

Multi-Objective Optimal Control with Safety as a Priority

Kendra Lesser
Dept. of Computer Science
University of Oxford
kendra.lesser@cs.ox.ac.uk

Alessandro Abate
Dept. of Computer Science
University of Oxford
aabate@cs.ox.ac.uk

ABSTRACT

This work develops a lexicographic approach to multi-objective optimal control on models for cyber-physical systems, encompassing in particular stochasticity, limited access to model variables (partial observations), and possibly hybrid (continuous and discrete) dynamics (with the finite state POMDP framework as a known special instance). The technique is showcased in two new case studies in the area of buildings automation systems. Technically, the main achievements of this work are: The application of the lexicographic framework to multi-objective optimization including quantitative probabilistic safety requirements, thus leading to a principled and scalable integration of correct-by-design synthesis for safety and optimal synthesis for performance; the novel extension of the lexicographic framework to partially-observed stochastic models with continuous (possibly hybrid) dynamics; the emphasis on computational aspects, including the use of compact and approximate representations of value functions combined with the quantification of error bounds on model abstractions.

1. INTRODUCTION

Cyber-physical systems (CPS) are frequently safety-critical, and hence require stringent guarantees that the embedded control processes governing their operation do not lead to unsafe configurations. Smart-energy buildings, infrastructure networks, and electricity grids are all examples of CPS that integrate digital control signals with underlying continuous physical processes, and for which failure to operate safely would be prohibitively costly. Rather than checking if safety specifications are met *after* a controller is designed (known as validation), one way to achieve safety guarantees is to enforce safety requirements during the control design process (called “correct-by-design” control), either by treating safety as a constraint that must be satisfied or by reducing it as the objective of an optimal control problem.

Because of their intrinsic complexity, CPS models are often subject to uncertainties, either due to inaccuracies in

the model or to environmental influences (weather-related, delays, or presence of humans in the loop) that cannot be predicted exactly. CPS are further characterised by information exchange to the controllers about the system, often done through sensors that a) may not be ubiquitous, and cannot measure all necessary information, and b) may introduce noise, and hence inaccuracies, into their measurements. *Partially observable* stochastic models, such as finite state partially observable Markov decision processes (POMDPs) or partially observable stochastic hybrid systems (POSHS), with both discretely and continuously evolving states, provide a fitting and realistic modelling framework for CPS.

The inclusion of stochasticity leads to regarding safety probabilistically - it is difficult, when not overtly conservative, to guarantee it with certainty, that is with probability equal to one. If we want to *design* a safe controller, we could synthesise a controller that maximises the probability of safety, as done in [2] for fully observable stochastic hybrid systems, and in [8, 13, 14] for POSHS. The major drawback to this approach is that safety is rarely the *only* concern when designing controllers for CPS: stabilisation, tracking, or costs minimisation are just some of other possible objectives of interest. We are thus led to a heterogeneous *multi-objective* optimisation setup, where we not only want to design a controller that maximises the safety of the system, but also takes into account these additional performance-related objectives. In order to tackle this synthesis goal, one option is to synthesise a permissive controller, which provides a set of control inputs that satisfy a given requirement (i.e. safety), so that some other objective can be optimised over the set of available inputs. This is done for finite state Markov decision processes (MDPs) in [9], and for partially observable (but otherwise deterministic) discrete event systems, in [12, 24], although none actually consider a second control objective. Another possibility is to treat safety as a constraint to be enforced in the optimisation of a different objective, as in [21], which specifically considers enforcement with probability one, and only for finite state MDPs. Multi-objective optimisation of stochastic games is considered in [6] for complex objectives, such as safety, although partial observability is not considered.

None of the above multi-objective approaches are easily extended to general partially-observable stochastic systems, which are more generally rarely considered in the context of verification and formal synthesis (notable exceptions are: [25] introducing verification of hidden Markov models, [11] considering synthesis over POMDPs for qualitative

(yes/no) specifications, and [5] discussing decidability results for POMDPs).

Over partially observable models, we would further like the possibility to explore the trade-off between safety and other objectives, which is not possible when we require safety to be enforced with probability one. We therefore take a lexicographic optimisation approach, and apply it to the formulation of safety as a unique specification presented for POSHS in [8,13]. Lexicographic optimisation utilises multiple objective functions and assigns a preference (or priority) to each. A set of control inputs is generated that produces outcomes within a tolerance from the optimal solution, and that subset is then used as the set of possible inputs over which the next objective is optimised, thus taking a hierarchical approach to controller synthesis. This technique was first introduced in [15], and more recently elaborated upon for MDPs in [23] and for POMDPs in [22]. In [22], however, results are only presented for finite state models, and some details are overlooked, so that the results are not applicable to complex objectives such as safety.

Our contributions are several. First, we extend the lexicographic optimisation framework to multi-objective optimisation with safety requirements, thus integrating correct-by-design synthesis and synthesis for optimal performance. Second, we apply the lexicographic framework to partially observable stochastic systems, with an extension to models with continuous or hybrid dynamics. Third, we provide an approximate computational technique to solve the lexicographic optimisation problem based on point-based value iteration (PBVI) [17], and quantify error bounds for this approximation, as well as error bounds on the abstraction necessary to reduce stochastic continuous or hybrid systems to finite state POMDPs. Finally, throughout the article we apply our theoretical work on two case studies in the area of buildings automation systems.

2. BACKGROUND

2.1 POMDP Model

We consider systems that can be modelled as partially observable Markov decision processes (POMDPs). A standard POMDP has discrete states, actions, and observations, and is studied over an additive cost function.

DEFINITION 1. A POMDP is a tuple $\mathcal{J} = (\mathcal{S}, \mathcal{U}, \mathcal{Y}, T, Y, R)$, where

1. \mathcal{S} is a finite set of states
2. \mathcal{U} is a finite set of possible control inputs
3. \mathcal{Y} is a finite set of observations
4. $T : \mathcal{S} \times \mathcal{S} \times \mathcal{U} \rightarrow [0, 1]$ is a state transition function that assigns a probability measure to state s_{n+1} given state s_n and control u_n for all n : $T(s_{n+1}|s_n, u_n)$
5. $Y : \mathcal{Y} \times \mathcal{S} \times \mathcal{U} \rightarrow [0, 1]$ is an observation function that assigns a probability measure to observation y_n given state s_n and control u_{n-1} for all n : $Y(y_n|s_n, u_{n-1})$
6. $R : \mathcal{S} \rightarrow [0, 1]$ is an initial probability measure over the state space \mathcal{S} : $R(s)$

In this model the state is in general not accessed, and is therefore not available for control. Instead, a controller has access to all past observations and control inputs, which are stored in an information vector $i_n = (u_0, \dots, u_{n-1}, y_1, \dots, y_n) \in \mathcal{I}_n = \mathcal{U}^n \times \mathcal{Y}^n$. A control input is thus selected via a policy π , which maps the available information to the set of possible controls – more precisely:

DEFINITION 2. A policy π for a POMDP \mathcal{J} over a time horizon N is a sequence of functions, $\pi = (\pi_0, \dots, \pi_{N-1})$, such that $\pi_n : \mathcal{I}_n \rightarrow \mathcal{U}$.

We consider only non-randomised policies (each input i_n produces a single output u_n) and the set of all possible non-randomised policies is denoted by Π . Typically, a control policy is obtained as the argument that minimises an expected sum of costs associated with the state and control input, namely

$$\pi^* = \arg \min_{\pi \in \Pi} \mathbb{E}^\pi \left[\sum_{n=0}^N C(s_n, u_n) \middle| R \right]. \quad (1)$$

The execution of a POMDP proceeds as follows. At time $n = 0$, state s_0 is produced from the initial distribution $R : s_0 \sim R(\cdot)$. At each subsequent time step $n > 0$, the state is s_n . An observation y_n is produced according to $y_n \sim Y(\cdot|s_n, u_{n-1})$. This observation is added to the list of past observations and control inputs to produce $i_n = (u_0, \dots, u_{n-1}, y_1, \dots, y_n)$. A new control input is chosen according to $u_n = \pi_n(i_n)$ (obtained solving (1)), and cost $C(s_n, u_n)$ is accrued. At the final time step, when no control input is needed, the cost function $C(s_N, u_N)$ is modified to $C(s_N, 0)$.

2.1.1 Case Study: Optimised Boiler Maintenance

To illustrate how a POMDP can model a physical system, consider a simplified maintenance problem for a single boiler that heats a building¹. The boiler operates with a certain level of efficiency that is captured through a degradation model. The degradation level slowly increases in time as a result of normal use, which we model using a stochastic difference equation as

$$s_{n+1} = s_n + v_n, \quad (2)$$

with s_n denoting the degradation level at discrete time step n , and v_n a random variable taking only non-negative values from the set $\mathcal{V}(s_n)$, i.e. $v_n \in \mathcal{V}(s_n)$ for all n , where the set is a function of the current state s_n . The probability that $v_n = \bar{v} \in \mathcal{V}(\bar{s})$ is defined through the probability mass function P_v , so that $\mathbb{P}(v_n = \bar{v}|s_n = \bar{s}) = P_v(\bar{v}, \bar{s})$. The degradation level is restricted to the finite set $s_n \in \{0, \dots, 100\} = \mathcal{S}$: at $s_n = 0$ the boiler is perfectly clean and operates efficiently, whereas at $s_n = 100$ the boiler cannot operate (e.g., it is fully clogged). The dependence of the domain of \mathcal{V} on s_n ensures that s_{n+1} does not leave \mathcal{S} , e.g. if $s_n = 99$, $\mathcal{V}(99) = \{0, 1\}$.

