

Optimizing Fixed-Size Stochastic Controllers for POMDPs and Decentralized POMDPs

Christopher Amato
Daniel S. Bernstein
Shlomo Zilberstein

*Department of Computer Science, University of Massachusetts,
Amherst, MA 01003 USA*

CAMATO@CS.UMASS.EDU
BERN@CS.UMASS.EDU
SHLOMO@CS.UMASS.EDU

Abstract

POMDPs and their decentralized multiagent counterparts, DEC-POMDPs, offer a rich framework for sequential decision making under uncertainty. Their computational complexity, however, presents an important research challenge. One approach that effectively addresses the intractable memory requirements of current algorithms is based on representing agent policies as finite-state controllers. In this paper, we propose a new approach that uses this representation and formulates the problem as a nonlinear program (NLP). The NLP defines an optimal policy of a desired size for each agent. This new representation allows a wide range of powerful nonlinear programming algorithms to be used to solve POMDPs and DEC-POMDPs. Although solving the NLP optimally is often intractable, the results we obtain using an off-the-shelf optimization method are competitive with state-of-the-art POMDP algorithms and outperform state-of-the-art DEC-POMDP algorithms. Our approach is easy to implement and it opens up promising research directions for solving POMDPs and DEC-POMDPs using nonlinear programming methods.

1. Introduction

Developing effective frameworks for reasoning under uncertainty is a thriving research area in artificial intelligence. Rapid progress in recent years has contributed to the development of several new solution techniques. Nevertheless, improving the scalability of existing algorithms is still an important challenge. When a single decision maker is considered, the partially observable Markov decision process (POMDP) offers a powerful and widely used framework. The POMDP model offers a rich language to describe single agent planning in domains that involve stochastic actions, noisy observations, and a variety of possible objective functions. POMDP applications include robot control (Simmons & Koenig, 1995), medical diagnosis (Hauskrecht & Fraser, 1998) and machine maintenance (Eckles, 1968). In the area of robot control, partial observability is used to model sensors that provide uncertain and incomplete information about the state of the environment. In a medical setting, the internal state of the patient is often not known with certainty. The machine maintenance problem – one of the earliest application areas of POMDPs – seeks to find a cost-effective strategy for inspection and replacement of parts in a domain where partial information about the internal state is obtained by inspecting the manufactured products. Numerous other POMDP applications are surveyed in (Cassandra, 1998b).

Developing effective algorithms for POMDPs is an active research area that has seen significant progress (Cassandra, 1998a; Hansen, 1998; Ji, Parr, Li, Liao, & Carin, 2007;

Littman, Cassandra, & Kaelbling, 1995; Meuleau, Kim, Kaelbling, & Cassandra, 1999; Pineau, Gordon, & Thrun, 2003; Poupart, 2005; Poupart & Boutilier, 2003; Smith & Simmons, 2004, 2005; Spaan & Vlassis, 2005). Despite this progress, it is generally accepted that exact solution techniques are limited to toy problems due to high memory requirements. Consequently, approximation algorithms are often necessary in practice. Approximation methods perform well in many problem domains, but have some known disadvantages. For instance, point-based methods (Ji et al., 2007; Pineau et al., 2003; Spaan & Vlassis, 2005; Smith & Simmons, 2004, 2005) perform better with a small reachable belief space. Other algorithms rely on local search methods based on gradient ascent (Meuleau et al., 1999) and linear programming (Poupart & Boutilier, 2003) or require fine tuning of heuristic parameters to produce high-valued solutions (Poupart, 2005).

Solving problems in which a group of agents must operate collaboratively in some environment based solely on local information is also an important challenge. When the agents must rely on uncertain partial knowledge about the environment, this is particularly difficult. This problem can be modeled as a decentralized partially observable Markov decision process (DEC-POMDP), which is a multiagent extension of the POMDP framework. In a DEC-POMDP each agent must cope with uncertainty about the other agents as well as imperfect information of the system state. The agents seek to optimize a shared objective function using solely local information in order to act.

Many real world applications in cooperative multiagent systems can be modeled as DEC-POMDPs. Some examples include distributed robot control (Becker, Zilberstein, Lesser, & Goldman, 2004; Emery-Montemerlo, Gordon, Schneider, & Thrun, 2004) and networking problems (Bernstein, Hansen, & Zilberstein, 2005; Hansen, Bernstein, & Zilberstein, 2004). In multiagent domains, robot sensors not only provide uncertain and incomplete information about their own state, but also about the location of the other robots. This lack of information leads to different perspectives on the state of the environment and complicates the planning process. Similarly, in a decentralized network, each node must make decisions about when and where to send packets without full awareness of the knowledge and actions of nodes in the rest of the network. Other applications of DEC-POMDPs include e-commerce and coordination of space exploration systems.

Although there has been some recent work on exact and approximate algorithms for DEC-POMDPs (Bernstein et al., 2005; Hansen et al., 2004; Nair, Pynadath, Yokoo, Tambe, & Marsella, 2003; Petrik & Zilberstein, 2007; Seuken & Zilberstein, 2007b; Szer, Charpillet, & Zilberstein, 2005; Szer & Charpillet, 2005), only two algorithms (Bernstein et al., 2005; Szer & Charpillet, 2005) are able to find solutions for the infinite-horizon discounted case. Another algorithm can solve a restricted subclass of infinite-horizon DEC-POMDPs, using an average reward formulation rather than the standard discounted objective (Petrik & Zilberstein, 2007). The infinite-horizon problem is particularly suitable for such domains as networking and some robot control problems, where the agents operate continuously. The two general algorithms for infinite-horizon problems seek to optimize fixed-size finite-state controllers. These approaches perform well in some problems, but Szer and Charpillet's algorithm (2005) is limited to very small solutions because of high memory and running time requirements. Bernstein *et al.*'s method (2005) is more scalable, but it often produces low-quality results. It is worth noting that many POMDP algorithms cannot be easily extended to apply to DEC-POMDPs. One reason for this is that the decentralized nature of

the DEC-POMDP framework results in a lack of a shared belief state, making it impossible to properly estimate the state of the system. Thus, a DEC-POMDP cannot be formulated as a continuous state MDP and algorithms that use this formulation do not generalize.

In this paper, we explore a new solution technique for POMDPs and DEC-POMDPs that optimizes fixed-size controllers. Our approach formulates the optimal memory-bounded solution as a nonlinear program (NLP), and exploits existing nonlinear optimization techniques to solve the problem. Nonlinear programming is an active field of research that has produced a wide range of techniques that can efficiently solve a variety of large problems (Bertsekas, 2004). In the POMDP case, parameters are optimized for a fixed-size controller which produces the policy for that problem. In the DEC-POMDP case, a set of fixed-size independent controllers is optimized, which when combined, produce the policy for the DEC-POMDP. The formulation provides a new framework for which future algorithms can be developed. We present an overview of how to solve these problems optimally, but as this would often be intractable in practice, we also evaluate an effective approximation technique using standard NLP solvers. Our approach facilitates scalability as it offers a tradeoff between solution quality and the available resources. That is, for a given amount of memory, we can search for a controller that is optimal for that size. Large controllers may be needed to provide a near optimal solution in some problems, but our experiments suggest that smaller controllers produce high quality solutions in a wide array of problems.

In order to increase the solution quality of our memory-bounded technique for DEC-POMDPs, we examine the benefits of introducing a *correlation device*, which is a shared source of randomness. This allows a set of independent controllers to be correlated in order to produce higher values, without sharing any additional local information. Correlation adds another mechanism in an effort to gain the most possible value with a fixed amount of space. This has been shown to be useful in order to increase value of fixed-size controllers (Bernstein et al., 2005) and we show that is also useful when combined with our NLP approach.

The rest of the paper is organized as follows. We first give an overview of the POMDP and DEC-POMDP models and explain how infinite-horizon solutions can be represented as stochastic controllers. We discuss some previous work on each model and then show how to define optimal controllers using a nonlinear program. We then provide experimental results using a generic off-the-shelf NLP solver in various POMDP and DEC-POMDP benchmark domains. For POMDP problems, our approach is competitive with other state-of-the-art solution methods in terms of solution quality and running time, but it generally produces higher quality when given less memory. When applied to DEC-POMDPs, our approach proves to be more valuable. In general, optimizing our NLP formulation produces controllers that outperform other DEC-POMDP approaches in a range of domains. In both cases, very concise representations are found. Our results show that by using the NLP formulation, small, high-valued controllers can be efficiently found for a large assortment of POMDPs and DEC-POMDPs. Furthermore, performance is likely to improve as more specialized and scalable NLP solvers are developed that can take advantage of the structure of the controller optimization problem.

2. Background

As the DEC-POMDP model is an extension of the POMDP model, we begin with an overview of POMDPs and then show how this framework can be extended to cooperative multiagent environments modeled as DEC-POMDPs.

