




A Semi-Decentralized Approach to Multiagent Control

Mahdi Al-Husseini 
Stanford University
Stanford, United States
mah9@stanford.edu

Mykel J. Kochenderfer 
Stanford University
Stanford, United States
mykel@stanford.edu

Kyle H. Wray 
Northeastern University
Boston, United States
k.wray@northeastern.edu




ABSTRACT

We introduce an expressive framework and algorithms for the semi-decentralized control of cooperative agents in environments with communication uncertainty. Whereas semi-Markov control admits a distribution over time for agent actions, semi-Markov communication, or what we refer to as semi-decentralization, gives a distribution over time for what actions and observations agents can store in their histories. We extend semi-decentralization to the partially observable Markov decision process (POMDP). The resulting SDec-POMDP unifies decentralized and multiagent POMDPs and several existing explicit communication mechanisms. We present recursive small-step semi-decentralized A* (RS-SDA*), an exact algorithm for generating optimal SDec-POMDP policies. RS-SDA* is evaluated on semi-decentralized versions of several standard benchmarks and a maritime medical evacuation scenario. This paper provides a well-defined theoretical foundation for exploring many classes of multiagent communication problems through the lens of semi-decentralization.

KEYWORDS

Semi-decentralization; Semi-Markov process; Multiagent communication; Admissible heuristic search; Planning

ACM Reference Format:

Mahdi Al-Husseini , Mykel J. Kochenderfer , and Kyle H. Wray . 2026. A Semi-Decentralized Approach to Multiagent Control. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/IFSK7108>

1 INTRODUCTION

Many complex real-world problems require the coordination of multiple cooperative agents to solve, but feature limited opportunities for information exchange. The decentralized partially observable Markov decision process (Dec-POMDP) formalizes multi-agent planning and control in settings where explicit communication is impossible [1, 4, 5]. Several model variants take advantage of existing, if limited, information structures and extend the Dec-POMDP to explicitly incorporate costly [10], delayed [19, 20], lossy [31], or intermittent [32] communication.

In this paper we pursue a general framework that unifies several multiagent communication mechanisms. We are interested in problems whose communication dynamics are conditioned with some

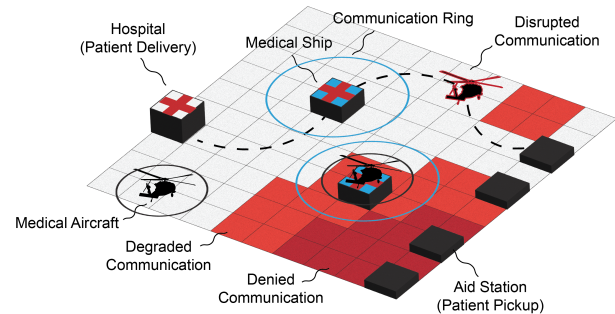


Figure 1: A semi-decentralized multiagent evacuation scenario with probabilistic restrictions on communication. Aircraft and watercraft must coordinate under communication constraints to move patients from aid stations to hospitals.

probability on the state, joint actions, or joint observations. An example domain is area-wide Global Positioning System (GPS) denial via jamming [11]. As seen in the evacuation scenario depicted in Figure 1, agents need to coordinate joint tasks in environments with degraded, denied, or disrupted communication. Agents must reason about which actions to take in light of available communication, the influence actions taken have on future communication, and future communication's influence on future actions. In decentralized systems, information may be action-gated in costly communication, constrained by channel capacity in lossy communication, temporally offset in intermittent or delayed communication, and either deterministic or stochastic in nature, while semi-decentralized systems generalize all of these through probabilistic information flow.

We begin by formally defining the *semi-decentralization property*. The key insight is that **semi-Markov systems for resolving agent control may be co-opted in semi-decentralized systems for agent communication**. Semi-decentralized systems may simultaneously contain a subset of agents that are necessarily decentralized and a subset of agents that are permissibly centralized. We introduce the semi-decentralized POMDP (SDec-POMDP), which unifies several existing multiagent models and explicit communication mechanisms. The SDec-POMDP reveals underlying mechanisms for memory propagation, selector functions, which we toggle here but are otherwise inherently present and constant in existing models. We then detail an exact optimal planning algorithm for solving SDec-POMDPs and apply to a set of new semi-decentralized benchmarks and a maritime medical evacuation application. Our contributions follow:



This work is licensed under a Creative Commons Attribution International 4.0 License.

- We formulate the semi-decentralization property by extending semi-Markov control concepts to communication. Semi-decentralization is then applied to multiagent POMDPs, forming the SDec-POMDP.
- We prove that the SDec-POMDP unifies the Dec-POMDP, the MPOMDP, k -steps delayed communication, and the Dec-POMDP-COM.
- We introduce Recursive Small-Step Semi-Decentralized A* (RS-SDA*), an exact algorithm for solving semi-decentralized problems. We evaluate performance in semi-decentralized variants of four Dec-POMDP benchmarks, then apply to a complex medical evacuation scenario with joint tasks.

Ultimately, we describe and validate the mechanisms of a novel yet foundational property and model for multiagent decision making in probabilistic communication environments.

2 RELATED WORK

Multiagent Models Centralized planning can occur when agents freely communicate in reliable networks with no latency or when individual agents possess full system observability. This allows a single planner to select joint actions over the global state [21]. A partially observable centralized multiagent system is represented by a multiagent partially observable Markov decision process (MPOMDP). MPOMDP planners benefit from conditioning joint actions on joint observations but scale poorly and are susceptible to communication failures. MPOMDPs therefore have limited application to many challenging real-world problems. By contrast, cooperative agents in decentralized systems can neither communicate explicitly nor observe the entire system state, and therefore must select their own actions in accordance with local observations [15]. The decentralized POMDP (Dec-POMDP) and MPOMDP share an underlying model but possess distinct policies and histories. Decentralized models have been used to support unmanned aerial vehicle formation flight [3], maritime traffic management [28], and wildfire surveillance [14]. Decentralized systems permit implicit communication, in which agents transmit information by means of taken actions and received observations [12]. Explicit communication, by contrast, endows agents with formalized communications actions.