The degradation level is determined by a number of factors, and can be estimated as the ratio of power demanded (as a function of the desired temperature) to the power

¹Optimal predictive maintenance is well motivated. Maintenance costs contribute between 15 and 60 percent of the cost of goods produced [16], and an efficient maintenance policy can reduce operational costs by 5 to 40 percent [1]. Although optimal maintenance is often treated as a simple case study in dynamic control texts [4], [10], there are few specialized publications [3].

output of the boiler. These quantities are not known exactly, but rather may be measured by possibly noisy sensors. We model this uncertainty through an *observed* degradation level,

$$y_n = s_n + w_n, \quad (3)$$

where the true degradation level s_n is corrupted by a noise term $w_n \in \mathcal{W}(s_n)$, where the noise is stochastic and takes values within the set $\mathcal{W}(s_n)$, again dependent on the actual degradation level. The probability that $w_n = \bar{w}$ given $s_n = \bar{s}$ is defined by the probability mass function P_w , namely $\mathbb{P}(w_n = \bar{w} | s_n = \bar{s}) = P_w(\bar{w}, \bar{s})$. Because we know that $s_n \in \{0, 100\}$, the observations are also restricted to the set $\mathcal{Y} = \{0, \dots, 100\}$.

The control actions available at each time step are twofold: either to clean the boiler, which resets its degradation level, or to do nothing. In the first instance, the difference equation (2) for a single time step changes to

$$s_{n+1} = z_n, \quad (4)$$

with z_n a random variable taking values in $\mathcal{Z} = \{0, \dots, 5\}$ with probabilities given by a distribution P_z .

We can therefore model the dynamical system (2), (3), and (4) as a POMDP with state space \mathcal{S} , observation space \mathcal{Y} , and control space $\mathcal{U} = \{\text{Clean}, \text{NoClean}\}$. The transition function T is derived from (2) and P_v , so that

$$T(s'|s, u) = \begin{cases} P_v(s' - s, s), & \text{if } s' - s \in \mathcal{V}(s), u = \text{NoClean} \\ P_z(s'), & \text{if } s' \in \mathcal{Z}, u = \text{Clean} \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

The observation function is derived from (3) and P_w as

$$Y(y|s) = \begin{cases} P_w(y - s, s), & \text{if } y - s \in \mathcal{W}(s) \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

We will assume that the initial distribution R is concentrated at a single point s^0 , i.e. $R(s^0) = \mathbb{P}(s_0 = s^0) = 1$, although any well-defined probability mass function over \mathcal{S} is acceptable.

2.2 Optimal Control of POMDPs

The optimal policy that minimises an expected sum of costs, as in (1), can be synthesised via dynamic programming, much like for a Markov decision process (MDP), in which the state s_n is known completely [4]. A value function over the fully observed state, which represents the expected sum of costs accrued from time n to time N given the state at time n , is set and optimised backward-recursively for each time step n . The optimal policy for a POMDP is found by redefining the state space as something that is fully observable, and thus reformulating the POMDP as an MDP [4].

One way to redefine the POMDP as an MDP is to treat the information vector i_n as the fully observed state of the system. The information vector, however, increases in size with n , and can thus be difficult to store. A common alternative is to instead use a belief state, which is a sufficient statistic for the information vector, and therefore condenses the information stored in i_n without sacrificing the ability to construct optimal policies. For an additive objective function like (1), the belief state is a distribution that describes the probability of being in state s , given all past observations and actions [4], namely $b(s_n) = \mathbb{P}[s_n | u_0, \dots, u_{n-1}, y_1, \dots, y_n]$.

By treating the belief state as the true state of the system, (1) can be equivalently solved by generating and recursively solving a value function over the belief state, namely

$$V_n^*(b) = \min_{u \in \mathcal{U}} \left\{ \sum_s C(s, u) b(s) + \sum_y V_{n+1}^*(M_{y,u} b) \mathbb{P}(y|b, u) \right\}, \quad (7)$$

initialised at N as $V_N^*(b) = \sum_s C(s, u) b(s)$, with the transition operator $M_{y,u} b$ that provides the next belief state b_{n+1} given the current observation, action, and belief state according to a Bayesian update

$$(M_{y,u} b)(s') = \frac{Y(y|s', u) \sum_{s \in \mathcal{S}} T(s'|s, u) b(s)}{\mathbb{P}(y|b, u)}, \quad (8)$$

with the likelihood of the observation given by

$$\mathbb{P}(y|b, u) = \sum_{s \in \mathcal{S}} b(s) \sum_{s' \in \mathcal{S}} T(s'|s, u) Y(y|s', u). \quad (9)$$

An optimal policy π^* for the POMDP is then defined in terms of the belief state, and maps beliefs to actions: $\pi^* : \mathfrak{B} \rightarrow \mathcal{U}$ (\mathfrak{B} denotes the space of all beliefs). The optimal policy π^* is obtained directly from the value function as $\pi_n^*(b) = \arg \min_{u \in \mathcal{U}} V_n^u(b)$.

The value function in (7) for a POMDP is often represented in an alternate form that makes it conceptually easier to solve. If we pose the problem as a maximization (i.e. maximize the sum of costs multiplied by -1), then for a finite horizon $N < \infty$, the value function at each time n is piecewise-linear and convex in the belief state [19], and thus can be expressed as

$$V_n^*(b) = \max_{\alpha_n \in \Gamma_n} \sum_s \alpha_n(s) b(s). \quad (10)$$

The functions $\alpha_n \in \mathbb{R}^{|\mathcal{S}|}$, or “ α -vectors”, represent a bounded portion of the value function, and characterise the current value of being in each state $s \in \mathcal{S}$, given a specific action u is taken, plus the expected sum of future rewards assuming that all subsequent actions are chosen optimally. Because each α -vector is associated with a specific action, by picking the α -vector maximising $\sum_s \alpha_n(s) b(s)$ we also define the optimal policy for belief b at time n . The collection of α -vectors needed to exactly represent the value function V_n at time n is finite and represented by a set denoted as Γ_n .

The α -vectors at time n are computed recursively from the α -vectors known at time $n+1$. For each action, we access one of $|\mathcal{Y}|$ observations (where $|\cdot|$ indicates the cardinality of the set), and for each observation there is a subsequent α -vector defined at time $n+1$, resulting in $|\mathcal{U}| |\Gamma_{n+1}|^{|\mathcal{Y}|}$ α -vectors at time n . Rather than computing and storing an exponentially growing set of α -vectors, one option is to use point-based value iteration (PBVI), which is a sampling-based method used to approximate value functions [18].

There are many variants of PBVI around a basic algorithm. A finite subset $B \subset \mathfrak{B}$ is generated through some form of sampling, and for each $b \in B$, the value function is estimated by producing a single α -vector for each sampled belief state at each iteration. Therefore, if $B = (b^0, b^1, \dots, b^m)$, then $\Gamma_n = (\alpha_n^0, \alpha_n^1, \dots, \alpha_n^m)$ for all n , and the collection of α -vectors does not increase in size at each iteration. For a finite horizon problem, we distinguish the belief state at different times, so that $B = \bigcup_{n=0}^N B_n$, with

$B_n = \{b_n^1, \dots, b_n^{m_n}\}$ and $B_0 = \{R\}$, the initial distribution of the POMDP of interest.

For an additive objective function, the α -vectors are computed as in [17]. First, for each belief state b^k and at each time step n , a set of intermediate α -vectors $\alpha_{n,y,u}^j$ are computed recursively from $\tilde{\Gamma}_{n+1}$

$$\alpha_{n,y,u}^j(s) = C(s, u) + \sum_{s'} \alpha_{n+1}^j(s') Y(y|s') T(s'|s, u) \quad (11)$$

for $\alpha_{n+1}^j \in \tilde{\Gamma}_{n+1}$. Then, for each y and b^k we define $y(k) = \arg \max_j \sum_s \alpha_{n,y,u}^j(s) b^k(s)$. The final α -vectors associated with a single control input and belief state are given as

$$\alpha_{n,u}^k(s) = \sum_y \alpha_{n,y,u}^{y(k)}(s) \quad (12)$$

and $\alpha_n^k = \alpha_{n,u^*}^k(s)$, with $u^* = \arg \max_{u \in \mathcal{U}} \sum_s \alpha_{n,u}^k(s) b^k(s)$, is added to $\tilde{\Gamma}_n$.

An α -vector α_n^j corresponding to b^j will likely apply to all belief points in a region around b^j (i.e. for any b in a neighborhood of b^j the same action will likely be optimal). Hence the value at some b not necessarily in B can be approximated by $V_n^*(b) \approx \max_{\alpha_n^i \in \tilde{\Gamma}_n} \sum_s \alpha_n^i(s) b(s)$ as in (10) but with the restricted set $\tilde{\Gamma}_n \subset \Gamma_n$. The PBVI algorithm then consists of selecting a set of belief points $\{B_n\}$, and computing α_n^k for each $b_n^k \in B_n$.

The convexity of the value function guarantees that by storing only a subset $\tilde{\Gamma}_n$ of α -vectors, the approximation provides a lower bound to the maximum expected cost at any belief state (or an upper bound to the minimum), i.e. $\max_{\alpha_n^i \in \tilde{\Gamma}_n} \sum_s \alpha_n^i(s) b(s) \leq \max_{\alpha_n^i \in \Gamma_n} \sum_s \alpha_n^i(s) b(s)$ for any $b \in \mathfrak{B}$.

2.3 Safety Objectives for POMDPs

Rather than only minimising an expected sum of costs as in (1), we consider the specific task of generating safety-preserving controllers (where safety is preserved with maximal probability), with the additional objective of minimising the associated costs. This is a multi-objective goal with heterogeneous components (safety and performance). The techniques presented can, however, be extended to more than two separate objectives, and also to other complex specifications beyond safety (reach-avoid, reachability, etc.).