2.1 The POMDP model

A POMDP can be defined with the following tuple: $M = \langle S, A, P, R, \Omega, O \rangle$, with

- S , a finite set of states with designated initial state distribution b_0
- A , a finite set of actions
- P , the state transition model: $P(s'|s, a)$ is the probability of transitioning to state s' if action a is taken in state s
- R , the reward model: $R(s, a)$ is the immediate reward for taking action a in state s
- Ω , a finite set of observations
- O , the observation model: $O(o|s', a)$ is the probability of observing o if action a is taken and the resulting state is s'

We consider the case in which the decision making process unfolds over an infinite sequence of stages. At each stage the agent selects an action, which yields an immediate reward and an observation. The agent must choose an action based on the history of observations seen. Note that because the state is not directly observed, it may be beneficial for the agent to remember the observation history. This will help the agent to identify the set of possible states at any step. The objective of the agent is to maximize the expected discounted sum of rewards received. Because we consider the infinite sequence problem, we use a discount factor, $0 \leq \gamma < 1$, to maintain finite sums.

Finite-state controllers can be used as an elegant way of representing POMDP policies using a finite amount of memory (Hansen, 1998). The state of the controller is based on the observation sequence seen, and in turn the agent's actions are based on the state of its controller. To help distinguish states of the finite-state controller from states of the POMDP, we will refer to controller states as nodes. These controllers address one of the main causes of intractability in POMDP exact algorithms by not storing whole observation histories. Thus states of the controller can encapsulate key information about the observation history in a fixed number of nodes. We also allow for stochastic transitions and action selection, as this can help to make up for limited memory (Singh, Jaakkola, & Jordan, 1994). We provide an example showing this phenomenon in a previous paper (Amato, Bernstein, & Zilberstein, 2007). The finite-state controller can formally be defined by the tuple $\langle Q, \psi, \eta \rangle$, where Q is the finite set of controller nodes, $\psi : Q \rightarrow \Delta A$ is the action selection model for each node, mapping nodes to distributions over actions, and $\eta : Q \times A \times O \rightarrow \Delta Q$ represents the node transition model, mapping nodes, actions and observations to distributions over the resulting nodes. The value of a node q at state s , given action selection and node transition probabilities $P(a|q)$ and $P(q'|q, a, o)$, is given by the following Bellman equation:

$$V(q, s) = \sum_a P(a|q) \left[R(s, a) + \gamma \sum_{s'} P(s'|s, a) \sum_o O(o|s', a) \sum_{q'} P(q'|q, a, o) V(q', s') \right]$$

2.2 The DEC-POMDP model

A DEC-POMDP involves multiple agents that operate under uncertainty based on different streams of observations. Like a POMDP, the infinite-horizon DEC-POMDP unfolds over an infinite sequence of stages. At each stage, every agent chooses an action based purely on its local observations, resulting in an immediate reward for the set of agents and an observation for each individual agent. Again, because the state is not directly observed, it may be beneficial for each agent to remember its observation history. A *local policy* for an agent is a mapping from local observation histories to actions while a *joint policy* is a set of policies, one for each agent in the problem. Like the POMDP case, the goal is to maximize the infinite-horizon total cumulative reward, beginning at some initial distribution over states b_0 . Again, a discount factor, γ , is used to maintain finite sums.

More formally, a DEC-POMDP can be defined by the tuple: $\langle I, S, \{A_i\}, P, R, \{\Omega_i\}, O \rangle$ where I is the finite set of agents, S is again the finite set of states, A_i and Ω_i are the finite sets of actions and observations for each agent and P , R and O are now defined over the set of agents. That is, P is the set of state transition probabilities: $P(s'|s, \vec{a})$, the probability of transitioning from state s to s' when actions \vec{a} are taken by the set of agents, R is the reward function: $R(s, \vec{a})$, the immediate reward for being in state s and agents taking actions \vec{a} , and O is the set of observation probabilities: $O(\vec{o}|s', \vec{a})$, the probability of agents seeing the set of observations \vec{o} given the set of actions \vec{a} has been taken which results in state s' .

Finite state controllers also provide an appealing way to model DEC-POMDP policies with finite memory. Each agent's local policy can be represented as a local controller and the resulting set of controllers supply the joint policy, called the joint controller. These local controllers are defined in the same way as the single agent controllers discussed above. The joint controller has a defined initial node for each local controller and transitions occur in the local controllers based on the action taken and the observation seen by the corresponding agent. The value for starting in nodes \vec{q} and at state s with action selection and node transition probabilities for each agent, i , is given by the following Bellman equation:

$$V(\vec{q}, s) = \sum_{\vec{a}} \prod_i P(a_i|q_i) \left[R(s, \vec{a}) + \gamma \sum_{s'} P(s'|\vec{a}, s) \sum_{\vec{o}} O(\vec{o}|s', \vec{a}) \sum_{\vec{q}'} \prod_i P(q_i|q_i, a_i, o_i) V(\vec{q}', s') \right]$$

Note that the values can be calculated offline in order to determine a controller for each agent that can then be executed online in a decentralized manner.

3. Previous work

A range of solution techniques for infinite-horizon POMDPs and DEC-POMDPs already exist. We first review POMDP controller-based approaches, as that is the focus of our work. These techniques seek to determine the best action selection and node transition parameters for a fixed-size controller. We also mention some other POMDP approximation methods that are based on improving the value function for a set of belief points (distributions over the state space). We then review the current infinite-horizon DEC-POMDP approximate algorithms and discuss the shortcomings of the current single and multiagent approximate methods.

<p>For a given node q and variables x_a and $x_{q',a,o}$</p> <p>Maximize ϵ, subject to</p> <p>Improvement constraints:</p> $\forall s V(q, s) + \epsilon \leq \sum_a \left[x_a R(s, a) + \gamma \sum_{s'} P(s' s, a) \sum_o O(o s', a) \sum_{q'} x_{q',a,o} V(q', s') \right]$ <p>Probability constraints:</p> $\sum_a x_a = 1, \quad \forall a \sum_{q'} x_{q',a,o} = x_a, \quad \forall a x_a \geq 0, \quad \forall q', a, o x_{q',a,o} \geq 0$

Table 1: The linear program used by BPI. Variables x_a and $x_{q',a,o}$ represent $P(a|q)$ and $P(q', a|q, o)$ for a given node, q .

3.1 POMDP approaches

Policy Iteration (PI) (Hansen, 1998) is a technique to find epsilon-optimal POMDP controllers that alternates between policy improvement and evaluation. Although PI was originally developed for deterministic controllers, it can be extended to stochastic ones as well. In the improvement phase, dynamic programming is used to enlarge the controller. Also referred to as full backup, this entails creating a node for each possible combination of action and transition for each observation to any of the nodes in the current controller. Thus a controller with $|Q|$ nodes before the backup has $|A||Q|^{|\Omega|}$ nodes afterwards. The controller is then evaluated and nodes that have lesser or equal value for all states are replaced by nodes that (pointwise) dominate them. The incoming edges of these nodes are redirected to the dominating nodes, guaranteeing at least equal value. The new controller is then evaluated to determine the updated value for each node and state. This process continues until the controller is no longer changed by the improvement phase. While this method guarantees that a controller arbitrarily close to optimal will be found an intractable amount of memory is often required. Even when a high quality controller is found, it may be very large with many unnecessary nodes generated along the way. This is exacerbated by the fact that the algorithm cannot take advantage of an initial state distribution and must attempt to improve the controller for any initial state.

Poupart and Boutilier (Poupart & Boutilier, 2003) developed a method called bounded policy iteration (BPI) that uses a one step dynamic programming lookahead to attempt to improve a POMDP controller without increasing its size. Like PI, this approach also alternates between policy improvement and evaluation. It iterates through the nodes in the controller and uses a linear program, shown in Table 1, to examine the value of probabilistically taking an action and then transitioning into the current controller. If an improvement can be found for all states, the action selection and node transition probabilities are updated accordingly. The controller is then evaluated and the cycle continues until no further improvement can be found. BPI guarantees to at least maintain the value of a provided controller, but it also does not use start state information resulting in larger than necessary controllers that are not likely to be optimal.

To increase the performance of BPI, two improvements were added by Poupart and Boutilier. First, they allowed the controller to grow by adding nodes when improvements

can no longer be made. This may permit a higher value to be achieved than what BPI can produce with a fixed-size controller. Second, they used a heuristic for incorporating start state knowledge and increasing the performance of BPI in practice (Poupart, 2005). In this extension, termed biased BPI, improvement is concentrated in certain node and state pairs by weighing each pair by the (unnormalized) occupancy distribution, which can be found by solving the following set of linear equations:

$$o(q', s') = bp_0(q', s') + \gamma \sum_{q, s, a, o} o(q, s) P(s'|s, a) P(a|q) O(o|s', a) P(q'|q, a, o)$$

for all states and nodes. The value bp_0 is the probability of beginning in a node state pair. A factor, δ , can also be included, which allows the value to decrease by that amount in each node and state pair. This makes changes to the parameters more likely, as a small amount of value can now be lost. As a result, value may be higher for the start state and node, but as value can be lost for any pair, it could wind up being lower instead.

Meuleau et al. (1999) have proposed another approach to improve a fixed-size controller. They used gradient ascent (GA) to change the action selection and node transition probabilities and increase value. A cross-product MDP is created from the controller and the POMDP by considering the states of the MDP to be all combinations of the states of the POMDP and the nodes of the controller while actions of the MDP are based on actions of the POMDP and deterministic transitions in the controller after an observation is seen. The value of the resulting MDP can be determined and matrix operations allow the gradient to be calculated. The gradient can then be followed in an attempt to improve the controller.