The state-of-the-art exact optimal algorithm for solving Dec-POMDPs, recursive small-step multi-agent A* (RS-MAA*) [16], relies on a combination of incremental expansion, clustering, variable-depth recursive heuristics, and heuristic memoization. We extend RS-MAA* to semi-decentralized systems.

Communication Schemes The literature features numerous explicitly modeled communication mechanisms and frameworks, owing to the sheer diversity of information structures in real-world problems. Goldman and Zilberstein formalize the decentralized POMDP with communication (Dec-POMDP-Com), which incorporates an alphabet of possible messages and a communication cost function [9]. The Dec-POMDP-Com models costly communication and generally assumes noise-free instantaneous broadcasting. Delayed communication occurs when agents learn the local observation received by others after one or more time-steps [21]. In one-step delayed communication, agents determine the latest joint action from the joint policy, but are unaware of the latest joint observation [19]. In k -steps delayed communication, a sufficient statistic

may be used in place of the past joint policy [20]. Other multiagent models consider noisy [31] or intermittent [32] communication channels.

In the costly, delayed, noisy, and intermittent communication cases, the domain environment is orthogonal to the communication channel, and the actions taken by agents and resulting state do not affect their future ability to coordinate. Therefore, while communication directly influences control, control does not in turn directly influence communication. This distinguishes semi-decentralized infrastructure from existing communication schemes.

Communication in Reinforcement Learning Our primary contribution is a foundational model that supports both planning and reinforcement learning in restricted communication environments. We further present an exact algorithm for semi-decentralized planning, in which the communication and model dynamics are known prior to execution [17]. Still, the growing body of multiagent deep reinforcement learning with communication (Comm-MADRL) research is replete with promising techniques for codifying communication and defining communication policy. Zhu et al. categorize Comm-MADRL approaches by communication policy, to include full communication, partial communication, individual control, and global control subcategories [33]. We are chiefly concerned with the individual control literature [13, 26, 27], and seek to learn optimal control policies in light of potential communication links between agents. In the sub-field of *learning tasks with communication*, policies generated using learning algorithms simultaneously maximize environmental rewards and determine effective communication protocols for agents [33]. This parallels recent efforts in planning to design joint communication and control strategies [29], which we also do.

3 PRELIMINARIES

3.1 Dec-POMDPs

The *decentralized partially observable Markov decision process* (Dec-POMDP) is a stochastic, decentralized multiagent model for sequential decision-making under partial observability characterized by tuple $\langle I, S, \bar{A}, \bar{O}, T, O, R \rangle$, where:

- I is a finite set of k agents,
- S is a finite set of states,
- $\bar{A} = \times_i A_i$ is a finite set of joint actions,
- $\bar{O} = \times_i O_i$ is a finite set of joint observations, and
- $T : S \times \bar{A} \times S \rightarrow [0, 1]$ is a state transition function where $T(s' | s, \bar{a})$ is the probability of being in state s' given joint action \bar{a} being performed in state s ,
- $O : \bar{O} \times S \times \bar{A} \rightarrow [0, 1]$ is a joint observation function where $O(\bar{o}' | s', \bar{a})$ specifies the probability of attaining joint observation \bar{o}' when joint action \bar{a} results in state s' , and
- $R : S \times \bar{A} \rightarrow \mathbb{R}$ is a reward function such that $R(s, \bar{a})$ is the immediate reward for performing joint action \bar{a} in state s .

Agents in the *decentralized partially observable Markov decision process* (Dec-POMDP) cannot explicitly share information. Each agent therefore selects actions in accordance with a local policy π_i informed by action observation history h_i , where $\bar{h} \in (\prod_{i \in I} A_i O_i)^*$. We are primarily concerned with deterministic policies, from which agent actions can be inferred using only observation histories \bar{o}_h . The collection of individual policies is the decentralized policy

set $\bar{\pi} : \langle \pi_1, \pi_2, \dots, \pi_n \rangle$. Assume an infinite horizon $h = \infty$ and time discount rate γ . The objective is to find a policy set $\bar{\pi}$ maximizing expected reward over states and observation histories:

$$V^{\bar{\pi}}(s, \bar{o}_h) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(s^t, \bar{\pi}(\bar{o}_h^t)) \mid b^0, \bar{\pi} \right]$$

$$\bar{\pi}^* \leftarrow \operatorname{argmax}_{\bar{a} \in \bar{A}} \left(R(s, \bar{a}) + \gamma \sum_{s' \in S} \sum_{\bar{o}' \in \bar{O}} \Pr(s', \bar{o}' \mid s, \bar{a}) V^*(s', \bar{o}'_h) \right).$$