Specifically, consider a safe or desired subset of the state space, denoted $K \subset \mathcal{S}$, in which we would like the state s_n to remain over a finite time horizon $n = 0, \dots, N$. We would like to design a control policy that maximises the probability that $s_n \in K$ for $n = 0, \dots, N$ and to determine the value of that maximal probability – this task is known as “quantitative verification.” In other words, we would like to find

$$p_{\text{safe}}^N(\pi, R; K) = \mathbb{P}^\pi [s_0 \in K, \dots, s_N \in K | R], \quad (13)$$

$$\pi_{\text{safe}}^* = \arg \max_{\pi \in \Pi} p_{\text{safe}}^N(\pi, R; K). \quad (14)$$

We can express (13) using standard stochastic optimal control notation (i.e. the cost objective is an expected value of a function of one-step costs) by recalling that for a random variable s and set K , the probability that $s \in K$ is equal to the expected value of the indicator function over that event,

$\mathbb{E}[\mathbf{1}_K(s)]$ [2]. Hence, (13) is equal to

$$p_{\text{safe}}^N(\pi, R; K) = \mathbb{P}^\pi \left[\prod_{n=0}^N \mathbf{1}_K(s_n) \middle| R \right].$$

To maximise the probability of safety is therefore equivalent to maximizing a *multiplicative* objective function (rather than additive, as discussed in Section 2.2). The belief state $b(s_n) = \mathbb{P}[s_n | i_n]$ is *no longer* sufficient when the objective is multiplicative, because the costs accrued at previous time steps must also be considered, as they affect the current possible cost (i.e. if any previous state left the safe set, the system can not in subsequent time steps be considered safe). Two options to construct an equivalent MDP and value function over a fully observable state to compute (14) are 1) To derive a sufficient statistic and to redefine the belief state directly over the multiplicative objective function, as in [13], or 2) To reformulate the POMDP so that (13) can be framed as an additive objective function, as in [7]. We will present the second option, as it will subsequently allow us to use a lexicographic approach for the multi-objective problem.

To express (13) as an expected value over a sum of costs, an additional binary state variable q_n is introduced (which is not observed),

$$q_n = \prod_{i=0}^{n-1} \mathbf{1}_K(s_i). \quad (15)$$

The variable q_n keeps track of whether the state of the system has remained within K up to the previous time step, thus keeping track of the previous “costs” incurred (where the cost is simply equal to one or zero, depending on whether the system has remained safe or not).

The new state of the system is $\bar{s} = (s, q) \in \mathcal{S} \times \{0, 1\} = \bar{\mathcal{S}}$, and the safety objective can be rewritten as

$$\max_{\pi \in \Pi} \mathbb{E}^\pi \left[\prod_{n=0}^N \mathbf{1}_K(s_n) \middle| R \right] = \max_{\pi \in \Pi} \mathbb{E}^\pi [\mathbf{1}_{K \times \{1\}}(s_N, q_N) | R]. \quad (16)$$

The cost function $C(s_n, u_n)$ from $n = 0, \dots, N-1$ is $C(\bar{s}_n, u_n) = 0$, and at time N , $C(\bar{s}_N, u_N) = \mathbf{1}_{K \times \{1\}}(s_N, q_N)$, and so (16) is an additive objective function with zero costs accrued until the final time step.

The transition probability function T is replaced by $\bar{T} : \bar{\mathcal{S}} \times \bar{\mathcal{S}} \times \mathcal{U} \rightarrow [0, 1]$, with

$$\bar{T}(s', q' | s, q, u) = \begin{cases} T(s' | s, u), & \text{if } q = q' = 0 \\ 0, & \text{if } q = 0, q' = 1 \\ \mathbf{1}_{\mathcal{S} \setminus K}(s) T(s' | s, u), & \text{if } q = 1, q' = 0 \\ \mathbf{1}_K(s) T(s' | s, u), & \text{if } q = q' = 1 \end{cases} \quad (17)$$

We also transform $Y(y|s)$ to $\bar{Y}(y|\bar{s})$, such that $\bar{Y}(y|s, q) = Y(y|s)$, which is independent of q , and similarly transform $R(s)$ to $\bar{R}(s, q) = R(s)$.

Therefore, the problem of maximizing the probability of safety for a POMDP \mathcal{J} can be posed equivalently as an additive cost problem over a different POMDP $\bar{\mathcal{J}} = (\bar{\mathcal{S}}, \mathcal{Y}, \mathcal{U}, \bar{T}, \bar{Y}, \bar{R})$. The belief state is now the conditional probability distribution of the current state of the system conditioned on all past observations and control inputs, only now the current state is \bar{s} rather than just s : $\bar{b}(\bar{s}_n) = \mathbb{P}[s_n, q_n | i_n]$.

3. PROBLEM FORMULATION

We would like to consider the problem of maximising the probability that the state of a POMDP remains safe over a finite time horizon, while simultaneously minimising an expected sum of costs. It may be the case that the control policy π^* from (14) that produces the highest probability that the system remains safe is simply not feasible to implement from a cost perspective. It may also be the case that a small sacrifice in the optimal safety level (or probability) of the system leads to a significant reduction in costs. Depending on the requirements on safety-criticality of the system at hand, we can define a tolerance η , which indicates how much we are willing to lower the probability of safety from the optimal one, in order to attempt to reduce the expected cost by implementing a (sub-)optimal policy.

We will denote by $\tilde{\Pi}$ the set of control policies that produce a probability of safety within the desired tolerance (no more than η below the maximum), namely

$$\tilde{\Pi} = \{\pi : p_{\text{safe}}^N(\pi_{\text{safe}}^*, R; K) - p_{\text{safe}}^N(\pi, R; K) \leq \eta\}. \quad (18)$$

We would then like to find the policy $\tilde{\pi}^* \in \tilde{\Pi}$ that minimises the expected sum of costs:

$$\tilde{\pi}^* = \arg \min_{\pi \in \tilde{\Pi}} \mathbb{E}^\pi \left[\sum_{n=0}^N C(s_n, u_n) \middle| R \right]. \quad (19)$$

We therefore have a multi-objective optimisation problem of a hierarchical nature, where one objective takes priority over the other, and whose cost functions are inherently different.

PROBLEM 1. *Given a POMDP \mathcal{J} , a safe set $K \subset \mathcal{S}$, a cost function $C(s, u)$, a time horizon $N < \infty$, and a tolerance η , we would like to synthesize a control policy $\tilde{\pi}^*$ that*

1. *Guarantees that the probability that the system state s_n remains in K for all $n = 0, \dots, N$ (13) is no less than η below the maximum possible safety probability, i.e.*

$$p_{\text{safe}}^N(\pi_{\text{safe}}^*, R; K) - p_{\text{safe}}^N(\tilde{\pi}^*, R; K) \leq \eta, \quad (20)$$

2. *Guarantees that the expected cost using policy $\tilde{\pi}^*$ is minimal over all policies that produce a safety probability within η of the optimal, i.e. $\tilde{\pi}^*$ is the solution to (19).*

3.0.1 Case Study Continued

Returning to the boiler maintenance problem, we have two main objectives. The first is to ensure that the boiler degradation level does not reach 80, at which point the boiler is essentially broken and would result in excessive maintenance costs and possible loss of heat over an extended period of time. To avoid this unsafe outcome, the preventive maintenance action ‘‘Clean’’ may be taken, which also incurs a cost but is less costly than if the boiler were to break.

In addition, the higher the degradation level, the less efficient the boiler is, and the more expensive it is (the more fuel required) for the boiler to maintain a set temperature. It is therefore preferable to do preventive maintenance, and to clean the boiler before the degradation level grows too high both to avoid a breakdown, and also to reduce fuel costs.

The total cost at each time step, assuming the degradation level has not reached 80, is given by the function

$$C(s, u) = 0.05 * s_n + 100 * \mathbf{1}_{\text{Clean}}(u). \quad (21)$$

The first term captures efficiency costs associated with the degradation level, and the second gives the cost of cleaning the boiler (100 Euros) if $u = \text{Clean}$ is selected.

Rather than assigning a cost to the event that the degradation level reaches 80, we treat the set $K^c = \{80, \dots, 100\}$ as an unsafe set that we would like to avoid, and we wish to find a cleaning policy (at what point the boiler should be cleaned) that maximises the probability that the unsafe set is not reached. We would further like to *know* this probability, so that we can provide guarantees on the safety of the boiler under the policy we propose.

If we were to only consider safety, an intuitive optimally safe policy would be to clean the boiler all the time, regardless of the degradation level. From a cost standpoint, however, this is not practical, and we should also factor in the fuel costs as well as costs associated with cleaning when designing a safe policy.

4. LEXICOGRAPHIC OPTIMAL CONTROL

We consider now multiple objective functions to be maximised concurrently: we shall employ the same belief state, which by definition does not depend on the cost function so long as the costs incurred are additive at each time step, as the input to multiple value functions, one for each cost objective. We can then set up a POMDP with lexicographic optimization criteria, introduced in [22], with k cost functions $\{C_1, \dots, C_k\}$ and with tolerance parameters $\{\eta_1, \dots, \eta_{k-1}\}$ associated with the first $k-1$ cost functions. The cost functions are ordered so that C_1 is the most important to optimise, and C_k the least important. Each cost function C_i is optimised using the value function $V_{n,i}$, defined as in (7) with C replaced by C_i .

The idea behind the lexicographic approach is to progressively limit the set of control inputs available to optimise the value functions $V_{n,i}(b)$ at each time step n , and for each belief state, according to the preference ordering. The restricted set of control inputs available to value function $V_{n,i+1}(b)$ for a specific belief state b , denoted $\mathcal{U}_{n,i}(b)$, is chosen such that the value function $V_{n,i}^u(b)$ (evaluated at a specific $u \in \mathcal{U}_{n,i-1}$) is within $\bar{\eta}_i$, the *one step* slack allowance (as opposed to η_i , which is the total slack allowance over the entire time horizon), of $V_{n,i}^*(b)$, and thus ultimately ensures that the policy $\tilde{\pi}^*$ generated by selecting the optimal control input $u_n^* \in \mathcal{U}_{n,k-1}(b)$ for any belief b at any time n , is within $\tilde{\Pi}$.