Other approximation approaches that have recently gained a high level of popularity are point-based methods such as point-based value iteration (PBVI) (Pineau et al., 2003), PERSEUS (Spaan & Vlassis, 2005), heuristic search value iteration (HSVI) (Smith & Simmons, 2004) and point-based policy iteration (PBPI) (Ji et al., 2007). These techniques are approximations of value iteration for finite-horizon POMDPs, but they can find solutions for sufficiently large horizons that they can effectively produce infinite-horizon policies.

Value iteration is an optimal dynamic programming algorithm for finding solutions to POMDPs (Sondik, 1971). This is done by transforming the POMDP into a continuous state $|S|$ -dimensional “belief state” MDP and building up the value function from the last step until the first. Policies, starting with just the action set on the last step, are backed up exhaustively by considering for each action and any observation that is seen taking any current policy from the next step on. A policy can then be pruned if for any belief state, there is some other policy with higher value. This procedure of backing up policies and pruning them continues until a desired finite-horizon is reached or for the infinite-horizon case, value does not change more than a given epsilon. The value of a policy can be represented as a vector in $|S|$ -dimensional space. The optimal value function for any given horizon is then the max over these vectors resulting in a piecewise linear and convex value function. The infinite-horizon value function can be approximated within any ϵ using a finite set of vectors.

PBVI, PERSEUS, HSVI and PBPI fix the number of belief points considered at each step of value iteration. That is, at each step, the value vector of each belief point under consideration is backed up and the vector that has the highest value at that point is retained. This permits the number of vectors to remain constant and the value function is still defined

over the entire belief space. If the right points are chosen, it may help concentrate the value iteration on certain regions in the belief space. PBVI selects the set of belief points by sampling from the action and observation sets of the problem while PERSEUS backs up only a randomly selected subset of these points. Thus, PERSEUS is often faster and produces a more concise solution than PBVI. PBVI has been shown to be optimal for sufficiently large and well constructed belief point sets, but this is intractable for reasonably sized problems. HSVI chooses points based on an additional upper bound that is incorporated into the algorithm. This algorithm will also converge to the optimal solution, but again this often requires an intractable amount of resources. PBPI uses Hansen’s PI, but with PBVI rather than exact improvement over the whole belief space in the improvement phase. This allows faster performance and sometimes higher values than the original PBVI.

3.2 DEC-POMDP approaches

As mentioned above, the two algorithms that can solve infinite-horizon discounted DEC-POMDPs are those of Bernstein et al. (2005) and Szer and Charpillet (2005). Bernstein et al.’s approach, called bounded policy iteration for decentralized POMDPs (DEC-BPI), is a multiagent extension of Poupart and Boutilier’s BPI algorithm. Like BPI, it is an approximate algorithm that seeks to improve stochastic finite-state controllers. Szer and Charpillet describe an approach that searches for optimal fixed-size deterministic controllers.

DEC-BPI, like BPI, improves a set of fixed-size controllers by using linear programming. This is done by iterating through the nodes of each agent’s controller and attempting to find an improvement. The algorithm and linear program are very similar to those of BPI except the optimization is done over not just the states of the problem, but also the nodes of the other agents’ controllers. That is, the linear program searches for a probability distribution over actions and transitions into the agent’s current controller that increases the value of the controller for any initial state and any initial node of the other agents’ controllers. If an improvement is discovered, the node is updated based on the probability distributions found. Each node for each agent is examined in turn and the algorithm terminates when no agent can make any further improvements.

Szer and Charpillet have developed a best-first search algorithm that finds deterministic finite-state controllers of a fixed size. This is done by calculating an approximate value for the controller given the current known deterministic parameters and filling in the remaining parameters one at a time in a best-first fashion. The authors prove that this technique will find the optimal deterministic finite-state controller of a given size.

3.3 Disadvantages of previous approximate algorithms

Because the optimal algorithms are not very practical, we focus on approximate methods. Our new approach addresses several deficiencies of existing techniques. We discuss these limitations, first considering POMDP methods and then DEC-POMDP approaches.

The linear program used by BPI may allow for controller improvement, but local maxima are likely to be found since it performs only a one step lookahead while keeping the controller values fixed. Also, because the improvement must be over all states it can be difficult to find parameters that meet this requirement. Adding the heuristics used in biased BPI reduce this necessity by optimizing over weighted states as well as allowing value to be lost. The

disadvantage of biased BPI is that the heuristics are not useful in all problems and for cases when they are, domain knowledge or a large amount of experimentation may be necessary to properly set them.

Gradient ascent has similar problems with local maxima as well other concerns. Meuleau et al. must construct a cross-product MDP from the controller and the underlying POMDP in a complex procedure to calculate the gradient. Also, the authors' representation does not take into account the probability constraints and thus does not calculate the true gradient of the problem. Due to this complex and incomplete gradient calculation, GA can be time consuming and error prone. Techniques more advanced than gradient ascent may be used to traverse the gradient, but these shortcomings remain.

Point-based methods perform very well on many problems, but rely on choosing points that are highly representative of the belief space. If there is a large amount of reachable belief space and the values of different areas of that space are diverse, many points would be needed to reasonably approximate the value. As more points are used, the algorithms become more similar to value iteration and thus become less tractable.

DEC-POMDP algorithms also have significant drawbacks. As DEC-BPI is an extension of BPI, its disadvantages are similar to those stated above. In general, it allows memory to remain fixed, but provides only a locally optimal solution. This is due to the linear program considering the old controller values from the second step on and the fact that improvement must be over all possible states and initial nodes for the controllers of the other agents. As the number of agents or size of controllers grows, this later drawback is likely to severely hinder improvement. Also, start state knowledge is not used in DEC-BPI. While a heuristic like that used in biased BPI could be added, its applicability often depends on the domain and a more principled solution is desirable.

Szer and Charpillet's approach is limited due to the fact that it searches for optimal deterministic controllers. Since it focuses on deterministic controllers, larger controller sizes may often be required in order to produce values comparable to those generated by a stochastic counterpart. Since the approach relies on search to find an optimal set of controllers of a fixed size, it has a very high space and time requirements. These drawbacks indicate that a new approach is needed to improve the scalability and solution quality of POMDP and DEC-POMDP approximate methods.

4. Optimizing fixed-size controllers

Unlike other controller based approaches for POMDPs and DEC-POMDPs, our formulation defines an optimal controller for a given size. This is done by creating a set of new variables that represent the values of each node (or set of nodes in the DEC-POMDP case) and state pair. Intuitively, this allows changes in the controller probabilities to be reflected in the values of the nodes of the controller. This is in contrast to backups used by other methods which iteratively improve the probabilities and easily get stuck in local optima. Our approach allows both the values and probabilities in the controller to be optimized in one step, thus representing the optimal fixed-size controller. To ensure that these values are correct given the action selection and node transition probabilities, nonlinear constraints (based on the Bellman equation) must be added. The resulting NLP is generally harder to solve, but many robust and efficient algorithms can be applied.

One premise of our work is that an optimal formulation of the problem facilitates the design of solution techniques that can overcome the limitations of previous controller-based algorithms and produce better stochastic controllers. The general nature of our formulation allows a wide range of solution methods to be used. This results in a search that is more sophisticated than previously used in controller-based methods. While no existing technique guarantees global optimality in a finite amount of time, experimental results show that our new formulation is advantageous. For instance, our results suggest that many POMDPs and DEC-POMDPs have small optimal controllers or can be approximated concisely with finite state-controllers. Thus, it is often unnecessary to use a large amount of memory in order to represent a good approximation. Our NLP is also able to take advantage of the start distribution of the problem, thus making better use of limited controller size. Lastly, because our method searches for stochastic controllers, it is able to find higher-valued, more concise controllers than search in the space of deterministic controllers.

Compared to point-based approaches, our formulation does not need to choose a set of points and may be able to cover the belief space better in some problems. That is, while point-based methods work well when there is a small reachable belief space or when the chosen points are very representative of the whole space, the NLP approach seeks to optimize the value of a controller for a specific initial point. For domains in which point-based methods cannot find representative belief points, our approach may still be able to construct high quality controllers. Also, since point-based methods rely on finite-horizon dynamic programming, it may be difficult for these methods to complete the number of backups necessary to approximate the value function well. As our approach uses finite-state controllers, it is more suitable for finding infinite-horizon policies.

In the rest of this section, we give a formal description of the NLP and prove that its optimal solution defines an optimal controller of a fixed size. We begin with the POMDP formulation and then show the extension to the DEC-POMDP case.

4.1 NLP formulation for POMDPs

Unlike BPI, which alternates between policy improvement and evaluation, our nonlinear program improves and evaluates the controller in one phase. The value of an initial node is maximized at an initial state distribution using parameters for the action selection probabilities at each node $P(a|q)$, the node transition probabilities $P(q'|q, a, o)$, and the values of each node in each state $V(q, s)$. To ensure that the value variables are correct given the action and node transition probabilities, nonlinear constraints must be added to the optimization. These constraints are the Bellman equations given the policy determined by the action selection and node transition probabilities. Linear constraints are used to maintain proper probabilities.