3.2 MPOMDPs

The *multiagent partially observable Markov decision process* (MPOMDP) is a stochastic, centralized multiagent model for sequential decision-making under partial observability also characterized by tuple $\langle I, S, \bar{A}, \bar{O}, T, O, R \rangle$. MPOMDP agents also lack a complete picture of the underlying system state, but can share individual actions a_i taken and observations received o_i . This permits a *sufficient statistic* in the form of a probability distribution over states called a *belief* $b \in \Delta^n$, where n is the number of states and Δ^n is the $n - 1$ simplex. We denote the multiagent belief using b_α . The belief over successor states s' is updated using the set of agent histories \bar{h} of actions taken and observation received:

$$b'_\alpha(s') = \eta O(\bar{o}' \mid s', \bar{a}) \sum_{s \in S} T(s' \mid s, \bar{a}) b_\alpha(s)$$

where η is a normalizing constant equal to $\Pr(\bar{o}' \mid b_\alpha, \bar{a})^{-1}$, and $\Pr(\bar{o}' \mid b_\alpha, \bar{a}) = \sum_{s' \in S} O(\bar{o}' \mid s', \bar{a}) \sum_{s \in S} T(s' \mid s, \bar{a}) b_\alpha(s)$. Belief $b'_{\alpha \bar{o} \bar{a}}$ is then generated by applying the belief update equation to all $s' \in S$. An initial belief b^0 is assumed to avoid considering the Bellman equation for infinite starting beliefs. From the agent perspective, reward must be calculated as a function of belief, such that $R(b_\alpha, \bar{a}) = \sum_{s \in S} R(s, \bar{a}) b_\alpha(s)$. In an MPOMDP, we seek a single joint policy $\pi_\alpha : \Delta^n \rightarrow \bar{A}$ that maps beliefs to actions. The expected value of b_α over an infinite horizon may be written as:

$$V^{\pi_\alpha}(b_\alpha) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(b_\alpha^t, \pi_\alpha(b_\alpha^t)) \mid b_\alpha^0 = b_\alpha, \pi_\alpha \right].$$

We seek a π_α^* that maximizes expected reward over time, determined by:

$$\pi_\alpha^* \leftarrow \operatorname{argmax}_{\bar{a} \in \bar{A}} \left(R(b_\alpha, \bar{a}) + \gamma \sum_{\bar{o}' \in \bar{O}} \Pr(\bar{o}' \mid b_\alpha, \bar{a}) V^*(b'_{\alpha \bar{o} \bar{a}}) \right).$$

3.3 Semi-Markov Processes

Definition 1. The *semi-Markov property for control* admits a distribution over time for state transitions aligned with agent actions.

Definition 2. *Sojourn control time* τ is the assignment of general continuous random variable \mathcal{T} .

$$Q(\mathcal{T} \leq \tau, s' \mid s, a) \quad (1)$$

The *semi-Markov decision process* (SMDP) [24] is a stochastic single agent model for sequential decision-making with sojourn system

control. The SMDP is characterized by state transition distribution Q , which is the probability that the next state transition occurs at or before τ and results in successor state s' . It is mathematically convenient to define Q as the product of $F(\tau \mid s, a)T(s' \mid s, a, \tau)$, or equivalently $F(\tau \mid s', a, s)T(s' \mid s, a)$, where F is the cumulative distribution function of τ . The distribution of τ is conditioned on the current state s and action taken a and is therefore Markov. The “semi” in semi-Markov reflects the arbitrary probability distribution followed by model transitions. When $\tau = 0$, a new action is taken to coincide with a *decision epoch*. The system natural process time is denoted as $\eta \in \mathbb{R}^+$, such that each decision epoch $t \in \mathbb{N}$ with corresponding sojourn time $\tau^t \in \mathbb{R}^+$ and state $s^t \in S$ occurs at η^t , where η^t is decision epoch start time within the natural process.

4 SEMI-DECENTRALIZATION

Definition 3. The *semi-Markov property for communication, or semi-decentralization*, admits a distribution over time for what information agents can store in memory.

Definition 4. *Sojourn communication time* τ is general continuous random variable representing the time for an agent to return to an information-sharing state. We here overload τ for sojourn communication time, distinct from sojourn control time.

$$Q(\bar{\mathcal{T}} \leq \bar{\tau}', s' \mid \bar{\tau}, s, \bar{a}, \bar{a}') \quad (2)$$

As with SMDPs, we can define Q as the product of $F(\bar{\tau}' \mid s', \bar{a}', \bar{\tau})$ and $T(s' \mid s, \bar{a}, \bar{\tau})$, where $\bar{\tau}'$ may be conditioned on the subsequent joint action set \bar{a}' . SMDPs have one agent with an implicit conditioned $\bar{\tau} = 0$. However, SDec-POMDPs have multiple agents with varied $\bar{\tau}$. Thus it is more general with $\bar{\tau}'$ conditioned on $\bar{\tau}$. Semi-decentralized models assume an initial $\bar{\tau}^0$, which can be interpreted as the communicating state of each agent when $\eta = 0$. When $\tau = 0$, information sharing can occur coinciding with a *communication epoch*. We assume noise-free instantaneous broadcast communication resulting in a single communicating agent set as in a blackboard [7, 8]. Semi-decentralization may however incorporate multiple distinct communicating sets. Whereas as semi-Markov control systems toggle model transition dynamics using τ , semi-decentralized systems toggle updating histories using $\bar{\tau}$.

5 THE SDEC-POMDP

The *semi-decentralized partially observable Markov decision process* (SDec-POMDP) is a stochastic, semi-decentralized multiagent model for sequential decision-making under partial observability characterized by tuple $\langle I, S, \bar{A}, \bar{O}, F, T, O, R \rangle$.