The actual procedure is as follows. Starting at the final time step N , we construct $V_{N,i}^u(b) = \sum_s C_i(s, u)b(s)$ for all $i = 1, \dots, k$. For each $b \in \mathfrak{B}$, we then construct the set of permissible control inputs $\mathcal{U}_{N,i}(b)$ for all $i = 1, \dots, k-1$, defined as follows:

$$\mathcal{U}_{N,i}(b) = \left\{ u \in \mathcal{U} : V_{N,i}^u(b) - \min_{\bar{u} \in \mathcal{U}_{N,i-1}} V_{N,i}^{\bar{u}}(b) \leq \bar{\eta}_i \right\}. \quad (22)$$

How to choose $\bar{\eta}_i$ under various conditions will be described subsequently. The set $\mathcal{U}_{N,0}$ is simply \mathcal{U} .

Once $\mathcal{U}_{N,k-1}(b)$ is computed, we find $u^* = \arg \min_{u \in \mathcal{U}_{N,k-1}} V_{N,k}^u(b)$. We then set $V_{N,i}^*(b) = \sum_s C_i(s, u^*)b(s)$. The process is repeated for all $b \in \mathfrak{B}$, and for all time steps $n = N-1, \dots, 0$, using the value function in (7) with each cost function C_i , and with the restricted control inputs $\mathcal{U}_{n,i}(b)$.

The α -vector representation of the value function associated with a POMDP allows us to alternatively characterise

the lexicographic preferences through sets of α -vectors, rather than the sets of control inputs $\mathcal{U}_{n,i}(b)$. Similar to the sets $\mathcal{U}_{n,i}(b)$, we generate sequential sets $\bar{\Gamma}_{n,i} \subset \Gamma_{n,i}$, with the difference being that the sets $\bar{\Gamma}_{n,i}$ are not generated separately for each belief state. For time N , we set $\bar{\Gamma}_{N,1} = \Gamma_{N,1}$, and for $n < N$, $\bar{\Gamma}_{n,1} = \hat{\Gamma}_{n,1}$, where $\hat{\Gamma}_{n,1}$ is the full set of α -vectors (for all possible control inputs) computed recursively from $\bar{\Gamma}_{n+1,1}$. For $i = 2, \dots, k$,

$$\bar{\Gamma}_{n,i} = \left\{ \alpha \in \hat{\Gamma}_{n,i} : \exists b \in \mathfrak{B} \text{ s.t. } \max_{\alpha' \in \bar{\Gamma}_{n,i-1}} \alpha' \cdot b - \alpha'' \cdot b \leq \bar{\eta}_i \right. \\ \left. \text{with } \alpha'' \in \bar{\Gamma}_{n,i-1}, \text{ and } u_\alpha = u_{\alpha''} \right\}. \quad (23)$$

The notation u_α refers to the control input associated with vector α . In other words, the control inputs allowed when generating the α -vectors for $\bar{\Gamma}_{n,i}$ are restricted to those associated with α -vectors in $\bar{\Gamma}_{n,i-1}$ such that there exists some $b \in \mathfrak{B}$ for which the inner product between the α -vector and b is within $\bar{\eta}_i$ of the optimal. We no longer compute and store all α -vectors for each value function $V_{n,i}$, but rather prune according to whether they meet the tolerance specification for the previous cost function.

4.1 Application to a Prioritised Safety Objective

The lexicographic optimal control approach can be used when the priority is a safety objective in (13) by using the additive objective formulation that relies on including the binary variable q_n as part of the state, as described in Section 2.3. This formulation allows us to use the same belief states for the value functions associated with both the safety and cost objectives. The value functions for the safety objective are

$$V_{n,1}^u(\bar{b}) = \sum_{s,q} \mathbf{1}_{K \times \{1\}}(s,q) \bar{b}(s,q), \\ V_{n,1}^u(\bar{b}) = \sum_{s,q} V_{n+1,1}^*(\bar{M}_{y,u} \bar{b}) \mathbb{P}(y|\bar{b},u), \quad (24)$$

with \bar{M} and $\mathbb{P}(y|\bar{b},u)$ defined as in (8) and (9), only with T replaced by \bar{T} , Y replaced by \bar{Y} , and summations over \bar{s} . The value function for minimizing the cost function C is

$$V_{N,2}^u(\bar{b}) = \sum_{s,q} C(s,u) \bar{b}(s,q), \\ V_{n,2}^u(\bar{b}) = \sum_{s,q} C(s,u) \bar{b}(s,q) + \sum_y V_{n+1,2}^*(\bar{M}_{y,u} \bar{b}) \mathbb{P}(y|\bar{b},u). \quad (25)$$

Notice that the variable q_n is essentially ignored in the value function $V_{n,2}$, and $\bar{b}(\bar{s})$ is reduced to $b(s)$ by treating $b(s)$ as the marginal distribution, namely $b(s) = \sum_q \bar{b}(s,q)$.

Problem 1 can then be solved in theory using the above belief state and value functions, and by applying the lexicographic approach described above with $k = 2$ separate cost functions, and a single tolerance parameter η (and also a single one-step tolerance parameter $\bar{\eta}$). As seen with a single objective function, however, (24) and (25) cannot be solved exactly, but instead approximated using a modification of PBVI.

4.2 Approximate Solution using PBVI

In order to use PBVI in a lexicographic setting with both a safety and a cost objective, we need to compute two sets of

α -vectors, one for the safety objective, and one for the cost, which we denote $\alpha_{n,1}^k$ and $\alpha_{n,2}^k$, respectively. The doubly restricted sets of α -vectors, denoted $\bar{\Gamma}_{n,i}$ (restricted because we compute less α -vectors for the PBVI algorithm, and because we store different α -vectors depending on the tolerance parameter η) are generated according to the standard PBVI algorithm, with two modifications.

The first modification is in computing the index $y(k)$. If we were to compute the exact solution, the α -vectors for *all possible* combinations of observations and control inputs would be generated, independently of the belief state, and therefore the selection of the optimal index $y(k)$ is not necessary. However, when applying PBVI, we must be careful to ensure that the index $y(k)$ computed for $\alpha_{n,y,u,1}^{y(k)}$ and $\alpha_{n,y,u,2}^{y(k)}$ is the same, i.e. the same α -vector $\alpha_{n+1}^{y(k)}$, and hence the same optimal policy at subsequent time steps, is associated with both $\alpha_{n,1}^k$ and $\alpha_{n,2}^k$. Otherwise, different optimal policies would be associated with the safety and cost objectives, implying the selection of different control inputs for the same belief state at time $n + 1$ depending on which objective is considered. Since the idea is to select a *single* control input to simultaneously optimize both objectives, this cannot be allowed.

Further, the term $\sum_{s'} \alpha_{n+1}^j(s') Y(y|s') T(s'|s,u)$ in (11) is proportional to $\sum_{s'} \alpha_{n+1}^j(s') b_{n+1}(s')$. Hence when $y(k)$ is selected, we are actually optimizing the value function at time $n + 1$. The b_{n+1} produced by (11) may not, however, be one of the elements in B_{n+1} , and so the exact optimal α_{n+1}^j likely has not been computed. The implication when performing the lexicographic optimization is that the (sub)optimal α_{n+1}^j designed to maximize the safety objective to within $\bar{\eta}$ of the optimal, and minimize the cost objective, likely has also not been computed. Therefore the lexicographic optimization must be redone for the belief state \bar{b}_{n+1} generated by $\alpha_{n,y,u,i}^j$, and the tolerance preference $\bar{\eta}$ must once again be enforced. The modified index, denoted $\bar{y}(k)$, is computed according to Algorithm 1, with inputs $\{\alpha_{n,y,u,1}^j\}_j$, $\{\alpha_{n,y,u,2}^j\}_j$, \bar{b}^k , and $\bar{\mathcal{S}}$. Although PBVI is mentioned as an approximation algorithm for the lexicographic optimization of POMDPs in [22], they fail to mention this subtle difference from the standard, single objective PBVI algorithm.

The α -vectors $\alpha_{n,1}^k$ and $\alpha_{n,2}^k$ are defined as follows:

$$\alpha_{N,u,1}^k(\bar{s}) = \mathbf{1}_{K \times \{1\}}(\bar{s}), \quad \forall \bar{b}^k \in \bar{B}_N \\ \alpha_{n,u,1}^k(\bar{s}) = \sum_{y \in \mathcal{Y}} \sum_{\bar{s}' \in \bar{\mathcal{S}}} \alpha_{n+1,1}^{\bar{y}(k)}(\bar{s}') \bar{Y}(y|\bar{s}') \bar{T}(\bar{s}'|\bar{s},u), \quad \forall \bar{b}^k \in \bar{B}_n, \\ \forall n = N - 1, \dots, 0 \quad (26)$$

$$\alpha_{N,u,2}^k(\bar{s}) = C(s,u), \quad \forall \bar{b}^k \in \bar{B}_N \\ \alpha_{n,u,2}^k(\bar{s}) = C(s,u) + \sum_{y \in \mathcal{Y}} \sum_{\bar{s}' \in \bar{\mathcal{S}}} \alpha_{n+1,u,2}^{\bar{y}(k)}(\bar{s}') \bar{Y}(y|\bar{s}') \bar{T}(\bar{s}'|\bar{s},u), \\ \forall \bar{b}^k \in \bar{B}_n, \forall n = N - 1, \dots, 0 \quad (27)$$

The second modification is the same as for the exact lexicographic solution described in Section 2.2. Once the complete sets of α -vectors $\alpha_{n,u,i}^k$ are computed according to (26) and (27), the sets $\bar{\Gamma}_{n,i}$ are generated by first constructing the set $\mathcal{U}_n(\bar{b}^k)$, which can be expressed in terms of the α -vectors