To reduce the representation complexity, the action selection and node transition probabilities are merged into one, with

$$P(q', a|q, o) = P(a|q)P(q'|q, a, o) \text{ and } \sum_{q'} P(q', a|q, o) = P(a|q)$$

This results in a quadratically constrained linear program. QCLPs may contain quadratic terms in the constraints, but have a linear objective function. They are a subclass of general nonlinear programs that has structure which algorithms can exploit. This produces

<p>For variables: $x(q', a, q, o)$ and $z(q, s)$</p> <p>Maximize $\sum_s b_0(s)z(q_0, s)$, subject to</p> <p>The Bellman constraints:</p> $\forall q, s \ z(q, s) =$ $\sum_a \left[\left(\sum_{q'} x(q', a, q, o) \right) R(s, a) + \gamma \sum_{s'} P(s' s, a) \sum_o O(o s', a) \sum_{q'} x(q', a, q, o) z(q', s') \right]$ <p>Probability constraints:</p> $\forall q, o \ \sum_{q', a} x(q', a, q, o) = 1, \quad \forall q', a, q, o \ x(q', a, q, o) \geq 0,$ $\forall q, o, a \ \sum_{q'} x(q', a, q, o) = \sum_{q'} x(q', a, q, o_k)$

Table 2: The NLP defining an optimal fixed-size POMDP controller. Variable $x(q', a, q, o)$ represents $P(q', a|q, o)$, variable $z(q, s)$ represents $V(q, s)$, q_0 is the initial controller node and o_k is an arbitrary fixed observation.

a problem that is often more difficult than a linear program, but simpler than a general nonlinear program. The QCLP formulation permits a large number of algorithms to be applied. Because the QCLP is a subclass of the general NLP and to use the same terminology as the DEC-POMDP case, we will refer to the QCLP formulation as an NLP.

Table 2 describes the NLP which defines an optimal fixed-size controller for the POMDP case. The value of a designated initial node is maximized given the initial state distribution and the necessary constraints. The first constraint represents the Bellman equation for each node and state which maintains correct values as probability parameters change. The second and third constraints ensure that the x variables represent proper probabilities and the last constraint guarantees that action selection does not depend on the resulting observation which has not yet been seen.

Theorem 1 *An optimal solution of the NLP in Table 2 results in an optimal stochastic controller for the given size and initial state distribution.*

Proof The optimality of the controller follows from the one-to-one correspondence between the objective function of the NLP and the value of the POMDP at the initial state distribution. The Bellman equation constraints restrict the value variables to be consistent with the chosen action selection and transition probabilities. The remaining constraints guarantee that the action selection and transition probabilities are selected from valid distributions. Hence, the optimal solution of this NLP represents the value of a fixed-size controller that is optimal for the given POMDP. \square

<p>For variables of each agent i: $x(q_i, a_i)$, $y(q_i, a_i, o_i, q'_i)$ and $z(\vec{q}, s)$</p> <p>Maximize $\sum_s b_0(s)z(q^{\vec{0}}, s)$, subject to</p> <p>The Bellman constraints:</p> $\forall \vec{q}, s \quad z(\vec{q}, s) =$ $\sum_{\vec{a}} \prod_i x(q_i, a_i) \left[R(s, \vec{a}) + \gamma \sum_{s'} P(s' s, \vec{a}) \sum_{\vec{o}} O(\vec{o} s', \vec{a}) \sum_{\vec{q}'} \prod_i y(q_i, a_i, o_i, q'_i) z(\vec{q}', s') \right]$ <p>And probability constraints for each agent i:</p> $\forall q_i \quad \sum_{a_i} x(q_i, a_i) = 1, \quad \forall q_i, o_i, a_i \quad \sum_{q'_i} y(q_i, a_i, o_i, q'_i) = 1$ $\forall q_i, a_i \quad x(q_i, a_i) \geq 0, \quad \forall q_i, o_i, a_i \quad y(q_i, a_i, o_i, q'_i) \geq 0$

Table 3: The NLP defining a set of optimal fixed-size DEC-POMDP controllers. For each agent i , variable $x(q_i, a_i)$ represents $P(a_i|q_i)$, variable $y(q_i, a_i, o_i, q'_i)$ represents $P(q'_i|q_i, a_i, o_i)$ and variable $z(\vec{q}, s)$ represents $V(\vec{q}, s)$ where $q^{\vec{0}}$ represents the initial controller node for each agent.

4.2 NLP formulation for DEC-POMDPs

The nonlinear program for DEC-POMDPs is an extension of the one defined above which incorporates multiple controllers. That is, the NLP seeks to optimize the value of a set of fixed-size controllers given an initial state distribution and the DEC-POMDP model. The variables for this problem are the action selection and node transition probabilities for each node of each agent’s controller as well as the joint value of a set of controller nodes. Hence, these variables are for each agent i , $P(a_i|q_i)$ and $P(q'_i|q_i, a_i, o_i)$ and for the set of agents and any state, $V(\vec{q}, s)$. Similar to the way our POMDP approach differs from BPI, this approach differs from DEC-BPI in that it explicitly represents the node values as variables, thus allowing improvement and evaluation to take place simultaneously. To ensure that the values are correct given the action and node transition probabilities, nonlinear constraints must be added to the optimization. These constraints are the DEC-POMDP Bellman equations given the policy determined by the action and transition probabilities. We must also ensure that the necessary variables are valid probabilities. This representation is shown formally for an arbitrary number of agents in Table 3.

Theorem 2 *An optimal solution of the NLP in Table 3 results in optimal stochastic controllers for the given size and initial state distribution.*

Proof Like the POMDP case, the optimality of the controllers follows from the NLP constraints and maximization of given initial nodes at the initial state distribution. The Bellman equation constraints restrict the value variables to valid amounts based on the

chosen probabilities and because the probability variables are independent, decentralized controller is ensured. The maximum value is found given these restrictions for the initial nodes and state distribution. Hence, optimal solution of this NLP represents optimal fixed-size controllers. \square

5. Incorporating correlation in the decentralized model

As an improvement to DEC-BPI, Bernstein et al. also allow the agents’ controllers to be correlated by using a shared source of randomness in the form of a *correlation device*. They show that higher-valued controllers can sometimes be achieved when incorporating this approach. As an example of a correlation device, imagine that before each action is taken, a coin is flipped and both agents have access to the outcome. Each agent can then use the new information to affect their choice of action. Along with stochasticity, correlation is another means of increasing value when memory is limited. For instance, assume the agents’ actions are randomized so that each chooses action A 50% of the time and action B 50% of the time. This results in the joint actions of (A, A) , (A, B) , (B, A) and (B, B) each 25% of the time. When a coin is added, the agents can correlate their actions based on what is seen, such as allowing each agent to perform action A when “heads” is seen and action B when “tails” is seen. As each signal occurs 50% of the time, this permits joint actions (A, A) and (B, B) to also take place 50% of the time. This type of policy is impossible without the correlation device showing that the policies can be randomized and correlated to allow higher value to be attained.

It is worth noting that there are many ways a correlation device can be implemented in practice. These implementations include each agent using a random number generator with a common seed, or agreeing before execution starts on a list of random numbers. In each of these examples, each agent has access to the random signal at each step, but no local information is shared and thus the execution remains decentralized.

More formally, a correlation device provides extra signals to the agents and operates independently of the DEC-POMDP. That is, the correlation device is a tuple $\langle C, \psi \rangle$, where C is a set of states and $\psi : C \rightarrow \Delta C$ is a stochastic transition function that we will represent as $P(c'|c)$. At each step of the problem, the device transitions and each agent can observe its state.

The independent local controllers defined above (in Table 3) can be modified to make use of a correlation device. This is done by making the parameters dependent on the signal from the correlation device. For agent i , action selection is then $P(a_i|q_i, c)$ and node transition is $P(q'_i|q_i, a_i, o_i, c)$. For n agents, the value of the correlated joint controller beginning in nodes \vec{q} , state s and correlation device state c is defined as

$$V(\vec{q}, s, c) = \sum_{\vec{a}} \prod_i P(a_i|q_i, c) \cdot \left[R(s, \vec{a}) + \gamma \sum_{s'} P(s'|\vec{a}, s) \sum_{\vec{o}} O(\vec{o}|s', \vec{a}) \sum_{\vec{q}'} \prod_i P(q_i|q_i, a_i, o_i, c) \sum_{c'} P(c'|c) V(\vec{q}', s', c') \right]$$

Our NLP formulation can be extended to include a correlation device. The optimization problem, shown in Table 4 *without* the probability constraints, is very similar to the previous

<p>For variables of each agent i: $x(q_i, a_i, c)$, $y(q_i, a_i, o_i, q'_i, c)$, $z(\vec{q}, s, c)$ and $w(c, c')$</p> <p>Maximize $\sum_s b_0(s)z(q^0, s, c)$, subject to</p> <p>The Bellman constraints:</p> $\forall \vec{q}, s \quad z(\vec{q}, s, c) = \sum_{\vec{a}} \prod_i x(q_i, a_i, c).$ $\left[R(s, \vec{a}) + \gamma \sum_{s'} P(s' s, \vec{a}) \sum_{\vec{o}} O(\vec{o} s', \vec{a}) \sum_{\vec{q}'} \prod_i y(q_i, a_i, o_i, q'_i, c) \sum_{c'} w(c, c')z(\vec{q}', s', c) \right]$

Table 4: The nonlinear program for including a correlation device in our DEC-POMDP NLP formulation. For each agent i , variable $x(q_i, a_i, c)$ represents $P(a_i|q_i, c)$, variable $y(q_i, a_i, o_i, q'_i, c)$ represents $P(q'_i|q_i, a_i, o_i, c)$, variable $z(\vec{q}, s, c)$ represents $V(\vec{q}, s, c)$, q^0 represents the initial controller node for each agent and $w(c, c')$ represents $P(c'|c)$. The other constraints are similar to those above with the addition of a sum to one constraint for the correlation device.

