^R$  correspond to binary decisions indicating whether each location was selected, with associated startup costs  $\mathbf{c}$ . Let  $\mathbf{m}$  be the inventory capacities for all locations. We assume that each location can initially order up to a maximum of their capacity. Let  $\mathbf{u}$  be the initial inventory ordered, with per unit costs  $\mathbf{b}$ . Then the objective for the design-only problem would be

$$\min_{\mathbf{x} \in \mathbb{B}^R, \mathbf{u} \in \mathbb{R}^R} \mathbf{c}^\top \mathbf{x} + \mathbf{b}^\top \mathbf{u},$$

with constraints

$$\mathbf{1}^\top \mathbf{x} = r \quad \text{and} \quad 0 \leq u_i \leq m_i x_i \quad \forall i \in \{1, \dots, R\}.$$

Design uncertainty may reflect variability in the economic environment and market demand by the time of completion. Once the locations are ready for use, the operational phase corresponds to inventory management at each location, which is modeled as an MDP; see (Puterman, 2014, Section 3.2) for an example. Let  $\mathbf{f}^k$  be the vector of fixed costs incurred at each location in every period. Given an inventory level  $s$  and order quantity  $a$ , let  $\mathbf{o}_{s,a}^k(\mathbf{x})$  be the ordering costs,  $\mathbf{h}_{s,a}^k(\mathbf{x})$  be the holding costs, and  $\mathbf{r}_{s,a}^k(\mathbf{x})$  be the total sale revenues. We assume that demand is not backlogged,

but lost demand incurs shortage costs  $\mathbf{p}_{s,a}^k(\mathbf{x})$ . We further assume that the costs and revenues are linear in  $\mathbf{x}$ . To see why this assumption is justified, fix  $i \in \{1, \dots, R\}$  and note that  $x_i$  indicates if the location is operational. Then, if the per unit ordering cost is  $o_i^k$ , the total ordering cost is given by

$$o_i^k x_i a_i + M_i^k (1 - x_i) a_i = (o_i^k a_i - M_i^k a_i) x_i + M_i^k a_i,$$

where  $M_i$  is a penalty parameter that enforces  $a_i = 0$  if the  $i$ -th location has not been selected in the design phase. Thus, the ordering cost is linear in the leader variable. Similar arguments apply to the other cost functions as well. However, there is an added subtlety here because the state and action spaces depend on  $\mathbf{x}$ . Specifically, the inventory at the  $i$ -th location in any period must be less than  $m_i x_i$ , and the order quantities in state  $s$  may not exceed  $m_i x_i - s$ . We present two ways of modeling these constraints within the bilevel framework.

1. Penalty term in cost function: The immediate cost function of the MDP may include additional terms of the form

$$M_1(m_i x_i - s_i)^+ \text{ and } M_2(m_i x_i - s - a)^+,$$

where  $M_1$  and  $M_2$  are (large) penalty parameters. In this case, however, the cost functions will be piecewise linear and convex in  $\mathbf{x}$ , which makes it more challenging to solve the bilevel problem.

2. Auxiliary variables in the dual problem (3): This is a similar idea as above but utilizes the dual formulation for the MDP. We define auxiliary variables  $\rho_s = (m_i x_i - s_i)^+$  and  $\mu_{s,a} = (m_i x_i - s - a)^+$  and add terms  $M_1 \rho_s$  and  $M_2 \mu_{s,a}$  to the objective function of (3). The definitions of  $\rho_s$  and  $\mu_{s,a}$  may be enforced via linear constraints. This results in a larger dual problem that continues to be linear. Thus, this is an instance where reformulating the bilevel problem as a single-level MIP may be advantageous.

Dependence of the operational phase on initial inventory  $\mathbf{u}$  may be modeled similarly. Thus, our framework may be applied here.

### 3.3 Queue Design and Control

Our third example considers the problem of optimally designing a queueing system. Several authors have studied the optimal selection of queue parameters, such as arrival rates and service times,

where performance measures of the system are expressible as functions of these parameters in closed form; see Tadj and Choudhury (2005) for a survey. Our framework allows for greater flexibility by explicitly modeling the operational phase as an MDP.