Algorithm 1 LexicoOpt

Input: $\{\alpha_1^i\}_{i=1}^M, \{\alpha_2^i\}_{i=1}^M, b, \bar{\eta}$ **Output:** i^* , optimal index associated with α_1, α_2

```

1:  $i^* \leftarrow \arg \max_i \sum_{s \in \mathcal{S}} \alpha_1^i(s)b(s)$ 
2:  $\max\text{Safe} \leftarrow \sum_{s \in \mathcal{S}} \alpha_1^{i^*}(s)b(s)$ 
3:  $\min\text{Cost} \leftarrow \sum_{s \in \mathcal{S}} \alpha_2^{i^*}(s)b(s)$ 
4: for  $i = 1, \dots, M$  do
5:    $\text{newSafe} \leftarrow \sum_{s \in \mathcal{S}} \alpha_1^i(s)b(s)$ 
6:    $\text{newCost} \leftarrow \sum_{s \in \mathcal{S}} \alpha_2^i(s)b(s)$ 
7:   if  $\max\text{Safe} - \text{newSafe} \leq \bar{\eta}$  then
8:     if  $\text{newCost} < \min\text{Cost}$  then
9:        $i^* \leftarrow i$ 
10:       $\min\text{Cost} \leftarrow \sum_{s \in \mathcal{S}} \alpha_2^{i^*}(s)b(s)$ 
11:    end if
12:  end if
13: end for

```

as

$$\mathcal{U}_n(\bar{b}^k) = \left\{ u \in \mathcal{U} : \max_{\bar{s} \in \mathcal{U}} \sum_{\bar{s}} \alpha_{n,\bar{u},1}^k(\bar{s}) \bar{b}^k(\bar{s}) - \sum_{\bar{s}} \alpha_{n,\bar{u},1}^k(\bar{s}) \bar{b}^k(\bar{s}) \leq \bar{\eta} \right\}. \quad (28)$$

The optimal control input u^* is chosen as

$$u^* = \arg \min_{u \in \mathcal{U}_n(\bar{b}^k)} \sum_{\bar{s}} \alpha_{n,u,2}^k(\bar{s}) \bar{b}^k(\bar{s}), \quad (29)$$

and $\alpha_{n,1}^k = \alpha_{n,u^*,1}^k$ and $\alpha_{n,2}^k = \alpha_{n,u^*,2}^k$ are added to $\tilde{\Gamma}_{n,1}$ and $\tilde{\Gamma}_{n,2}$, respectively. Algorithm 2 summarizes the entire procedure for solving Problem 1 using PBVI and a lexicographic approach.

Note that Algorithm 2 returns the sets $\{\tilde{\Gamma}_{n,1}\}_{n=0}^N$ and $\{\tilde{\Gamma}_{n,2}\}_{n=0}^N$, which in turn provide the multi-objective policy $\tilde{\pi}^*$. The optimal control input at time n , for any belief state $\bar{b} \in \mathfrak{B}$, is found again according to Algorithm 1, with inputs $\tilde{\Gamma}_{n,1}$, $\tilde{\Gamma}_{n,2}$, and \bar{b} . Using the returned index i^* , u^* is equal to the control input associated with both $\alpha_{n,1}^{i^*}$ and $\alpha_{n,2}^{i^*}$, which is guaranteed to be the same based on the structure of Algorithm 2.

The probability that the system remains safe, if policy $\tilde{\pi}^*$ is implemented, and the expected cost associated with the policy, are bounded by

$$p_{\text{safe}}^N(\tilde{\pi}^*, R; K) \geq \sum_{\bar{s}} \alpha_{0,1}^1(\bar{s}) \bar{R}(\bar{s}) \quad (30)$$

$$\mathbb{E}^{\tilde{\pi}^*} \left[\sum_{n=0}^N C(s_n, \tilde{\pi}_n^*(\bar{b}_n)) \middle| R \right] \leq \sum_{\bar{s}} \alpha_{0,2}^1(\bar{s}) \bar{R}(\bar{s}),$$

since $\bar{b}^1 = \bar{R}$.

4.3 Bound on the One-Step Tolerance

We next discuss how to select the one step tolerance $\bar{\eta}$ to ensure that the approximate safety probability returned by PBVI is within tolerance η of the exact optimal safety probability, $p_{\text{safe}}^N(\pi_{\text{safe}}^*, R; K)$. The proof is similar to that in [22], although without the modifications to the PBVI algorithm discussed in the previous section, the claims in [22]

Algorithm 2 LexicoPBVI

Input: $\{B_n\}_{n=0}^N, K, C(\cdot, \cdot), N, \mathcal{J}, \bar{\eta}$ **Output:** $\{\tilde{\Gamma}_{n,1}\}_{n=0}^N, \{\tilde{\Gamma}_{n,2}\}_{n=0}^N$

```

1:  $\tilde{\Gamma}_{N,1} = \emptyset, \tilde{\Gamma}_{N,2} = \emptyset$ 
2:  $\alpha_{N,1}(\cdot) = \mathbf{1}_K(\cdot), \tilde{\Gamma}_{N,1} = \tilde{\Gamma}_{N,1} \cup \{\alpha_{N,1}\}$ 
3: for all  $b \in B_N$  do
4:    $u^* = \arg \min_{u \in \mathcal{U}} \sum_{s \in \mathcal{S}} C(s, u)b(s)$ 
5:    $\alpha_{N,2}(\cdot) = C(\cdot, u^*), \tilde{\Gamma}_{N,2} = \tilde{\Gamma}_{N,2} \cup \{\alpha_{N,2}\}$ 
6: end for
7: for  $n = N - 1, \dots, 0$  do
8:    $\tilde{\Gamma}_{n,1} = \emptyset, \tilde{\Gamma}_{n,2} = \emptyset$ 
9:   for all  $b^k \in B_n$  do
10:    for all  $u \in \mathcal{U}$  do
11:       $\alpha_{n,u,1}^k = \mathbf{0}, \alpha_{n,u,2}^k = \mathbf{0}$ 
12:      for all  $y \in \mathcal{Y}$  do
13:        for  $j = 1, \dots, |\tilde{\Gamma}_{n+1,1}|$  do
14:          Compute  $\alpha_{n,y,u,1}^j, \alpha_{n,y,u,2}^j$  according to (26)
          and (27), respectively
15:        end for
16:         $j^* = \text{LexicoOpt}(\{\alpha_{n,y,u,1}^j\}_j, \{\alpha_{n,y,u,2}^j\}_j, b^k, \bar{\eta})$ 
17:         $\alpha_{n,u,1}^k = \alpha_{n,u,1}^k + \alpha_{n,y,u,1}^{j^*}$ 
18:         $\alpha_{n,u,2}^k = \alpha_{n,u,2}^k + \alpha_{n,y,u,2}^{j^*}$ 
19:      end for
20:    end for
21:     $u^* = \text{LexicoOpt}(\{\alpha_{n,u,1}^k\}_u, \{\alpha_{n,u,2}^k\}_u, b^k, \bar{\eta})$ 
22:     $\tilde{\Gamma}_{n,1} = \tilde{\Gamma}_{n,1} \cup \{\alpha_{n,u^*,1}^k\}, \tilde{\Gamma}_{n,2} = \tilde{\Gamma}_{n,2} \cup \{\alpha_{n,u^*,2}^k\}$ 
23:  end for
24: end for

```

are incorrect. We first derive the one-step tolerance needed to guarantee an overall tolerance η , if we were able to solve the lexicographic optimization exactly (without PBVI).

PROPOSITION 1. *For time horizon N , initial distribution R , and total desired tolerance η , setting the one step tolerance $\bar{\eta}$ to $\bar{\eta} = \frac{\eta}{N}$ ensures that the safety probability found by exactly solving the lexicographic POMDP optimization, denoted $V_{0,1}^\eta(\bar{R})$ satisfies $p_{\text{safe}}^N(\pi^*, R; K) - V_{0,1}^\eta(\bar{R}) \leq \eta$.*

As an extension, in order to ensure that the selected tolerance is not exceeded when using PBVI, we must also take into account the error introduced in the PBVI algorithm when maximising the safety objective. This error is given in [14] and is a straightforward extension of the standard PBVI error presented in [17]. First, we employ parameter δ^B to denote how densely the belief space \mathfrak{B} has been sampled to produce \bar{B} , defined as the maximum Hausdorff distance with respect to n between \bar{B}_n and \mathfrak{B}_n :

$$\delta^B = \max_n \left\{ \max \left\{ \sup_{\bar{b}_n \in \bar{B}_n} \inf_{\bar{b}_n \in \mathfrak{B}_n} \|\bar{b}_n - \bar{b}_n\|_1, \sup_{\bar{b}_n \in \mathfrak{B}_n} \inf_{\bar{b}_n \in \bar{B}_n} \|\bar{b}_n - \bar{b}_n\|_1 \right\} \right\}. \quad (31)$$

Denoting V_n^B as the optimal value function for a single objective at time n computed using PBVI, and V_n^* as the exact optimal value function, then

$$V_n^*(\bar{b}) - V_n^B(\bar{b}) \leq (N - n)\delta^B \quad (32)$$

for all $\bar{b} \in \bar{\mathfrak{B}}$, and for all $n = 0, \dots, N$. We obtain:

PROPOSITION 2. *For time horizon N , initial distribution R , and total desired tolerance η , if $\frac{\eta}{N} - \delta^B \geq 0$, setting the one step tolerance $\bar{\eta}$ to $\bar{\eta} = \frac{\eta}{N} - \delta^B$ ensures that the safety probability found by solving the lexicographic POMDP optimisation using PBVI, denoted $V_{0,1}^{\eta,B}(\bar{R})$ satisfies $p_{\text{safe}}^N(\pi^*, K; R) - V_{0,1}^{\eta,B}(\bar{R}) \leq \eta$.*

Proposition 2 guarantees that the safety probability returned using Algorithm 2 will be within η of $p_{\text{safe}}^N(\pi^*, K; R)$, as long as the error $N\delta^B$ incurred as a result solely from the PBVI algorithm is not too large (i.e. does not exceed η). Fortunately, the PBVI error is tunable, and δ^B decreases to zero as the sampled set B approaches \mathfrak{B} . Unfortunately, the more points $b \in B$ that are generated, the more computation time is required for Algorithm 2, so there is an inherent trade-off between the accuracy of the PBVI algorithm and hence our ability to satisfy the tolerance specification η , and the amount of time required to compute the optimal policy for both the safety and cost objectives.