NLP. A new variable is added for the transition function of the correlation device and the other variables now include the signal from the device. The Bellman equation incorporates the new correlation device signal at each step, but the other constraints remain the same. A new probability constraint is added to ensure that the transition probabilities for each state of the correlation device sum to one.

6. Methods for solving the NLP

Many efficient constrained optimization algorithms can be used to solve large NLPs. Constrained optimization seeks to minimize or maximize an objective function based on equality and inequality constraints. When the objective and all constraints are linear, this is called a linear program (LP). As our POMDP formulation has a linear objective function, but contains some quadratic constraints, it is a quadratically constrained linear program. The DEC-POMDP representation also has a linear objective, but has some nonlinear constraints whose degree depends on the number of agents in the problem. Unfortunately, both of these problems are nonconvex. Essentially, this means that there may be multiple local maxima as well as global maxima, thus finding globally optimal solution is often very difficult. In the next two subsections, we discuss a method to find these globally optimal solutions as well as a more practical method for finding locally optimal solutions.

6.1 Finding a globally optimal solution for our NLP

One method that could be used to find optimal solutions for our NLP is *difference of convex functions* (DC) programming (Horst & Tuy, 1996). DC programming uses a general formulation that allows nonlinear and nonconvex constraints and is studied within the field

of global optimization. This formulation requires an objective function and constraints that are a difference of two convex functions. Under mild conditions, it is possible to devise DC programming algorithms that converge to the globally optimal solution.

More formally, a function f is called DC if it is real-valued and defined over a convex set $C \subset \mathbb{R}^n$ and $\forall x \in C$, f can be decomposed such that $f(x) = p(x) - q(x)$ where p and q are convex functions on C . A *DC programming problem* then has the form

$$\min f(x) \text{ such that } x \in C \text{ and } g_i(x) \leq 0$$

where $C \subset \mathbb{R}^n$ is convex and objective function f and each constraint g_i are DC on C .

Proposition 1 *The POMDP and DEC-POMDP NLP formulations can be defined as DC programming problems.*

Proof First, the maximization in our formulations can be changed to a minimization by multiplying by a factor of -1, and our probability and value bounds create a convex region over which the variables are defined. Also, we can add slack variables to the equality constraints in order to match the inequality constraints of the DC programming formulation. Our linear objective function and linear constraints in both the POMDP and DEC-POMDP case are trivially DC, but we must check the nonlinear constraints. Horst and Tuy show that any function whose second partial derivatives are continuous is also DC (corollary I.1). It is simple (but lengthy) to show that the nonlinear constraints in both our POMDP and DEC-POMDP formulations are twice differentiable and these partial derivatives are continuous. Thus, the variables are defined over a convex set and the objective and constraints are DC in both formulations. This allows our NLP representations to be written as DC programming problems. \square

Therefore, POMDPs and DEC-POMDPs can be formulated as DC programming problems. This has mostly theoretical significance; we doubt that this formulation could be effective in practice. In principle, global optimization algorithms such as branch and bound can be used, but because of the large size of the resulting DC programming problem, it is unlikely that current optimal solvers can handle even small POMDPs or DEC-POMDPs. Nevertheless, it would be interesting to identify classes of these problems for which DC optimization is practical.

6.2 Finding a locally optimal solution for our NLP

Since it may not be possible or feasible to solve the NLP optimally, locally optimal methods are often more useful in practice. A wide range of nonlinear programming algorithms have been developed that are able to efficiently solve nonconvex problems with many variables and constraints. Locally optimal solutions can be guaranteed, but at times, globally optimal solutions can also be found. For example, merit functions, which evaluate a current solution based on fitness criteria, can be used to improve convergence and the problem space can be made convex by approximation or domain information. These methods are much more robust than simpler methods such as gradient ascent, while retaining modest efficiency in many cases.

For this paper, we used a freely available nonlinearly constrained optimization solver called *snopt* (Gill, Murray, & Saunders, 2005) on the NEOS server (www-neos.mcs.anl.gov). The algorithm finds solutions by a method of successive approximations called sequential quadratic programming (SQP). SQP uses quadratic approximations which are then solved with quadratic programming (QP) until a solution to the more general problem is found. A QP is typically easier to solve, but must have a quadratic objective function and linear constraints. In *snopt*, the objective and constraints are combined and approximated to produce the QP. A merit function is also used to guarantee convergence from any initial point.

We are also exploring other techniques that take advantage of the structure of POMDPs and DEC-POMDPs in order to increase the scalability of certain solution methods. One promising idea is to use constraint partitioning (Wah & Chen, 2005) to break up a problem and solve the resulting pieces in such a way that allows a feasible solution to be found. For this approach, we are exploring different methods for breaking up our representations and more efficient ways of solving the subproblems. We are also examining ways to simplify our nonlinear program formulations and determine solution techniques that match these representations well, thereby increasing performance.

7. Experiments

For experimental comparison, we present an evaluation of the performance of our formulation using an off-the-shelf nonlinear program solver, *snopt*, as well as leading POMDP and DEC-POMDP approximation techniques. In these experiments we seek to determine how well our formulation performs when used in conjunction with a generic solver such as *snopt*. The formulation is very general and many other solvers may be applied. As mentioned above, we are currently developing a customized solver that would take further advantage of the inherent structure of the NLPs and increase scalability. While the POMDP results are intended to show that a simple solution of the NLP is competitive with current solution methods, the DEC-POMDP results show consistent improvement in solution quality over previous approaches.

7.1 POMDP results

In this section, we compare the results obtained using our new formulation with those of biased BPI, PBVI, PERSEUS, HSVI and PBPI. GA was also implemented, but produced significantly worse results and required substantially more time than the other techniques. Thus we omit the details of GA and focus on the more competitive techniques.

To improve scalability, we also use a version of our NLP formulation in which the action at each node is fixed. As the transitions between the nodes remain stochastic, this only reduces the expressiveness of the controller by a small amount due to the first action being chosen deterministically. While we could have allowed this first action to be stochastic, it was found that choosing this action greedily based on immediate reward at the initial belief often performed reasonably well. To set the other actions of the controller, we cycled through the available actions and assigned the next available action to the next node. That is, if a problem has 4 actions and 9 nodes then each action is represented twice in the controller except for the greedy first action which is represented 3 times. Given a fixed

controller size, fixing the action selection will usually perform worse than allowing stochastic actions, but the problem is simpler and thus larger controller sizes can be solved. When there were multiple highest-valued actions, one of these actions is randomly assigned to the first node. The hope is that although fixing actions will result in less concise controllers, higher-valued, larger controllers may be found. Throughout this section, we will refer to the optimization of our NLP using snopt as *NLO* for the fully stochastic case and as *NLO fixed* for the optimization with fixed actions.

Each of our NLP methods was initialized with ten random deterministic controllers and we report mean values and times after convergence. To slightly increase the performance of the solver upper and lower bounds were added. These represent the value of taking the highest and lowest valued actions respectively for an infinite number of steps. All results were found by using the NEOS server which provides a set of machines with varying CPU speeds and memory limitations.

BENCHMARK PROBLEMS

We first provide a comparison of our NLP formulations with leading POMDP approximation methods on common benchmark problems. The first three domains, which were introduced by Littman, Cassandra and Kaelbling (Littman et al., 1995), are grid problems in which an agent must navigate to a goal square. The set of actions consists of turning left, right or around, moving forward and staying in place. The observations consist of the different views of the walls in the each domain based on each agent’s ability to observe walls in four directions and a unique observation for the goal state. There are also three additional observable landmarks in the Hallway problem. Both the observations and transitions are extremely noisy and the start state is a random non-goal state. Only the goal has a reward of 1 and the discount factor used was 0.95. Following the convention of previously published results, we consider the versions of these problems that stop after the goal has been reached.

The other benchmark problem is a larger grid problem in which the goal is for the agent to catch and tag an opponent which attempts to move away (Pineau et al., 2003). The tagging agent has perfect knowledge of its state, but it cannot observe the opponent unless it is located in the same square. The agent can deterministically move in each of four directions and can use the “tag” action. The opponent attempts to move away from the agent in a fixed way, but may stay in the same location with probability 0.2. If the agent succeeds in tagging the opponent, a reward of 10 is given, but if the opponent is not in the same square when the tag action is taken a reward of -10 is given. All other actions result in a -1 reward and like the other benchmarks, the problem stops once the goal is reached and a discount factor of 0.95 was used.

Table 5 shows the results from previously published algorithms and our NLP formulations. We provide the mean values and times for the largest controller size that is solvable with less than (approximately) 400MB of memory and under eight hours on the NEOS server. The values of PBPI in the table are the highest values reported for each problem, but the authors did not provide results for the Hallway problem. Because the experiments were conducted on different computers, solution times give a general idea of the speed of the algorithms. It is also worth noting that while most of the other approaches are highly optimized, a generic solver is used with our approach in these experiments.