Model The **SDec-POMDP model** introduces selector functions f , g , and h to propagate agent memories, actions, and observations, respectively, to M_i and M_c through selector sub-nodes conditioned on $\bar{\tau}$. Figure 2 depicts an SDec-POMDP dynamic decision network where selector sub-nodes are contained within $Z_i^{\text{sel}} = \langle M_{ci}^{\text{sel}}, \bar{A}_i^{\text{sel}}, \bar{O}_i^{\text{sel}} \rangle$ and $Z_c^{\text{sel}} = \langle \bar{M}^{\text{sel}}, \bar{A}^{\text{sel}}, \bar{O}^{\text{sel}} \rangle$. The selector infrastructure defines information-sharing configurations and enables the model to simultaneously maintain sets of centralized and decentralized agents. The set compositions change with changes to $\bar{\tau}$ following each state transition. Define $M_c = (\prod_{i \in I} (\{\emptyset\} \cup A_i O_i))^*$, and $M_i = (A_i O_i)^* \times M_c$. The selector functions follow:

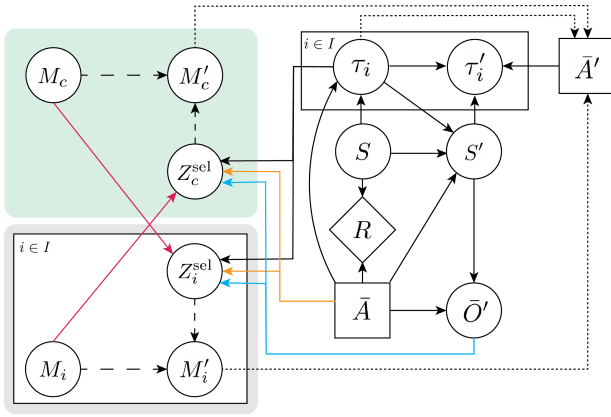


Figure 2: The SDec-POMDP dynamic decision network, with the policy infrastructure on the left and model on the right. The green backdrop contains the blackboard with memory M_c generated from the histories of communicating agents. The gray backdrop with plate notation includes the individual agent memories M_i . Z selector nodes are selectively toggled by $\bar{\tau}$ to facilitate memory propagation η , represented by dashed lines. Policy ψ edges are represented by dotted lines. The SDec-POMDP framework can be modified to capture the structural characteristics of different problem domains.

$$\forall i, f(m_{ci}^{sel}) = \begin{cases} m_c & \text{if } \tau_i = 0 \\ \emptyset & \text{if } \tau_i > 0 \end{cases}$$

$$f(m_c^{sel}) = \begin{cases} \forall i, m_i & \text{if } \tau_i = 0 \\ \emptyset & \text{if } \tau_i > 0 \end{cases}$$

$$\forall i, g(a_i^{sel}) = \begin{cases} \{a_j \in \bar{a}, \forall j \mid \tau_j = 0\} & \text{if } \tau_i = 0 \\ \{a_i\} & \text{if } \tau_i > 0 \end{cases}$$

$$g(a_c^{sel}) = \{a_i \in \bar{a}, \forall i \mid \tau_i = 0\}$$

$$\forall i, h(o_i^{sel}) = \begin{cases} \{o_j \in \bar{o}, \forall j \mid \tau_j = 0\} & \text{if } \tau_i = 0 \\ \{o_i\} & \text{if } \tau_i > 0 \end{cases}$$

$$h(o_c^{sel}) = \{o_i \in \bar{o}, \forall i \mid \tau_i = 0\}.$$

Policy For generality, assume an **SDec-POMDP policy** with deterministically stored $a_i o_i, \forall i$ and deterministically selected \bar{a} , parameterized by both $\pi_i : M_i \rightarrow A$ and $\pi_c : M_c \rightarrow \bar{A}$. Memory propagation η and policy ψ probability functions follow:

$$\psi(\bar{a}) = \begin{cases} 1, \forall i \begin{cases} a_i = (\pi_c(m_c))_i & \text{if } \tau_i = 0 \\ a_i = \pi_i(m_i) & \text{if } \tau_i > 0 \end{cases} \\ 0, \text{ otherwise} \end{cases}$$

$$\eta_c(m'_c) = \begin{cases} 1, \text{ if } m'_c = m_c \bar{m}^{sel} \bar{a}^{sel} \bar{o}^{sel} \\ 0, \text{ otherwise} \end{cases}$$

$$\forall i, \eta(m'_i) = \begin{cases} 1, \text{ if } m'_i = m_i m_{ci}^{sel} a_i^{sel} o_i^{sel} \\ 0, \text{ otherwise} \end{cases}.$$

Objective Function The **SDec-POMDP objective** is to identify the combination of memory propagation and policy functions that will maximize expected reward. Consider the infinite horizon case:

$$J(\psi, \eta_c, \bar{\eta}) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(s^t, \bar{a}^t) \mid b^0 \right]$$

$$J^* = \operatorname{argmax}_{\psi, \eta_c, \bar{\eta}} J(\psi, \eta_c, \bar{\eta}).$$

Rewrite in terms of policies $\bar{\pi}$ and π_c and blackboard and agent memories m_c and \bar{m} :

$$\begin{aligned} V^{\pi_c, \bar{\pi}}(m_c, \bar{m}) &= \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(s^t, \bar{a}^t) \mid b^0, m_c, \bar{m}, \pi_c, \bar{\pi} \right] \\ &= \sum_{s \in S} V^{\pi_c, \bar{\pi}}(m_c, \bar{m}, s) b^0(s) \end{aligned}$$

$$\begin{aligned} V^{\pi_c, \bar{\pi}}(m_c, \bar{m}, s) &= \sum_{\bar{a} \in \bar{A}} \psi(\bar{a} \mid m_c, \bar{m}) \left[R(s, \bar{a}) + \gamma \sum_{\bar{\tau} \in \bar{T}} F(d\bar{\tau} \mid s, \bar{a}) \right. \\ &\quad \sum_{s' \in S} T(s' \mid s, \bar{a}) \sum_{\bar{o}' \in \bar{O}} O(\bar{o}' \mid s', \bar{a}) \\ &\quad \sum_{m'_c \in M_c} \eta_c(m'_c \mid m_c, f(\bar{m}), g(\bar{a}), h(\bar{o}')) \\ &\quad \left. \sum_{\bar{m}' \in \bar{M}} \prod_{i \in I} \eta(m'_i \mid m_i, f(m_c), g(\bar{a}), h(\bar{o}')) V^{\pi_c, \bar{\pi}}(m'_c, \bar{m}', s') \right]. \end{aligned}$$

6 THEORETICAL ANALYSIS

Definition 5. Models X_θ and X_ϕ are equivalent if they reduce to one another via mapping function f , such that $X_\theta \leq X_\phi$ and $X_\phi \leq X_\theta$.