5. EXTENSION TO CPS MODELS

By definition, cyber-physical systems typically comprise processes evolving in a continuous space (the physical system), in combination with processes evolving in a discrete fashion (the “cyber”, or computational component embedded in the physical process). They are therefore often modeled within a hybrid system framework, with a state space $\mathcal{S} = \mathcal{X} \times \mathcal{Q}$, $\mathcal{X} \subseteq \mathbb{R}^n$, and $\mathcal{Q} = \{q_1, \dots, q_l\}$ a finite set of modes. A partially observable hybrid system may also have a combination of continuous and discrete observations, so that $\mathcal{Y} = \mathcal{Y}^x \times \mathcal{Y}^q$, $\mathcal{Y}^x \subset \mathbb{R}^m$ and $\mathcal{Y}^q = \{y_1^q, \dots, y_p^q\}$.

5.0.1 Case Study: Parsimonious Temperature Regulation

As an example of a CPS, consider the task of controlling a boiler to switch on or off to heat a single room in a building. The safe temperature range is between 17.5 and 22 degrees Celsius, and the safety objective is to construct a control policy to maximize the probability that the temperature in the room does not leave this range. Simultaneously we would like to limit the amount of time that the heater is on, to reduce operational costs, i.e. we want to minimize the cost function $\sum_{n=0}^N C(u_n)$, with $C(0) = 0$ and $C(1) = 1$.

The temperature evolution is approximated by an affine stochastic difference equation, as presented in [2]:

$$x_{n+1} = (1 - a)x_n + bu_n + ax_a + v_n, \quad (33)$$

where x_n is the temperature at time step n , x_a is the ambient temperature (assumed constant), a and b are constants representing the heat loss rate to the external environment and the rate of heat supplied by the heater respectively, and $\{v_n\}$ is a sequence of i.i.d. Gaussian random variables with zero mean and variance ν^2 representing stochastic disturbances to the temperature.

While the temperature evolution is a continuous process, and the state takes values in $\mathcal{X} = \mathbb{R}$, the control input is a discrete signal $u \in \mathcal{U} = \{0, 1\}$ telling the boiler to switch on or off; 0 indicates the boiler is off, and 1 that it is on. The control input effectively renders the temperature evolution as a hybrid process that switches between two modes, one

when the heater is on, and the other when it is off. So we can alternately think of the hybrid state space $\mathcal{S} = \mathbb{R} \times \{0, 1\}$.

Additionally, the temperature is measured and communicated to the controller by a sensor that is subject to noise, represented by $y_n = x_n + w_n$, with $\{w_n\}$ an i.i.d. sequence of Gaussian random variables with zero mean and variance ω^2 . The control input is assumed known, and hence we only consider the continuous observation process, $\mathcal{Y} = \mathcal{Y}^x = \mathbb{R}$.

5.1 Abstraction to POMDP

The discussed PBVI algorithm requires as input a POMDP with a finite set of states, observations, and control inputs. In order to apply Algorithm 2 to a partially observable hybrid system, we must first create a *finite state abstraction* of the original hybrid system, to approximately represent it as a POMDP as in Definition 1. We are interested in the use of abstractions with guaranteed behaviours, since the abstraction introduces an additional error into the evaluation of safety probabilities using the PBVI algorithm, because we are now evaluating safety probabilities over a *different* system. Formal abstractions are discussed in [20] and, for models with partial observations, in [14].

For safety critical systems, quantifying the probability that the system remains safe is crucial. When approximations are introduced, we must quantify the effect of the approximation on the safety probabilities being computed. We have already done this in Section 4.3, where the one step tolerance bound $\bar{\eta}$ is adjusted as a function of the error introduced by the PBVI algorithm, to ensure that the safety probability returned by Algorithm 2 is no more than a pre-specified distance η from the maximum. We can repeat this process to get a new one step tolerance $\bar{\eta}$ that also takes into account the error from the abstraction as in [14].

To create an abstraction of a partially observable stochastic hybrid system, the continuous spaces \mathcal{X} and \mathcal{Y}^x are partitioned into cells, and a representative point is assigned to each cell. For a certain class of stochastic hybrid systems, the error introduced by this partitioning, in combination with a PBVI algorithm, has been shown to be a linear function of the size of the partitions, and hence decreases to zero as the partitions are made smaller. A detailed description of the abstraction process, as well as the error it incurs, is provided in [14].

Here we will simply assume that the error introduced by the abstraction is known, and equal to δ^a , i.e. we know that $|V_0^B(\bar{R}) - V_{0,\delta}^B(\bar{R}_\delta)| \leq \delta^a$. The subscript δ indicates that the abstraction is being used (both when computing the value functions, and also because the probability densities for the original system must be mapped to probability mass functions, hence $\bar{R} \rightarrow \bar{R}_\delta$).

PROPOSITION 3. *For time horizon N , a POMDP abstraction \mathcal{J}_δ of a partially observable stochastic hybrid system \mathcal{H} with abstraction error δ^a , and total desired tolerance η , if $\frac{\eta}{N} - \frac{2\delta^a}{N} - \delta^B \geq 0$, setting the one step tolerance to $\bar{\eta} = \frac{\eta - 2\delta^a}{N} - \delta^B$ ensures that the safety probability found by solving the lexicographic POMDP optimization using PBVI and the abstraction \mathcal{J}_δ , denoted $V_{0,1,\delta}^{\eta,B}(\bar{R}_\delta)$ satisfies $p_{\text{safe}}^N(\pi^*, K; R) - V_{0,1,\delta}^{\eta,B}(\bar{R}) \leq \eta$.*

6. APPLICATIONS TO BUILDINGS AUTOMATION SYSTEMS

We now consider two case studies. The first is the finite state boiler maintenance problem introduced in 2.1.1, and the second is the hybrid state temperature regulation problem introduced in 5.0.1. In each case we must utilize sensor measurements to design safe, cost-efficient control policies for building operation.

6.1 Optimised Boiler Maintenance

We apply the lexicographic approach in combination with PBVI to the boiler maintenance problem described in Sections 2.1.1 and 3.0.1. We would like to maximise the probability that the boiler does not reach a degradation level of 80 or above, while simultaneously minimising the costs associated with maintaining and operating the boiler. To use Algorithm 2, we use inputs $K = \{0, \dots, 79\}$, the cost function C given in (21), and the POMDP \mathcal{J} presented in Section 2.1.1 extended to include binary variable q discussed in Section 2.3. We consider a time horizon $N = 30$.

A set of 50 belief states B_n is generated sequentially for each time step. At initial time 0, the 50 belief states are initialised by randomly selecting the state in which the system starts with probability one, i.e. $b_0^i(s) = \mathbf{1}_{s^i}(s)$ with s^i selected uniformly at random from $\{0, \dots, 79\}$ (we are not interested in starting from the unsafe set). The action u_0^i is also chosen at random, and an observation y^i is sampled from $\mathbb{P}(y|b_0^i, u_0^i)$ (9). The next belief state b_1^i is generated according to $M_{y,u,b}$ in (8). This is done repeatedly for all $i = 1, \dots, 50$ and $n = 1, \dots, 30$.

Next, the one step tolerance must be computed in order to guarantee that the overall tolerance is $\eta = 0.1$, i.e. that the probability that the boiler does not break down while minimising costs is not more than 0.1 below the maximum probability, obtained when costs are disregarded. To provide an exact bound on the tolerance requires incorporating the error δ^B from PBVI, which is difficult to compute and can be quite conservative. We therefore disregard the error δ^B , for two reasons. First, PBVI guarantees a lower bound to the actual value function when we use the lexicographic approach, and therefore although we do not have an exact guarantee on the distance from the optimal safety probability, we know that the safety probability is at least as high as the estimate we produce. Second, for this example PBVI returns safety probabilities that are equal to one, and therefore the PBVI solution is exact, since the probabilities are both lower and upper bounded by one, thus $\delta^B = 0$. Hence, we set $\bar{\eta} = \frac{\eta}{N} = \frac{1}{30}$.

The probability that the degradation level s_n does not reach 80 or above over the considered time horizon, given varying starting degradation levels s_0 , is presented in Fig. 1a, using the lexicographic PBVI algorithm (Alg. 2) as well as PBVI for the single safety objective (without minimizing costs). Similarly, Fig. 1b shows the expected costs over 30 days when only safety is considered versus using the lexicographic algorithm. Fig. 1b shows that the policy returned by Algorithm 2 (lexicographic case) is a threshold policy, since the expected cost increases by 100 (the cost of cleaning) for $s_0 > 26$, meaning that over 30 time steps, if the initial degradation level is below 26, we should not clean the boiler, and if it is greater than 26, the boiler should be cleaned only once. In a region below the threshold level of 26, there is a small sacrifice in the probability of safety (as seen in Fig. 1a) that does not exceed the tolerance $\eta = 0.1$, and is therefore deemed acceptable. If we were to shrink

the tolerance level, the threshold degradation level at which cleaning should be performed would be lower ($s_0 < 26$).

We also test each policy (safety only and safety plus costs) by simulating 1000 trajectories of the degradation level under each policy. Fig. 2 shows the average control input (clean or do not clean) at each time step over the 1000 trajectories under each policy, starting from $s_0 = 30$. Under the safety only policy, the boiler is cleaned at least every other day (which clearly does not make sense from a practical standpoint) whereas the lexicographic policy cleans the boiler once, at the first day, and does not clean it again.