Tiger-grid $ S = 36, A = 5, \Omega = 17$				Hallway $ S = 60, A = 5, \Omega = 21$			
	value	size	time		value	size	time
HSVI	2.35	4860	10341	PBVI	0.53	86	288
PERSEUS	2.34	134	104	HSVI2	0.52	147	2.4
HSVI2	2.30	1003	52	HSVI	0.52	1341	10836
PBVI	2.25	470	3448	biased BPI	0.51	43	185
PBPI	2.24	3101	51	PERSEUS	0.51	55	35
biased BPI	2.22	120	1000	BPI	0.51	1500	249730
NLO fixed	2.20	32	1823	NLO fixed	0.49	24	330
BPI	1.81	1500	163420	NLO	0.47	12	362
NLO	1.79	14	1174	Q_{MDP}	0.27	n.a.	1.34
Q_{MDP}	0.23	n.a.	2.76				

Hallway2 $ S = 93, A = 5, \Omega = 17$				Tag $ S = 870, A = 5, \Omega = 30$			
	value	size	time		value	size	time
PERSEUS	0.35	56	10	PBPI	-5.87	818	1133
HSVI2	0.35	114	1.5	PERSEUS	-6.17	280	1670
PBPI	0.35	320	3.1	HSVI2	-6.36	415	24
HSVI	0.35	1571	10010	HSVI	-6.37	1657	10113
PBVI	0.34	95	360	biased BPI	-6.65	17	250
biased BPI	0.32	60	790	BPI	-9.18	940	59772
NLO fixed	0.29	18	240	PBVI	-9.18	1334	180880
NLO	0.28	13	420	NLO fixed	-10.48	5	2117
BPI	0.28	1500	274280	NLO	-13.94	2	5596
Q_{MDP}	0.09	n.a.	2.23	Q_{MDP}	-16.9	n.a.	16.1

Table 5: Values, representation sizes and running times (in seconds) for the set of benchmark problems. Results for other algorithms were taken from the following sources: BPI (Poupart & Boutilier, 2003), biased BPI (Poupart, 2005), HSVI (Smith & Simmons, 2004), HSVI2 (Smith & Simmons, 2005), PERSEUS (Spaan & Vlassis, 2005), PBPI (Ji et al., 2007), PBVI (Pineau et al., 2003), Q_{MDP} (Spaan & Vlassis, 2005)

In general, we see that the nonlinear optimization approach is competitive both in running time and value produced, but does not outperform the other techniques. Our method achieves 94%, 92%, 83% of max value in the first three problems, but does so with much smaller representation size. A key aspect of the NLP approach is that high quality solutions can be found with very concise controllers. This allows limited representation size to be utilized very well, likely better than the other approaches in most problems.

The values in the table as well as those in Figure 1 show that our approach is currently unable to find solutions for large controller sizes. For example, in the Tag problem, a solution could only be found for a two node controller in the general case and a five node controller when the actions were fixed. As seen in the figures, there is a near monotonic increase of value in the Hallway problem when compared with either controller size or time.

As expected, fixing the actions at each node results in faster performance and increased scalability. This allows higher valued controllers to be produced, but scalability remains limited. As improved solvers are found higher values and larger controllers will be solvable by both approaches. The other benchmark problems display similar trends of near monotonic improvement and the need for increased scalability to outperform other methods on these problems.




Figure 1: Hallway problem values for various controllers sizes (left) and running times (right) using the NLP formulation with and without fixed actions at each node. Bars are used in the left figure to represent the minimum and maximum values for each size.

OTHER DOMAINS

We also examined the performance of our NLP approach and selected approximation approaches in two other domains, the machine maintenance and aloha problems (Cassandra, 1998a). For HSVI2 and PERSEUS, software was used from web sites of Trey Smith and Matthijs Spaan respectively. For biased BPI results, we used our own implementation and like our NLP results, the mean values and times of optimizing ten fixed-size deterministic controllers are reported. The heuristic parameters were found experimentally by choosing a wide range of values and reporting the best results. Similarly, for PERSEUS, several methods for choosing a set of belief points of size 10000 were used and again, the best results are reported.

The machine maintenance domain has 256 states, 4 actions and 16 observations. There are four independent components in a machine used to manufacture a part. Each component may be in good, fair or bad condition as well as broken and in need of replacement. At each step, four actions are possible. The machine can be used to manufacture parts or we can inspect, repair, or replace the machine. The manufacture action produces good parts based on the condition of the components. Good components always produce good parts,

and broken components always produce bad parts. Components in fair or bad condition raise the probability of producing bad parts. The condition of the resulting part is fully observed. Inspecting the machine causes a noisy observation of whether each component is good or bad. Components in good or fair condition are more likely to be seen as good, and those in fair or broken are more likely to be seen as bad. Repairing the machine causes parts that are not broken to improve by one condition with high probability. The replace action transitions all components to good condition and components may degrade each step unless repaired or replaced. Rewards for each action are: 1 for manufacturing good parts for the day, -1 for inspecting, -3 for repairing and -15 for producing bad parts. A discount factor of 0.99 was used.

The aloha problem is a simplified networking problem using the slotted Aloha protocol with 90 states, 30 actions and 3 observations. In this problem, time is divided into slots where the agent may receive a message and must decide what the probability that it will send one. If two transmitters send at the same time a collision takes place and the messages are backlogged. The number of messages that arrive at any given slot is determined by a Poisson distribution, but limited to at most 10. The probability of sending a message is discretized to $1/m$ with m being a positive integer less than M , the maximum number of backlogged messages. For this problem M was set to be 30. The agent cannot observe the number of messages in its queue only that the channel was idle, a message was transmitted or a collision took place. These observations are deterministic and the state of the system is a combination of the number of backlogged messages and the previous observation. Message sending success probabilities are defined for each action and rewards are given for each message under the maximum limit that the queue is currently at. A discount factor of 0.999 was used.

On these problems, which we provide results for in Table 6, our NLP formulation provides results that are nearly identical (over 99% of the value) on the machine problem (NLO fixed) and the aloha problem (NLO). Both of our methods also outperform our versions of biased BPI and PERSEUS in almost all evaluation criteria in both problems. While these two algorithms may be able to increase their performance by adjusting heuristics, we believe that our approach will continue to produce results that will match this performance without the need to adjust heuristics. One reason for this is that these problems are difficult for point-based approaches to solve optimally. This is due to the high discount factor causing many backups to be required and the need for many vectors to closely approximate the optimal value function. For instance, in the aloha domain, after 15 hours HSVI2 was able to produce a value of 1212.92 which required over 15000 vectors to represent. In contrast, while the mean value produced by our nonlinear optimization was 1211.67, a controller with value 1217.11 was found with only six nodes in about 20 minutes. This shows the value of our controller-based technique as a more efficient alternative to point-based methods in problems such as these.

7.2 DEC-POMDP results

We also tested our nonlinear programming approach in four DEC-POMDP domains. In each experiment, we compare our NLP solutions using snopt with Bernstein et al.’s DEC-BPI and Szer and Charpillet’s BFS for a range of controller sizes. We also implemented

Machine $ S = 256, A = 4, \Omega = 16$				Aloha 30 $ S = 90, A = 29, \Omega = 3$			
	value	size	time		value	size	time
HSVI2	63.17	575	317	HSVI2	1212.15	2909	1841
NLO fixed	62.65	20	3963	NLO	1211.67	6	1134
NLO	61.74	10	7350	NLO fixed	1076.49	26	2014
biased BPI	59.75	20	30831	biased BPI	993.22	22	5473
PERSEUS	39.28	86	2508	PERSEUS	853.24	114	2512

Table 6: Values, representation sizes and running times (in seconds) for the machine maintenance and aloha problems.

the NLP and DEC-BPI techniques with a correlation device of size two. Like the POMDP version, to improve scalability we also solved the NLP with fixed actions at each node. The assignments were made as described above with each action in turn fixed to each node and the set of actions which returns the highest immediate reward given the start state provides the first action of each controller. Also like the POMDP case, if an agent had multiple actions with equally high immediate values, a random one was chosen.

Each NLP and DEC-BPI algorithm was run until convergence was achieved with ten different random deterministic initial controllers, and the mean values and running times are reported. Unlike the figure in the POMDP section, the figures in this section do not include maximum and minimum value bars. While these could have been provided, they are very similar for each of our NLP methods and are omitted to increase readability. Also, we do not always show the results for the largest solvable controller once the values plateau. The optimal deterministic controller returned by BFS is also reported if it could be found in under eight hours. The times reported for each NLP method can only be considered estimates due to running each algorithm on external machines with uncontrollable load levels, but we expect that they vary by only a small constant factor. Throughout this section we will refer to the optimization of our NLP with snopt as *NLO*, the optimization with the correlation device with two states as *NLO corr*, and the these optimizations with fixed actions as *NLO fixed* and *NLO fixed corr*.