Definition 6. Model-policy structures XY_θ and XY_ϕ are equivalent if they reduce to one another via mapping function g , such that $XY_\theta \leq XY_\phi$ and $XY_\phi \leq XY_\theta$.

Definition 7. Model-policy-objective structures Z_θ and Z_ϕ are equivalent if $\forall \bar{h} V_{XY_\theta}(\bar{h}, b^0) = V_{XY_\phi}(\bar{h}, b^0)$.

6.1 MPOMDP

LEMMA 1. *SDec-POMDP and MPOMDP models are equivalent.*

PROOF. Demonstrate that 1. $X_{\text{MPOMDP}} \leq X_{\text{SDec-POMDP}}$ and 2. $X_{\text{SDec-POMDP}} \leq X_{\text{MPOMDP}}$

1. $X_{\text{MPOMDP}} \leq X_{\text{SDec-POMDP}}$

Let $I' = I$, $S' = S$, $\bar{A}' = \bar{A}$, $R' = R$, and $\bar{O}' = \bar{O}$, where prime notation indicates the SDec-POMDP for purposes of relating models. The state transition and observation functions are defined to reproduce the MPOMDP dynamics independently of $\bar{\tau}$, such that $T'(s' \mid s, \bar{a}, \bar{\tau}) = T(s' \mid s, \bar{a})$ and $O'(\bar{o}' \mid s', \bar{a}, \bar{\tau}) = O(\bar{o}' \mid s', \bar{a})$. Assume a deterministic communication sojourn time function F , where τ' for each agent is fixed to one, resulting in complete centralization at subsequent decision epochs; $\forall i, F'(\tau'_i < 1 \mid s, a_i, \tau_i) = 0$ and $F'(\tau'_i = 1 \mid s, a_i, \tau_i) = 1$. Selector functions therefore return joint actions and observations at each time-step; $\forall i, g'(a_i^{sel}) = g(a_c^{sel}) = \bar{a}$ and $\forall i, h'(o_i^{sel}) = h(o_c^{sel}) = \bar{o}$.

2. $X_{\text{SDec-POMDP}} \leq X_{\text{MPOMDP}}$

Let $R' = R$ and $\bar{A}' = \bar{A}$. The agent set is extended to include an additional blackboard agent with an independent memory, so that $I' = I \cup I_c$. We assume without loss of generality that $|I_c| = 1$, and that the communication sojourn time for this agent satisfies $\tau_i = 0 \forall i \in I_c$. The state space is augmented to include $\bar{\tau}$, such that $S' = S \times (\mathbb{R}^+)^n$. Similarly, the joint observation space is expanded to include the sequence of all action-observation pairs,

$$\bar{O}' = \bar{O} \times \prod_{i \in I} (A_i O_i)^*.$$

The transition and observation functions adopt the factored state space: $T'(\langle s', \bar{\tau}' \rangle | \langle s, \bar{\tau} \rangle, \bar{a}, \bar{a}') = F(\bar{\tau}' | s', \bar{a}', \bar{\tau})T(s' | s, \bar{a}, \bar{\tau})$ and $O'(\bar{o}' | \langle s', \bar{\tau}' \rangle, \langle s, \bar{\tau} \rangle, \bar{a}) = O(\bar{o}' | s', s, \bar{a}, \bar{\tau}, \bar{\tau})$ at its most general. \square

LEMMA 2. *SDec-POMDP and MPOMDP model-policy structures are equivalent.*

PROOF. Demonstrate that 1. $XY_{\text{MPOMDP}} \leq XY_{\text{SDec-POMDP}}$ and 2. $XY_{\text{SDec-POMDP}} \leq XY_{\text{MPOMDP}}$

1. $XY_{\text{MPOMDP}} \leq XY_{\text{SDec-POMDP}}$

Let each memory selector function return the joint memory, such that $\forall i, f'(m_{ci}^{\text{sel}}) = f'(m_c^{\text{sel}}) = \bar{m}$. The update rule for the joint memory is deterministic and concatenates the prior shared memory with the complete set of agent actions and observations from the current time-step, as represented by:

$$\eta'(\bar{m}') = \eta'_c(m'_c) = \begin{cases} 1, & \text{if } \bar{m}' = \bar{m} \bar{a}^{\text{sel}} \bar{o}^{\text{sel}} \\ 0, & \text{otherwise.} \end{cases}$$

Agent actions are selected using a policy over the joint memory:

$$\psi'(\bar{a}) = \begin{cases} 1, & \bar{a} = \pi(\bar{m}) \\ 0, & \text{otherwise.} \end{cases}$$