In the design phase, an agent seeks to determine the number and types of servers to be installed. Given  $n$  server types, let  $\mathbf{x} \in \mathbb{Z}_+^n$  denote the number of servers of each type that are selected. Let  $\mathbf{r}$  be the maximum allowable number of each type of server, and  $t$  be the maximum total number of servers. Then, if the one-time recruitment and training costs per server are given by  $\mathbf{c}$ , the design problem takes the form

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{Z}_+^n} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \mathbf{x} \leq \mathbf{r} \\ & \mathbf{1}^\top \mathbf{x} \leq t. \end{aligned}$$

The integrated framework additionally considers the uncertainty in the outcome of the design phase and the expected long-run operational cost. Design uncertainty in this setting may correspond to the realized service rates of each server type, as well as the servers' ability to successfully resolve the needs of various customer types. In the operational phase,  $J$  types of customers arrive according to some time-discretized arrival process with rate  $\eta_j$ ,  $j = 1, \dots, J$ . The state  $s$  of the MDP represents the number of customers of each type that are present in the queue. We assume that the system has a finite capacity, so that the state space is finite. However, a penalty  $p$  is incurred when the system is at capacity and customers are turned away.

In every decision epoch, the action  $a$  determines the number of customers of each type that are assigned to each server-type. Let  $p_{ij}$  be the probability with which a server of type  $i$  successfully serves a customer of type  $j$ , earning a reward  $r_{ij}$ . We assume that customers exit the system after service regardless of whether service was successful. Let  $o_i$  be the operating cost per period of a type- $i$  server. As in Section 3.2, we may express the costs and rewards as linear functions of the design variables. Then, the agent seeks to determine an assignment strategy that minimizes the total discounted cost less the expected reward. As such, the long-run cost of operating the queueing system can be included in the design decisions using our bilevel framework.

$n$	$m$	$ \mathcal{K} $	$ \mathcal{S} $	$ \mathcal{A} $	Avg. solve time (s)
<b>20</b>	40	20	10	20	3.384
<b>40</b>	40	20	10	20	5.060
<b>80</b>	40	20	10	20	7.732
<b>160</b>	40	20	10	20	12.584
<b>320</b>	40	20	10	20	34.005
80	<b>10</b>	20	10	20	7.431
80	<b>20</b>	20	10	20	7.464
80	<b>40</b>	20	10	20	7.439
80	<b>80</b>	20	10	20	7.640
80	<b>160</b>	20	10	20	7.664
80	40	<b>5</b>	10	20	1.323
80	40	<b>10</b>	10	20	2.849
80	40	<b>20</b>	10	20	7.239
80	40	<b>40</b>	10	20	18.315
80	40	<b>80</b>	10	20	50.925
80	40	20	<b>2</b>	20	0.572
80	40	20	<b>4</b>	20	1.453
80	40	20	<b>8</b>	20	4.620
80	40	20	<b>16</b>	20	25.932
80	40	20	<b>32</b>	20	237.272
80	40	20	10	<b>5</b>	0.772
80	40	20	10	<b>10</b>	2.524
80	40	20	10	<b>20</b>	8.005
80	40	20	10	<b>40</b>	29.119
80	40	20	10	<b>80</b>	95.440

**Table 1:** Average solve times (in seconds) over 100 instances for 25 different combinations of parameter values. The  $n$  leader variables are evenly split between binary variables and general integer variables.

Parameter	Mean	Variance
Leader integer variable upper bound	12	1
Leader constraint coefficients	$\bar{a} = 12$	$\sigma = 0.5$
Leader constraint upper bound	$\bar{b} = \frac{(n_1 + n_2)}{\bar{a} + 2\sigma}$	$\bar{b}/6$
Scenario relative probability	1	0.2
Initial state relative probability	1	0.2
Transition relative probability	1	0.4

**Table 2:** Parameters of the normal distributions used for instance generation in the numerical experiments. Discount factors were sampled uniformly at random over the interval  $[0.92, 0.97]$ , leader objective coefficients over  $[10, 100]$ , cost function coefficients over  $[-1, 1]$ , and cost function constant terms over  $[10, 40]$ .