RECYCLING ROBOTS

As a first comparison, we have extended the Recycling Robot problem (Sutton & Barto, 1998) to the multiagent case. The robots have the task of picking up cans in an office building. They have sensors to find a can and motors to move around the office in order to look for cans. The robots are able to control a gripper arm to grasp each can and then place it in an on-board receptacle. Each robot has three high level actions: (1) search for a small can, (2) search for a large can or (3) recharge the battery. In our two agent version, the larger can is only retrievable if both robots pick it up at the same time. Each agent can decide to independently search for a small can or to attempt to cooperate in order to receive a larger reward. If only one agent chooses to retrieve the large can, no reward is given. For each agent that picks up a small can, a reward of 2 is given and if both agents cooperate

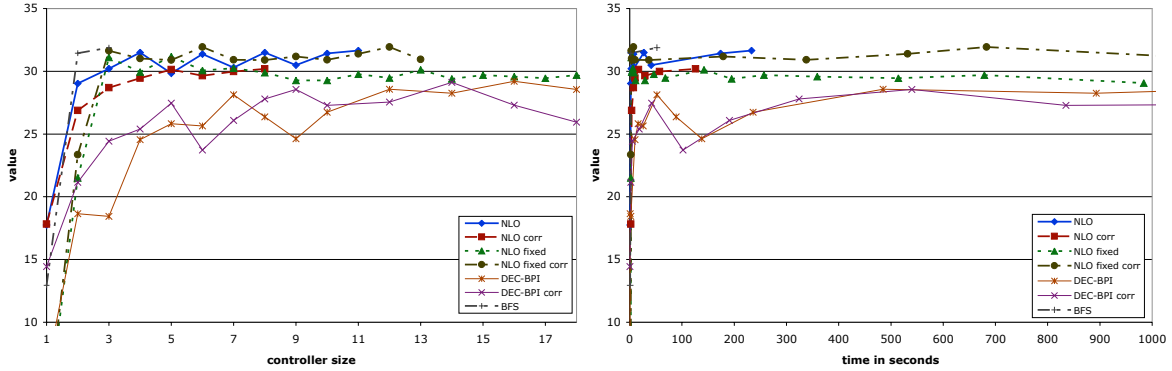


Figure 2: Recycling robots values for various controller sizes (left) and running times (right) using NLP methods and DEC-BPI with and without a 2 node correlation device and BFS

to pick the large can, a reward of 5 is given. The robots have the same battery states of high and low, with an increased likelihood of transitioning to a low state or exhausting the battery after attempting to pick up the large can. Each robot’s battery power depends only on its own actions and each agent can fully observe its own level, but not that of the other agent. If the robot exhausts the battery, it is picked up and plugged into the charger then continues to act on the next step with a high battery level. The two robot version used in this paper has 4 states, 3 actions and 2 observations. Also, a discount factor of 0.9 was used for all problems.

As seen in Figure 2, in this domain all the methods performed fairly well. While the DEC-BPI techniques converge over time to a value around 27, the other approaches converge to at least 30. BFS, our general formulation (NLO) and our formulation with fixed actions and a correlation device (NLO fixed corr) perform particularly well on this problem, converging to a value near 32 with very small controllers. We also see that these values were found quickly, often taking less than a minute for most methods. BFS could solve at most a controller of size three, but in this problem that is sufficient to produce a high-valued solution. We also see that while correlation helps in the case of fixed actions (NLO fixed corr), it does not improve solution quality in the general case (NLO corr). Fixing actions in our NLP approach (NLO fixed) allows larger controllers to be solved, but an increase in value is not found.

MULTIAGENT TIGER PROBLEM

Another domain with 2 states, 3 actions and 2 observations called the multiagent tiger problem was introduced by Nair et al. (2003). In this problem, there are two doors. Behind one door is a tiger and behind the other is a large treasure. Each agent may open one of the doors or listen. If either agent opens the door with the tiger behind it, a large penalty is given. If the door with the treasure behind it is opened and the tiger door is not, a reward is given. If both agents choose the same action (i.e., both opening the same door) a larger




Figure 3: Multiagent tiger problem values for various controller sizes (left) and running times (right) using NLP methods with and without a 2 node correlation device and BFS

positive reward or a smaller penalty is given to reward this cooperation. If an agent listens, a small penalty is given and an observation is seen that is a noisy indication of which door the tiger is behind. While listening does not change the location of the tiger, opening a door causes the tiger to be placed behind one of the door with equal probability.

In this problem, DEC-BPI performed very poorly (never producing values higher than -50) causing us to remove the values from Figure 3 for the sake of clarity. We can also see in the figures the limitation of BFS. While it can do very well for three nodes, it cannot find solutions for larger controllers. In contrast, the NLP approaches demonstrate a trend of higher-valued solutions as controller size increases. This suggests that as more scalable solution methods for our formulation are found, the increase in solution quality will persist. The NLP approaches also make better use of time, quickly producing high solution quality and then value increases more slowly. We also see that both correlation and fixing actions are helpful. In fact, the highest value produced by any method is given by a 19 node controller with fixed actions (NLO fixed) and the second highest value was produced by a seven node controller with a correlation device (NLO corr). While displaying an upward trend, the solutions for this problem, particularly with fixed actions, are somewhat irregular. This is likely due to the fact that there is a large penalty for opening the wrong door which can significantly lower the average presented in the figure. Thus, there is a large variance in solution quality but the high performance of the NLP methods is clear.

MEETING IN A GRID

A larger domain with 16 states, 5 actions and 2 observations was also introduced by Bernstein et al. In this problem, two agents must meet in a 2x2 grid with no obstacles. The agents start diagonally across from each other and available actions are move left, right, up, or down and stay in place. Only walls to the left and right can be observed, resulting in each agent knowing only if it is on the left or right half. The agents cannot observe each other and do not interfere with other. Action transitions are noisy with the agent possibly

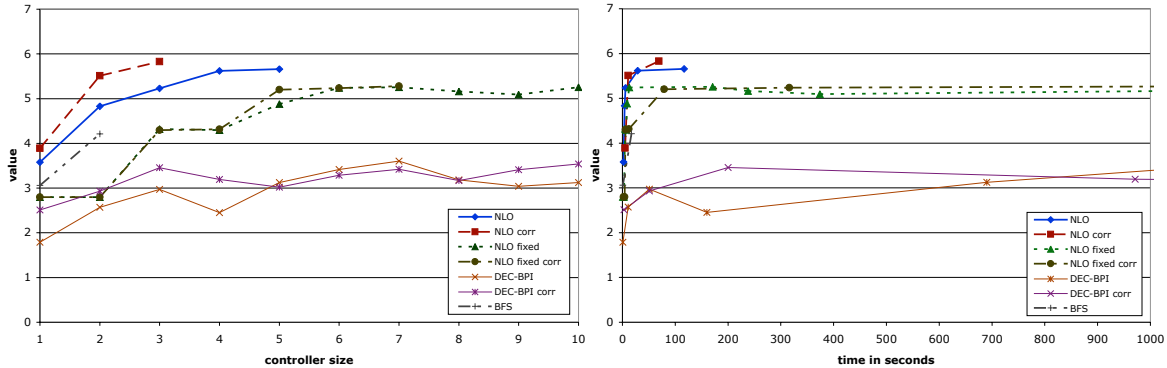


Figure 4: Meeting in a grid problem values for various controller sizes (left) and running times (right) using NLP methods and DEC-BPI with and without a 2 node correlation device and BFS

moving in another direction or staying in place and a reward of 1 is given for each step the agents share the same square.

The results for this example can be seen in Figure 4. The figures show that, in general, our NLP techniques and BFS outperform DEC-BPI. We also see that the standard NLO methods (NLO and NLO corr) perform especially well, producing higher values than the other methods for all solvable controller sizes. The fixed action NLP methods also perform well, but even though they can solve larger controller sizes than the standard approaches, the quality produced is not as high. It is likely that a greedy first action is not the best choice for this problem. The figure on the right shows that for all methods, a near maximum value was reached quickly by all methods (~ 1 minute) and then, if larger controllers could be solved, value increased very slightly thereafter. We also see that correlation increases solution quality for the general formulation of this problem (NLO corr), but not for the fixed action formulation (NLO fixed corr).

BOX PUSHING

A much larger domain was introduced by Seuken and Zilberstein (2007a). This problem, with 100 states, 4 actions and 5 observations consists of two agents that can gain reward by pushing different boxes. The agents begin facing each other in the bottom corners of a 4x3 grid with the available actions of turning right, turning left, moving forward or staying in place. There is a 0.9 probability that the agent will succeed in moving and otherwise will stay in place, but the two agents can never occupy the same square. The middle row of the grid contains two small boxes and one large box. This large box can only be moved by both agents pushing at the same time. The upper row of the grid is considered the goal row. The possible deterministic observations consist of seeing an empty space, a wall, the other agent, a small box or the large box. A reward of 100 is given if both agents push the large box to the goal row and 10 is given for each small box that is moved to the goal row. A

Value

size	NLO	NLO corr	NLO fixed	NLO fixed corr	DEC-BPI	DEC-BPI corr	BFS
1	-1.580	6.267	-2	-2	-10.367	-10.012	-2
2	31.971	39.825	19.660	19.657	3.293	4.486	x
3	46.282	50.834	44.916	44.579	9.442	12.504	x
4	50.640	x	54.230	48.623	7.894	x	x

Time

size	NLO	NLO corr	NLO fixed	NLO fixed corr	DEC-BPI	DEC-BPI corr	BFS
1	20	21	7	11	26	87	1696
2	115	314	15	27	579	3138	x
3	683	1948	128	214	4094	15041	x
4	5176	x	1824	1995	11324	x	x

Table 7: Box pushing problem values and running times (in seconds) using NLP methods and DEC-BPI with and without a 2 node correlation device and BFS

penalty of -5 is given for each agent that cannot move and -0.1 is given for each time step. Once a box is moved to the goal row, the environment resets to the original start state.