2. $XY_{\text{SDec-POMDP}} \leq XY_{\text{MPOMDP}}$

Construct policy $\pi^M(m'_i) = \langle a_1, \dots, a_n \rangle$, simulating where the action for agent i is determined to be:

$$a_i = \begin{cases} (\pi_c(\bar{m}_c))_i & \text{if } \tau_i = 0 \\ \pi_i(m_i) & \text{if } \tau_i > 0 \end{cases}$$

as if each agent draws from the blackboard's policy when its communication sojourn time is zero and otherwise following a local policy. The memory update rule is defined to extend the current memory with the observed joint outcome:

$$\eta'(\bar{m}') = \begin{cases} 1, & \text{if } \bar{m}' = \bar{m} \bar{o} \\ 0, & \text{otherwise.} \end{cases}$$

Joint action selection is consistent with the constructed joint policy:

$$\psi'(\bar{a}) = \begin{cases} 1, & \forall i, a_i = \pi^M(m'_i) \\ 0, & \text{otherwise.} \end{cases}$$

\square

LEMMA 3. *SDec-POMDP and MPOMDP model-policy-objective structures are equivalent.*

PROOF. Show $\forall \bar{h} V_{XY_{\text{MPOMDP}}}(\bar{h}, b_0) = V_{XY_{\text{SDec-POMDP}}}(\bar{h}, b_0)$

Show that the semi-decentralized value function reduces to the standard value function under the original joint policy. Begin with,

$$V^{\pi_c, \bar{\pi}}(m_c, \bar{m}, s) = \underbrace{\sum_{\bar{a} \in \bar{A}} \psi(\bar{a} | m_c, \bar{m})}_{1, m_c = \bar{h}} \left[R(s, \bar{a}) + \gamma \underbrace{\sum_{\bar{\tau} \in \bar{T}} F(d\bar{\tau} | s, \bar{a})}_{1, \text{all RV } \perp \text{ of } d\bar{\tau}} \right. \\ \left. \sum_{s' \in S} T(s' | s, \bar{a}) \sum_{\bar{o}' \in \bar{O}} O(\bar{o}' | s', \bar{a}) \right. \\ \left. \underbrace{\sum_{m'_c \in M_c} \eta_c(m'_c | m_c, f(\bar{m}), g(\bar{a}), h(\bar{o}'))}_{1, m'_c = m_c \bar{a}^{\text{sel}} \bar{o}^{\text{sel}}} \right. \\ \left. \underbrace{\sum_{\bar{m}' \in \bar{M}} \prod_{i \in I} \eta(m'_i | m_i, f(m_c), g(\bar{a}), h(\bar{o}'))}_{1, \forall i, m'_i = m_i a_i^{\text{sel}} o_i^{\text{sel}}} V^{\pi_c, \bar{\pi}}(m'_c, \bar{m}', s') \right].$$

By construction, each under-braced term evaluates deterministically: the blackboard's memory update enforces $m'_c = m_c \bar{a}^{\text{sel}} \bar{o}^{\text{sel}}$, each agent's local memory update yields $m'_i = m_i a_i^{\text{sel}} o_i^{\text{sel}}$, and the distribution over sojourn times collapse to one. Moreover, since $\psi(\bar{a} | m_c, \bar{m}, s) = 1$ whenever $\bar{a} = \pi(\bar{h})$, we obtain $V^{\pi_c, \bar{\pi}}(m_c, \bar{m}, s) = V^\pi(s, \bar{h})$, which results in standard value recursion:

$$R(s, \pi(\bar{h})) + \gamma \sum_{s' \in S} \sum_{\bar{o}' \in \bar{O}} \Pr(s', \bar{o}' | s, \pi(\bar{h})) V^\pi(s', \bar{h}').$$

Further observe that $\bar{m}_c = \bar{m}$ and

$$\sum_{s' \in S} T(s' | s, \bar{a}) \sum_{\bar{o}' \in \bar{O}} O(\bar{o}' | s', \bar{a}) = \sum_{s' \in S} \sum_{\bar{o}' \in \bar{O}} \Pr(s', \bar{o}' | s, \pi(\bar{h})).$$

\square

PROPOSITION 1. *The SDec-POMDP and MPOMDP are equivalent.*

6.2 Dec-POMDP

LEMMA 4. *SDec-POMDP and Dec-POMDP models are equivalent.*

PROOF. 1. Demonstrate that $X_{\text{Dec-POMDP}} \leq X_{\text{SDec-POMDP}}$ and 2. $X_{\text{SDec-POMDP}} \leq X_{\text{Dec-POMDP}}$

1. $X_{\text{Dec-POMDP}} \leq X_{\text{SDec-POMDP}}$

Again set $I' = I$, $S' = S$, $\bar{A}' = \bar{A}$, $R' = R$, and $\bar{O}' = \bar{O}$. Let $T'(s' | s, \bar{a}, \bar{\tau}) = T(s' | s, \bar{a})$ and $O'(\bar{o}' | s', \bar{a}, \bar{\tau}') = O(\bar{o}' | s', \bar{a})$. Action and observation selection are specified so that for every agent i , $g'(a_i^{\text{sel}}) = a_i$ and $h'(o_i^{\text{sel}}) = o_i$. Assign to the blackboard memory the null set, such that $g'(a_c^{\text{sel}}) = h'(o_c^{\text{sel}}) = \emptyset$. By the construction of g' and h' , τ has no impact and F can be disregarded.