## 4 Numerical Results

In this section, we illustrate how existing bilevel optimization algorithms may be used to solve realistic instances of the bilevel programs that arise in our framework. We present results for a general form of the problem, analyzing the performance of the bilevel MIP solver of Fischetti et al. (2017) (available at <https://msinnl.github.io/pages/bilevel.html>) to solve random instances of (2) for a range of parameter values. For each set of parameter values in Table 1, we generated and solved 100 test problems. The cost and constraint coefficients and other parameters for these problems were randomly generated given fixed mean values and variances that are listed in Table 2. Relative probabilities were randomly generated as noted in Table 2 and then normalized to sum to 1. We assumed that all leader variables are integer or binary and the number of integer and binary variables remained the same in each case.

The average solve times for each combination of parameter values are tabulated in Table 1. Trends in average solve time are illustrated in Figure 1. The number of leader constraints appear to have little effect on the average solve time. However, average solve time increase with the number of MDP states, number of MDP actions per state, number of scenarios, and number of leader variables. Preliminary tests also showed that varying the mean values used to generate right-hand sides for the leader constraints does not significantly affect average solve times.

## 5 Conclusion

We presented a modeling framework to integrate strategic and operational decisions made by the same agent. Given a static design phase modeled as a MIP and a dynamic operational phase rep-



**Figure 1:** Trends in average solve time (in seconds) over 100 instances upon varying the number of leader variables, leader constraints, scenarios, MDP states and MDP actions per state.

represented by an infinite-horizon MDP, we derived a bilevel optimization formulation that captures the temporal hierarchy of these decisions. The bilevel program, which has a mixed-integer linear leader problem and continuous linear follower problems, can directly be solved using existing computational methods. We provided three examples to illustrate how the modeling approach may be applied in practice, and presented numerical results to illustrate that realistic problem instances can be solved using an existing bilevel solver. Future work will explore tailored solution approaches that leverage the MDP structure and can lead to improved computational performance.

## Acknowledgments

This research was supported by the National Science Foundation grant CMMI 1826323 and the National Cancer Institute grant 5R01CA257814-02. The authors thank Tyler Perini of Rice University for his helpful comments.

## References

- G. Anandalingam. A mathematical programming model of decentralized multi-level systems. *Journal of the Operational Research Society*, 39(11):1021–1033, 1988.
- C. Audet, P. Hansen, B. Jaumard, and G. Savard. Links between linear bilevel and mixed 0-1 programming problems. *Journal of Optimization Theory and Applications*, 93(2):273–300, 1997.

- Styliani Avraamidou and Efstratios N Pistikopoulos. A bi-level formulation and solution method for the integration of process design and scheduling. In *Computer Aided Chemical Engineering*, volume 47, pages 17–22. 2019.
- Oluwadare Badejo and Marianthi Ierapetritou. Integrating tactical planning, operational planning and scheduling using data-driven feasibility analysis. *Computers & Chemical Engineering*, 161:107759, 2022.
- M. D. Bailey, S. M. Shechter, and A. J. Schaefer. SPAR: stochastic programming with adversarial recourse. *Operations Research Letters*, 34(3):307–315, 2006.
- J. F. Bard. Some properties of the bilevel programming problem. *Journal of Optimization Theory and Applications*, 68(2):371–378, 1991.
- Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge University Press, 2004.
- Baris Burnak, Nikolaos A Diangelakis, Justin Katz, and Efstratios N Pistikopoulos. Integrated process design, scheduling, and control using multiparametric programming. *Computers & Chemical Engineering*, 125:164–184, 2019.
- B. Colson, P. Marcotte, and G. Savard. Bilevel programming: A survey. *4OR*, 3(2):87–107, 2005.
- Stephan Dempe, Vyacheslav Kalashnikov, Gerardo A. Pérez-Valdés, and Nataliya Kalashnykova. *Bilevel Programming Problems: Theory, Algorithms and Applications to Energy Networks*. Energy Systems. Springer, 2015.
- C. Derman and M. Klein. Some remarks on finite horizon Markovian decision models. *Operations Research*, 13(2):272–278, 1965.
- M. Fischetti, I. Ljubić, M. Monaci, and M. Sinnl. Intersection cuts for bilevel optimization. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 77–88, 2016.
- M. Fischetti, I. Ljubić, M. Monaci, and M. Sinnl. A new general-purpose algorithm for mixed-integer bilevel linear programs. *Operations Research*, 65(6):1615–1637, 2017.
- José Fortuny-Amat and Bruce McCarl. A representation and economic interpretation of a two-level programming problem. *Journal of the Operational Research Society*, 32(9):783–792, 1981.