The results for this problem, which is an order of magnitude larger than previously published infinite-horizon problems, are given in Table 7. Because this problem is so large and given the time and memory limitations, BFS is only able to provide a solution for a one node controller and both DEC-BPI and our general approach with a correlation device (NLO corr) could not solve the formulation for a four node controller. Nonetheless, our techniques perform very well, never producing lower value than DEC-BPI or BFS and often producing much higher value. Fixing actions allows solutions to be found very quickly while still outperforming DEC-BPI for each controller size. The standard NLO approaches generally perform better, with NLO and NLO corr generating controllers with higher values than the other methods for all controller sizes except NLO fixed with four nodes. All NLP methods required much less time than the DEC-BPI techniques when the problems were solvable. Also, correlation provides a slight increase in value for the general formulation, but with fixed actions values are nearly identical for the first three controller sizes and slightly less for four node controllers.

8. Conclusions

In this paper, we introduced a new approach for solving POMDPs and DEC-POMDPs by defining optimal fixed-size solutions as nonlinear programs. This permits a wide range of powerful NLP solution methods to be applied to these problems. Because our approach optimizes a stochastic controller in the single agent case or a set of stochastic controllers for the multiagent case, it is able to make better use of the limited representation space than deterministic approaches. Thus, this new formulation may allow very concise high quality solutions to a large number POMDPs and DEC-POMDPs.

We showed that by using a generic nonlinear program solver, our formulation can provide POMDP solutions that are competitive with leading techniques in terms of running time and solution quality. Given the complexity of the problems that we considered, it is not surprising that no single approach dominates all the others. Nevertheless, our approach has some distinctive characteristics and advantages. In particular, it can produce good value with significantly smaller memory. And because the controller is more compact, in some cases the approach uses less time despite the use of nonlinear optimization. This work opens up new promising research directions that could produce further improvement in both quality and efficiency. As the solver in our experiments has not been optimized for our NLP it is likely that even higher performance can be achieved in the future. Also, as a controller-based approach, our method may be able to provide high quality solutions to some problems in which point-based methods have difficulty.

In the DEC-POMDP case, we demonstrated that our approach is able to produce solutions that are often significantly higher-valued than the other infinite-horizon DEC-POMDP algorithms for a range of problems. In all problems, a variant of our NLP approach produced the highest value that was generated for any controller size of any method. This included a problem that is an order of magnitude larger than any used previously in infinite-horizon DEC-POMDP literature. The combination of start state knowledge and more advanced optimization techniques allows us to make efficient use of the limited space of the controllers. In addition, incorporating a correlation device as a shared source of randomness can further increase solution quality. The running times needed to produce solutions for each problem also show that our NLP approach is frequently the most efficient method to achieve any given value. These results show that our approach can quickly find compact, high-valued solutions to be found for several different types of DEC-POMDPs. Lastly, as the number of agents in a given problem grows, the number of local maxima is likely to increase. Because solution of our NLP can avoid many of these suboptimal points, our representation should also perform well when the number of agents is increased beyond two.

In the future, we plan to explore more specialized algorithms that can be tailored for our optimization problem. While the performance achieved with a standard nonlinear optimization algorithm is good, specialized solvers might be able to further increase solution quality and scalability. Different methods may be able to take advantage of the specific structure inherent in POMDPs and DEC-POMDPs. Also, important subclasses could be identified for which globally optimal solutions can be efficiently provided. We also plan to characterize the circumstances under which introducing a correlation device is cost effective. Both of these improvements may allow optimal or near optimal fixed-size solutions to be found for important classes of POMDPs and DEC-POMDPs.

Acknowledgments

We would like to thank Marek Petrik for his helpful comments. Support for this work was provided in part by the National Science Foundation under Grant No. IIS-0535061 and by the Air Force Office of Scientific Research under Agreement No. FA9550-05-1-0254.

References

- Amato, C., Bernstein, D. S., & Zilberstein, S. (2007). Solving POMDPs using quadratically constrained linear programs. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pp. 2418–2424, Hyderabad, India.
- Becker, R., Zilberstein, S., Lesser, V., & Goldman, C. V. (2004). Solving transition-independent decentralized Markov decision processes. *Journal of AI Research*, 22, 423–455.
- Bernstein, D. S., Hansen, E., & Zilberstein, S. (2005). Bounded policy iteration for decentralized POMDPs. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pp. 1287–1292, Edinburgh, Scotland.
- Bertsekas, D. P. (2004). *Nonlinear Programming*. Athena Scientific.
- Cassandra, A. R. (1998a). *Exact and Approximate Algorithms for Partially Observable Markov Decision Processes*. Ph.D. thesis, Brown University, Providence, RI.
- Cassandra, A. R. (1998b). A survey of POMDP applications. In *AAAI Fall Symposium: Planning with POMDPs*, Orlando, FL.
- Eckles, J. E. (1968). Optimum maintenance with incomplete information. *Operations Research*, 16, 1058–1067.
- Emery-Montemerlo, R., Gordon, G., Schneider, J., & Thrun, S. (2004). Approximate solutions for partially observable stochastic games with common payoffs. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 136–143, New York, NY.
- Gill, P. E., Murray, W., & Saunders, M. (2005). Snopt: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47, 99–131.
- Hansen, E. A. (1998). Solving POMDPs by searching in policy space. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, pp. 211–219, Madison, WI.
- Hansen, E. A., Bernstein, D. S., & Zilberstein, S. (2004). Dynamic programming for partially observable stochastic games. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, pp. 709–715, San Jose, CA.
- Hauskrecht, M., & Fraser, H. (1998). Modeling treatment of ischemic heart disease with partially observable Markov decision processes. In *Proceedings of American Medical Informatics Association annual symposium on Computer Applications in Health Care*, pp. 538–542, Orlando, Florida.
- Horst, R., & Tuy, H. (1996). *Global Optimization: Deterministic Approaches*. Springer.
- Ji, S., Parr, R., Li, H., Liao, X., & Carin, L. (2007). Point-based policy iteration. In *Proceedings of the Twenty-Second National Conference on Artificial Intelligence*, Vancouver, Canada.
- Littman, M. L., Cassandra, A. R., & Kaelbling, L. P. (1995). Learning policies for partially observable environments: Scaling up. Tech. rep. CS-95-11, Brown University, Department of Computer Science, Providence, RI.

- Meuleau, N., Kim, K.-E., Kaelbling, L. P., & Cassandra, A. R. (1999). Solving POMDPs by searching the space of finite policies. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, Stockholm, Sweden.
- Nair, R., Pynadath, D., Yokoo, M., Tambe, M., & Marsella, S. (2003). Taming decentralized POMDPs: Towards efficient policy computation for multiagent settings. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pp. 705–711, Acapulco, Mexico.
- Petrik, M., & Zilberstein, S. (2007). Average-reward decentralized markov decision processes. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pp. 1997–2002, Hyderabad, India.
- Pineau, J., Gordon, G., & Thrun, S. (2003). Point-based value iteration: an anytime algorithm for POMDPs. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, Acapulco, Mexico.
- Poupart, P. (2005). *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes*. Ph.D. thesis, University of Toronto.
- Poupart, P., & Boutilier, C. (2003). Bounded finite state controllers. In *Advances in Neural Information Processing Systems*, 16, Vancouver, Canada.
- Seuken, S., & Zilberstein, S. (2007a). Improved memory-bounded dynamic programming for decentralized POMDPs. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, Vancouver, Canada.
- Seuken, S., & Zilberstein, S. (2007b). Memory-bounded dynamic programming for DEC-POMDPs. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pp. 2009–2015, Hyderabad, India.
- Simmons, R., & Koenig, S. (1995). Probabilistic navigation in partially observable environments. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, Montreal, Canada.
- Singh, S., Jaakkola, T., & Jordan, M. (1994). Learning without state-estimation in partially observable Markovian decision processes. In *Proceedings of the Eleventh International Conference on Machine Learning*, New Brunswick, NJ.
- Smith, T., & Simmons, R. (2004). Heuristic search value iteration for POMDPs. In *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence*, Banff, Canada.
- Smith, T., & Simmons, R. (2005). Point-based POMDP algorithms: Improved analysis and implementation. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, Edinburgh, Scotland.
- Sondik, E. J. (1971). *The Optimal Control of Partially Observable Markov Processes*. Ph.D. thesis, Stanford University.
- Spaan, M. T. J., & Vlassis, N. (2005). Perseus: Randomized point-based value iteration for POMDPs. *Journal of AI Research*, 24, 195–220.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press.

- Szer, D., & Charpillet, F. (2005). An optimal best-first search algorithm for solving infinite horizon DEC-POMDPs. In *Proceedings of the Sixteenth European Conference on Machine Learning*, Porto, Portugal.
- Szer, D., Charpillet, F., & Zilberstein, S. (2005). MAA*: A heuristic search algorithm for solving decentralized POMDPs. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, Edinburgh, Scotland.
- Wah, B. W., & Chen, Y. (2005). Solving large-scale nonlinear programming problems by constraint partitioning. In *Proceedings of the Eleventh International Conference on Principles and Practice of Constraint Programming*.