2. $X_{\text{SDec-POMDP}} \leq X_{\text{Dec-POMDP}}$

Reference Lemma 1 proof 2, as $X_{\text{Dec-POMDP}} = X_{\text{MPOMDP}}$. \square

LEMMA 5. *SDec-POMDP and Dec-POMDP model-policy structures are equivalent.*

PROOF. Demonstrate that 1. $XY_{\text{Dec-POMDP}} \leq XY_{\text{SDec-POMDP}}$ and 2. $XY_{\text{SDec-POMDP}} \leq XY_{\text{Dec-POMDP}}$

1. $XY_{\text{Dec-POMDP}} \leq XY_{\text{SDec-POMDP}}$

Let each memory selector function return the null set, such that $\forall i$, $f'(m_c^{\text{sel}}) = f'(m_c^{\text{sel}}) = \emptyset$. The update rule for each agent's memory is deterministic and concatenates the prior memory with the set of individual agent actions and observations from the current time-step, as represented by:

$$\eta'(m'_i) = \begin{cases} 1, & \text{if } m'_i = m_i a_i^{\text{sel}} o_i^{\text{sel}} \\ 0, & \text{otherwise} \end{cases}$$

Agent actions are selected using a policy over the agent's memory:

$$\psi'(\bar{a}) = \begin{cases} 1, & \forall i, a_i = \pi_i(m_i) \\ 0, & \text{otherwise} \end{cases}$$

By the construction of ψ' , η'_c can be disregarded.

2. $XY_{\text{SDec-POMDP}} \leq XY_{\text{Dec-POMDP}}$

Construct simulated policy $(\pi^D)_i(m'_i) = a_i$, where:

$$a_i = \begin{cases} (\pi_c(m_c))_i & \text{if } \tau_i = 0 \\ \pi_i(m_i) & \text{if } \tau_i > 0 \end{cases}$$

as if each agent draws from the blackboard's policy when its communication sojourn time is zero and otherwise following a local policy. Each agent's memory update rule appends the taken action and observation to the current memory:

$$\eta'(m'_i) = \begin{cases} 1, & \text{if } m'_i = m_i a_i o_i \\ 0, & \text{otherwise} \end{cases}$$

Agent action selection is consistent with the constructed policy:

$$\psi'(\bar{a}) = \begin{cases} 1, & \forall i, a_i = (\pi^D)_i(m'_i) \\ 0, & \text{otherwise} \end{cases}$$

□

LEMMA 6. *SDec-POMDP and Dec-POMDP model-policy-objective structures are equivalent.*

PROOF. Show $\forall \bar{h} V_{XY_{\text{Dec-POMDP}}}(\bar{h}, b_0) = V_{XY_{\text{SDec-POMDP}}}(\bar{h}, b_0)$

The semi-decentralized value function reduces to the decentralized value function under the original policy set, beginning with:

$$\begin{aligned} V^{\pi_c, \bar{\pi}}(\underbrace{m'_c}_{M_c=\emptyset}, \bar{m}, s) &= \sum_{\bar{a} \in \bar{A}} \underbrace{\psi(\bar{a} \mid \underbrace{m'_c}_{M_c=\emptyset}, \bar{m})}_1 \left[R(s, \bar{a}) + \right. \\ &\underbrace{\gamma \sum_{\bar{\tau} \in \bar{T}} F(d\bar{\tau} \mid s, \bar{a}) \sum_{s' \in S} T(s' \mid s, \bar{a}) \sum_{\bar{o}' \in \bar{O}} O(\bar{o}' \mid s', \bar{a})}_{1, \text{ all RV } \perp \text{ of } d\bar{\tau}} \\ &\underbrace{\sum_{\bar{m}' \in \bar{M}} \prod_{i \in I} \eta(m'_i \mid m_i, \underbrace{f(m'_c), g(\bar{a}), h(\bar{o}')}_{M_c=\emptyset})}_{1, \forall i, m'_i = m_i a_i^{\text{sel}} o_i^{\text{sel}}} \\ &\left. \underbrace{\sum_{m'_c \in M_c} \eta_c(m'_c \mid m_c, \underbrace{f(\bar{m}), g(\bar{a}), h(\bar{o}')}_{M_c=\emptyset}) V^{\pi_c, \bar{\pi}}(\underbrace{m'_c}_{M_c=\emptyset}, \bar{m}', s')}_{M_c=\emptyset} \right]. \end{aligned}$$

By construction, each under-braced term evaluates deterministically: the blackboard's memory remains $m'_c = \emptyset$, each agent's local

memory update yields $m'_i = m_i a_i^{\text{sel}} o_i^{\text{sel}}$, and the distribution over sojourn times collapses to one. As before, since $\psi(\bar{a} \mid m_c, \bar{m}, s) = 1$ whenever $\bar{a} = \bar{\pi}(\bar{h})$, we obtain $V^{\pi_c, \bar{\pi}}(m_c, \bar{m}, s) = V^{\bar{\pi}}(s, \bar{h})$, which results in standard value recursion:

$$R(s, \bar{\pi}(\bar{h})) + \gamma \sum_{s' \in S} \sum_{\bar{o}' \in \bar{O}} \Pr(s', \bar{o}' \mid s, \bar{\pi}(\bar{h})) V^{\bar{\pi}}(s', \bar{h}')$$

Again observe that,

$$\sum_{s' \in S} T(s' \mid s, \bar{a}) \sum_{\bar{o}' \in \bar{O}} O(\bar{o}' \mid s', \bar{a}) = \sum_{s' \in S} \sum_{\bar{o}' \in \bar{O}} \Pr(s', \bar{o}' \mid s, \bar{\pi}(\bar{h})).$$

□

PROPOSITION 2. *The SDec-POMDP and Dec-POMDP are equivalent*

Corollary 1. The SDec-POMDP is the same complexity class as a Dec-POMDP (NEXP-complete)

PROPOSITION 3. *The SDec-POMDP and k-steps-delayed communication are equivalent*

PROPOSITION 4. *The SDec-POMDP and Dec-POMDP-Com are equivalent*

We provide proofs for Propositions 3 and 4 in the technical appendix.