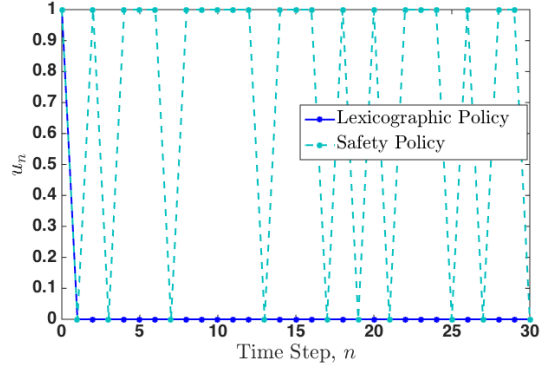


Figure 2: Simulated policy averaged over 1000 trials, starting from $s_0 = 30$, generated using Algorithm 2 (blue, solid line) and PBVI with a single safety objective (turquoise, dashed line). When costs are not considered, the boiler is cleaned more than half of the time, versus a single cleaning using the lexicographic policy.

6.2 Parsimonious Temperature Regulation

Switching from boiler maintenance, consider now the parsimonious temperature regulation problem introduced in Section 5.0.1.

To use Algorithm 2, we first abstract the system to a finite state POMDP by discretizing \mathcal{S} and \mathcal{Y} (for details see [?]). We only need to discretize the set K rather than all of \mathbb{R} , because once the temperature leaves K , we are not interested in its actual value. For the observations, we discretize $\bar{K} = [16, 23.5]$, because the probability of observing y outside of \bar{K} given $s \in K$ is $\epsilon \ll 1$ (specifically, $\epsilon = .0027$ for $\omega^2 = 25$). We use a grid spacing of $\delta^s = .01$ to create the finite set of states $\mathcal{S} = \{17.5, 17.51, \dots, 21.99\}$, and a grid spacing of $\delta^y = .25$ to create the finite set of observations $\mathcal{Y} = \{16, 16.25, \dots, 23.25\}$.

The probability distributions T and Y are constructed as for the boiler maintenance example, only now $T(s'|s, u)$ is derived from a Gaussian density function with mean $(1 - a)s + bu + as_a$ and variance ν^2 , and $Y(y|s)$ is derived from a Gaussian density with mean s and variance ω^2 . For instance, $T(s'|s, u) = \int_{s'}^{s'+\delta^s} \phi(x; (1 - a)s + bu + as_a, \nu^2) dx$ and similarly for Y (here $\phi(x; \mu, \sigma^2)$ is the Gaussian density with mean μ and variance σ^2 evaluated at x).

We compute the error bound δ^a on the abstraction using formulas derived in [14]. We only quantify the error from the state space discretisation; discretising the observations leads to an under-approximation of the safety probability (or over-approximation of the expected cost) similarly to the PBVI algorithm itself. As in the previous example, rather than

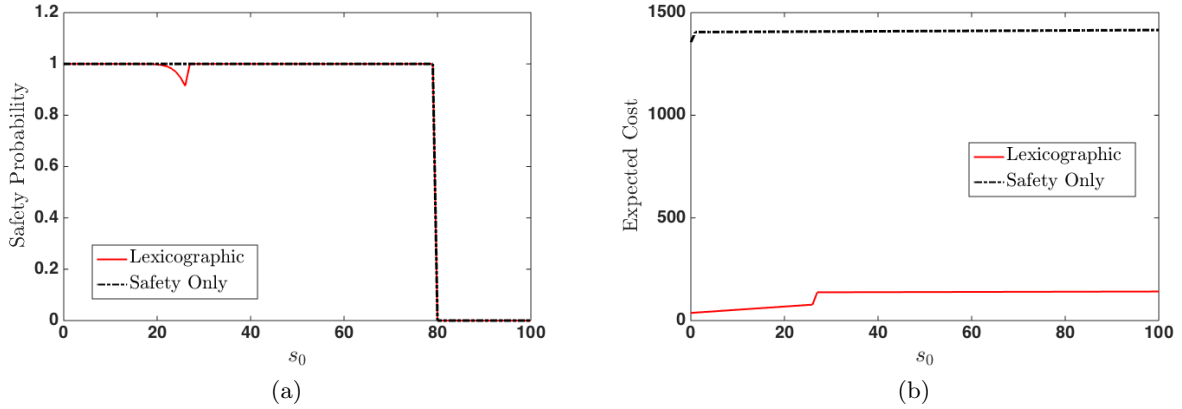


Figure 1: Comparison of safety probabilities (a) and expected total operational plus cleaning costs (b), over varying initial degradation levels s_0 and for $N = 30$ days, without considering the cost of fuel/cleaning (black, dashed line) and using Algorithm 2 to also minimize costs (red, solid line). The policy produced by maximizing safety alone leads to cleaning the boiler many times, whereas the lexicographic algorithm returns a policy that cleans the boiler at most once, with minimal effect on the probability of safety.

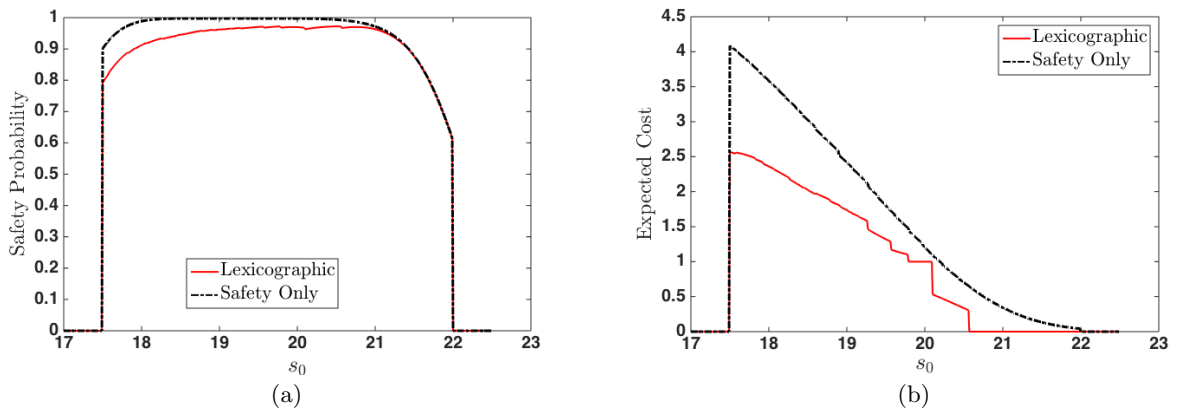


Figure 3: Comparison of safety probabilities (a) and expected number of time steps the heater is turned on (b), over varying initial temperatures s_0 and for $N = 5$ time steps, without considering the cost of turning the heater on (black, solid line) and using Algorithm 2 to also minimize the number of time steps the heater is on (blue, dashed line). Allowing the probability of remaining within the desired temperature range to drop slightly can lead to using the heater almost half as often.

incorporating the error from under-approximation explicitly, we rely on the fact that having a lower bound on the actual safety probability is desirable in itself.

For $\delta^s = 0.01$, and setting $\nu^2 = 0.2$, $\omega^2 = 0.25$, we obtain $\delta^a = 0.027$. For a time horizon of $N = 5$ time steps, and desired tolerance $\eta = 0.3$, the one step tolerance is $\bar{\eta} = 0.006$. We generate a set of 75 belief states in the same manner as the previous example, and apply Algorithm 2. Fig. 3a compares the maximal probability of safety (as underestimated by PBVI) to the under-estimated safety probability when using the lexicographic approach as a function of the initial temperature. Although the desired tolerance was made large in order to guarantee a non-zero one-step tolerance, Fig. 3a shows the actual distance from the maximal safety probability never exceeds 0.1 (because the error estimate δ^a is conservative). In fact, the safety probabilities produced by the lexicographic approach are quite close to the optimal, and even equal when the initial temperature is greater than 21, while the number of time steps the heater is on is lowered (cf. Fig. 3b). For an initial temperature between 17.5 and 18.5, the drop in likelihood of safety is around 0.1,

while the number of time steps the heater is likely to be on drops by more than one.

7. CONCLUSIONS AND FUTURE WORK

With focus on partially observable stochastic models, possibly with continuous and hybrid dynamics, this work has discussed the application of a lexicographic framework for multi-objective optimisation to problems with safety requirements, thus integrating correct-by-design synthesis for safety and optimal synthesis for performance. Computationally, the work has leveraged the use of compact representation of belief states, of sampling-base techniques (PBVI), and of formal abstractions, and showcased the results on two case studies in the area of buildings automation systems.

Methodologically, we are interested in exploring further connections with algorithms for approximate dynamic programming, and in furthering the usability of error bounds for formal abstractions [14,20]. In the area of modelling and control for smart buildings, we are interested in extensions to hybrid models of boiler degradation, and in developing and employing realistic models derived from real data.

APPENDIX

PROOF PROOF OF PROPOSITION 1. We first show by induction that $V_n^*(\bar{b}) - V_{n,1}^\eta(\bar{b}) \leq (N - n)\bar{\eta}$ for any $\bar{b} \in \bar{\mathfrak{B}}$, with V_n^* the value function that returns the true maximal safety probability, i.e. $V_0^*(\bar{R}) = p_{\text{safe}}^N(\pi^*, K; R)$. At time N , $V_N^*(\bar{b}) = \sum_{\bar{s}} \mathbf{1}_{K \times \{1\}}(\bar{s})\bar{b}(\bar{s})$, which does not depend on the control input, and so $V_N^*(\bar{b}) = V_{N,1}^\eta(\bar{b})$ for all \bar{b} , and the result is satisfied.