- Sebastián González, Felipe Feijoo, Franco Basso, Vignesh Subramanian, Sriram Sankaranarayanan, and Tapas K Das. Routing and charging facility location for EVs under nodal pricing of electricity: A bilevel model solved using special ordered set. *IEEE Transactions on Smart Grid*, 13(4):3059–3068, 2022.
- Ignacio Grossmann. Enterprise-wide optimization: A new frontier in process systems engineering. *AIChE Journal*, 51(7):1846–1857, 2005.
- Pierre Hansen, Brigitte Jaumard, and Gilles Savard. New branch-and-bound rules for linear bilevel programming. *SIAM Journal on scientific and Statistical Computing*, 13(5):1194–1217, 1992.
- Thomas Kleinert and Martin Schmidt. Why there is no need to use a big-M in linear bilevel optimization: A computational study of two ready-to-use approaches. *Computational Management Science*, 20(1):3, 2023.
- Thomas Kleinert, Martine Labbé, Ivana Ljubić, and Martin Schmidt. A survey on mixed-integer programming techniques in bilevel optimization. *EURO Journal on Computational Optimization*, 9:100007, 2021.
- Atul Kumar, Veeraruna Kavitha, and Nandyala Hemachandra. Finite horizon risk sensitive MDP and linear programming. In *2015 54th IEEE Conference on Decision and Control*, pages 7826–7831. IEEE, 2015.
- Ludger Leenders, Dörthe Franzisca Hagedorn, Hatim Djelassi, André Bardow, and Alexander Mitsos. Bilevel optimization for joint scheduling of production and energy systems. *Optimization and Engineering*, 24(1):499–537, 2023.
- M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.
- G. K. Saharidis and M. G. Ierapetritou. Resolution method for mixed integer bi-level linear problems based on decomposition technique. *Journal of Global Optimization*, 44(1):29–51, 2009.
- G. K. D. Saharidis, A. J. Conejo, and G. Kozanidis. Exact solution methodologies for linear and (mixed) integer bilevel programming. In *Metaheuristics for Bi-level Optimization*, pages 221–245. 2013.
- L. Tadj and G. Choudhury. Optimal design and control of queues. *Top*, 13(2):359–412, 2005.

- S. Tahernejad, T. K. Ralphs, and S. T. DeNegre. A branch-and-cut algorithm for mixed integer bilevel linear optimization problems and its implementation. *COR@L Technical Report 16T-015-R3, Lehigh University*, 2016.
- Z. Tao, R. B. Corotis, and J. H. Ellis. Reliability-based structural design with Markov decision processes. *Journal of Structural Engineering*, 121(6):971–980, 1995.
- Rafael D. Tordecilla, Jairo R. Montoya-Torres, Carlos L. Quintero-Araujo, Javier Panadero, and Angel A. Juan. The location routing problem with facility sizing decisions. *International Transactions in Operational Research*, 30(2):915–945, 2023.
- D. J. White. A survey of applications of Markov decision processes. *Journal of the Operational Research Society*, 44(11):1073–1096, 1993.
- P. Xu and L. Wang. An exact algorithm for the bilevel mixed integer linear programming problem under three simplifying assumptions. *Computers & Operations Research*, 41:309–318, 2014.