Next assume $V_{n+1}^*(\bar{b}) - V_{n+1,1}^\eta(\bar{b}) \leq (N - n - 1)\bar{\eta}$. At time n ,

$$V_n^*(\bar{b}) - V_{n,1}^\eta(\bar{b}) = V_n^*(\bar{b}) - V_{n,1}^{u^*}(\bar{b}) + V_{n,1}^{u^*}(\bar{b}) - V_{n,1}^\eta(\bar{b})$$

with $V_{n,1}^{u^*}$ the value function that uses the same optimal control input u^* at time n as V_n^* , but is defined recursively through the lexicographic value function $V_{n,1}^\eta$ (this is the same notation as used in the definition of $\mathcal{U}_{N,i}(b)$ (22)). Then,

$$\begin{aligned} V_n^*(\bar{b}) - V_{n+1}^\eta(\bar{b}) &\leq V_n^*(\bar{b}) - V_{n,1}^{u^*}(\bar{b}) + \bar{\eta} \\ &\leq \sum_y V_{n+1}^*(\bar{M}_{y,u^*}\bar{b})\mathbb{P}(y|\bar{b}, u^*) \\ &\quad - \sum_y V_{n+1,1}^\eta(\bar{M}_{y,u^*}\bar{b})\mathbb{P}(y|\bar{b}, u^*) + \bar{\eta} \\ &\leq (N - n - 1)\bar{\eta} \sum_y \mathbb{P}(y|\bar{b}, u^*) + \bar{\eta} \\ &\leq (N - n)\bar{\eta}. \end{aligned}$$

The first line follows from the definition of the one step tolerance (see (22)), the second to last line by the induction hypothesis, and the last line because $\sum_y \mathbb{P}(y|\bar{b}, u) = 1$. Since $V_0^*(\bar{b}) - V_{0,1}^\eta(\bar{b}) \leq N\bar{\eta}$ for all \bar{b} , we can achieve $V_0^*(\bar{b}) - V_{0,1}^\eta(\bar{b}) \leq \eta$ by setting $\bar{\eta} \leq \frac{\eta}{N}$. ■

PROOF PROOF OF PROPOSITION 2. First, because the PBVI algorithm is designed to produce a value function that is a lower bound on the actual value function, $V_{0,1}^{\eta,B} \leq V_{0,1}^\eta$. Further, by (32), we know that $V_0^*(\bar{b}) \leq N\delta^B + V_0^B(\bar{b})$. Then,

$$\begin{aligned} V_0^*(\bar{R}) - V_{0,1}^\eta(\bar{R}) &\leq V_0^*(\bar{R}) - V_{0,1}^{\eta,B} \\ &\leq V_0^B(\bar{R}) + N\delta^B - V_{0,1}^{\eta,B}(\bar{R}). \end{aligned}$$

Next, because we are enforcing the one step tolerance $\bar{\eta}$ between $\max_{u \in \mathcal{U}} V_{n,1}^{\eta,B}(\bar{b})$ and $V_{n,1}^\eta(\bar{b})$, we can use the same argument as in the proof of Prop. 1 to show that $V_0^*(\bar{R}) - V_{0,1}^{\eta,B}(\bar{R}) \leq N\bar{\eta}$. Therefore

$$V_0^*(\bar{R}) - V_{0,1}^{\eta,B}(\bar{R}) \leq N\delta^B + N\bar{\eta}$$

and we can guarantee that $V_0^*(\bar{R}) - V_{0,1}^{\eta,B}(\bar{R}) \leq \eta$ by setting $\bar{\eta} \leq \frac{\eta}{N} - \delta^B$. Intuitively, as the error from PBVI decreases to zero, the one step tolerance approaches that required in the case of having an exact solution. ■

PROOF PROOF OF PROPOSITION 3. As in the proof of Proposition 2, recall that $V_{n,1}^{\eta,B} \leq V_{n,1}^\eta$ and $V_0^*(\bar{b}) \leq N\delta^B + V_0^B(\bar{b})$. Also note that when we enforce the one step tolerance, we are enforcing it between $\max_{u \in \mathcal{U}} V_{n,1}^{\eta,B}$ and $V_{n,1}^{\eta,\delta}$, and therefore by the same argument as for Proposition 1, $V_{0,\delta}^B(\bar{R}) -$

$V_{0,1}^{\eta,B}(\bar{R}) \leq N\bar{\eta}$. Then,

$$\begin{aligned} V_0^*(\bar{R}) - V_{0,1}^\eta(\bar{R}) &\leq V_0^B(\bar{R}) + N\delta^B - V_{0,1}^{\eta,B}(\bar{R}) \\ &\leq V_0^B(\bar{R}) + N\delta^B - V_{0,\delta}^B(\bar{R}) + V_{0,\delta}^B(\bar{R}) \\ &\quad - V_{0,1,\delta}^{\eta,B}(\bar{R}) + V_{0,1,\delta}^{\eta,B}(\bar{R}) - V_{0,1}^{\eta,B}(\bar{R}) \\ &\leq \left| V_0^B(\bar{R}) - V_{0,\delta}^B(\bar{R}) \right| + V_{0,\delta}^B(\bar{R}) - V_{0,1,\delta}^{\eta,B}(\bar{R}) \\ &\quad + \left| V_{0,1,\delta}^{\eta,B}(\bar{R}) - V_{0,1}^{\eta,B}(\bar{R}) \right| + N\delta^B \\ &\leq \delta^a + N\bar{\eta} + \delta^a + N\delta^B \end{aligned}$$

Hence $V_0^*(\bar{R}) - V_{0,1}^\eta(\bar{R}) \leq \eta$ if we can enforce $\bar{\eta} \leq \frac{\eta - 2\delta^a}{N} - \delta^B$, which is true if $\frac{\eta - 2\delta^a}{N} - \delta^B \geq 0$. ■

A. REFERENCES

- [1] Studies show: Hvac system maintenance saves energy, 2012.
- [2] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008.
- [3] J. Berka and K. Macek. Effective maintenance of stochastic systems via dynamic programming. In *Proceedings of 19th Technical Computing Prague Conference*, Prague, Czech Republic, 2011.
- [4] D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, 2005.
- [5] K. Chatterjee, M. Chmelik, and M. Tracol. What is decidable about partially observable Markov decision processes with omega-regular objectives. In *Computer Science Logic 2013 (CSL 2013)*, *CSL 2013, September 2-5, 2013, Torino, Italy*, pages 165–180, 2013.
- [6] T. Chen, V. Forejt, M. Z. Kwiatkowska, A. Simaitis, and C. Wiltsche. On stochastic games with multiple objectives. In *Mathematical Foundations of Computer Science 2013 - 38th International Symposium, MFCS 2013, Klosterneuburg, Austria, August 26-30, 2013*, pages 266–277, 2013.
- [7] J. Ding, A. Abate, and C. Tomlin. Optimal control of partially observable discrete time stochastic hybrid systems for safety specifications. In *American Control Conference*, pages 6231–6236, 2013.
- [8] J. Ding and C. Tomlin. Robust reach-avoid controller synthesis for switched nonlinear systems. In *IEEE Conference on Decision and Control*, pages 6481–6486, 2010.
- [9] K. Dräger, V. Forejt, M. Z. Kwiatkowska, D. Parker, and M. Ujma. Permissive controller synthesis for probabilistic systems. In *Tools and Algorithms for the Construction and Analysis of Systems - 20th International Conference, TACAS 2014*, pages 531–546, 2014.
- [10] J. Dupačová and K. Sladký. Comparison of multistage stochastic programs with recourse and stochastic dynamic programs with discrete time. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, 82(11-12):753–765, 2002.
- [11] S. Giro and M. N. Rabe. Verification of partial-information probabilistic systems using counterexample-guided refinements. In *Proc. on*

Automated Technology for Verification and Analysis, LNCS, pages 333–348. Springer, 2012.

- [12] W. Kuijper and J. van de Pol. Compositional control synthesis for partially observable systems. In M. Bravetti and G. Zavattaro, editors, *CONCUR 2009 - Concurrency Theory*, volume 5710 of *Lecture Notes in Computer Science*, pages 431–447. Springer Berlin Heidelberg, 2009.
- [13] K. Lesser and M. Oishi. Reachability for partially observable discrete time stochastic hybrid systems. *Automatica*, 50(8):1989–1998, 2014.
- [14] K. Lesser and M. Oishi. Finite state approximation for verification of partially observable stochastic hybrid systems. In *Hybrid Systems: Computation and Control*, 2015.
- [15] L. G. Mitten. Preference order dynamic programming. *Management Science*, 21(1):43–46, 1974.
- [16] R. Mobley. *An Introduction to Predictive Maintenance*. Plant Engineering. Elsevier Science, 2002.
- [17] J. Pineau, G. Gordon, and S. Thrun. Anytime point-based approximations for large POMDPs. *Journal of Artificial Intelligence Research*, 27:335–380, 2006.
- [18] G. Shani, J. Pineau, and R. Kaplow. A survey of point-based POMDP solvers. *Autonomous Agents and Multi-Agent Systems*, 27(1):1–51, 2013.
- [19] E. Sondik. *The optimal control of partially observable Markov processes*. PhD thesis, Stanford University, 1971.
- [20] S. Soudjani and A. Abate. Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems*, 12(2):921–956, 2013.
- [21] M. Svorenova, I. Cerna, and C. Belta. Optimal control of MDPs with temporal logic constraints. In *Proceedings of the 52nd IEEE Conference on Decision and Control, CDC 2013, December 10-13, 2013, Firenze, Italy*, pages 3938–3943, 2013.
- [22] K. H. Wray and S. Zilberstein. Multi-objective POMDPs with lexicographic reward preferences. In *International Joint Conference on Artificial Intelligence*, 2015. To Appear.
- [23] K. H. Wray, S. Zilberstein, and A.-I. Mouaddib. Multi-objective MDPs with conditional lexicographic reward preferences. In *International Conference on Artificial Intelligence*, pages 3418–3424, 2015.
- [24] X. Yin and S. Lafortune. Synthesis of maximally permissive non-blocking supervisors for partially observed discrete event systems. In *IEEE Conference on Decision and Control*, pages 5156–5162, 2014.
- [25] L. Zhang, H. Hermanns, and D. N. Jansen. Logic and model checking for hidden Markov models. In *Proc. on Formal Techniques for Networked and Distributed Systems*, pages 98–112. Springer, 2005.