7 RECURSIVE SMALL-STEP SEMI-DECENTRALIZED A*

Recursive small-step semi-decentralized A* (RS-SDA*) is an exact planning algorithm for optimizing expected reward in SDec-POMDP problems. RS-SDA* modifies RS-MAA* [16] by maintaining a stage-specific partition of decentralized and centralized joint observation histories per probabilistic communication dynamics. Like RS-MAA*, RS-SDA* relies on incremental expansion of a small-step search tree, clustering, recursive heuristics, memoization, and last stage modifications. RS-SDA* generates a fully-specified policy set $\bar{\pi} \in \Pi$ using offline planning. A fully-specified policy contains both fully-specified local policies and, if appropriate for the problem, a blackboard policy $\bar{\pi} = \langle \pi_1, \dots, \pi_n, \pi_c \rangle$ where $\pi_i : \mathcal{O}^h \rightarrow A_i$ and $\pi_c : \bar{\mathcal{O}}^h \rightarrow \bar{A}$. Similarly, $\bar{\varphi} = \langle \varphi_1, \dots, \varphi_n, \varphi_c \rangle$, $\varphi_i : \mathcal{O}^{\leq h-1} \rightarrow A_i$, and $\bar{\varphi} : \bar{\mathcal{O}}^{\leq h-1} \rightarrow \bar{A}$. RS-SDA* is outlined in Algorithm 1.

Small-Step Search Tree We adopt the small-step approach first introduced by Cazenave for A* [6] and used by Kooops et al. in RS-MAA* [16] to limit search tree outdegree. Small-step search can be used with centralized and decentralized components of each policy node, depicted in Figure 3. As shown in Table 1, small-step search provides RS-SDA* with mixed component policies, pre-clustering, a lower bound (complete decentralization) and an upper bound (complete centralization) on the number of considered nodes per stage t . When $F(\tau \mid s, \bar{a}) = F(\tau \mid \bar{a})$ and the policy is deterministic, we can consolidate the centralized and decentralized components of policies and significantly reduce the search tree size. This results in an RS-SDA* lower bound that is equivalent to RS-MAA* and an RS-SDA* upper bound that has $|O_\star|^{t^n}$ levels per stage and considers $|A_\star|^n$ joint actions per level, where $|O_\star|$ is the size of the largest observation set and $|A_\star|$ is the size of the largest action set.

t	RS-MAA* (RS-SDA* lower bound)		RS-SDA* upper bound		Classical MAA*	
	levels	nodes/stage	levels	nodes/stage	levels	nodes/stage
0	$n O_* ^t$	$n O_* ^t A_* $	$ O_* ^t$	$ O_* ^t A_* ^n$	1	$ A_* ^n O_* ^t$
1	2	6	1	9	1	9
2	4	12	4	36	1	81
3	8	24	16	144	1	6,561
4	16	48	64	576	1	>1E6
5	32	96	256	2,304	1	>1E6
6	64	192	1,024	9,216	1	>1E6
7	128	384	4,096	36,864	1	>1E6
8	256	768	16,384	147,456	1	>1E6
8	512	1536	65,536	589,824	1	>1E6

Table 1: SDEC-TIGER levels and nodes by stage for RS-MAA*, RS-SDA*, and MAA*. $|A_*| = 3$, $|O_*| = 2$, and $n = 2$.

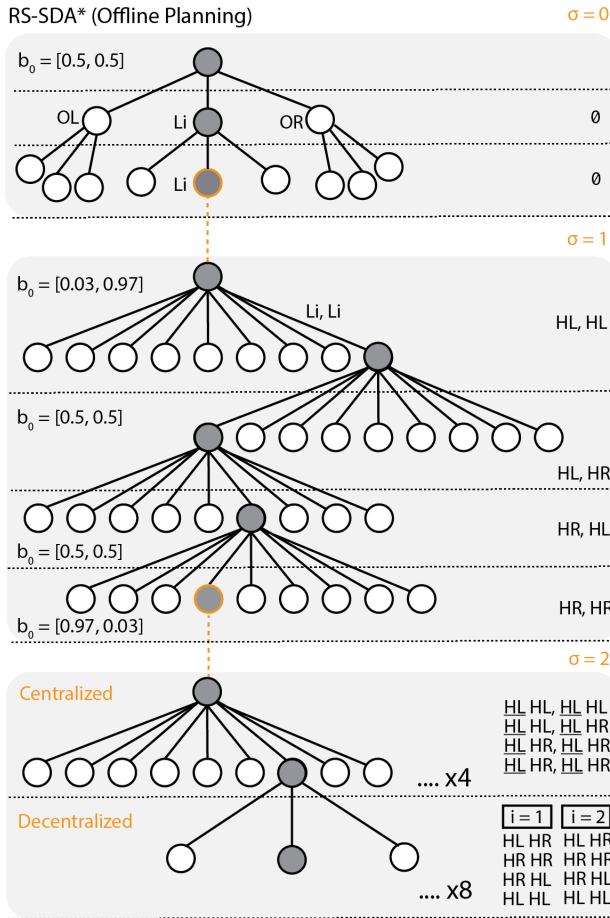


Figure 3: Illustrating RS-SDA* applied to SDEC-TIGER using mixed component policies through stage $\sigma = 2$.

Dynamic Programming We apply backward induction over beliefs to rapidly determine the value of centralized policy components. This bypasses expensive recursive heuristic calculations for large portions of the search. For each remaining horizon r and belief b , we compute $V_r(b)$ and $Q_r(b, \bar{a})$ and memoize both V and Q under keys (r, b) and (r, b, \bar{a}) to enable extensive reuse during A^* expansion. Similarly, observation likelihoods $P(\bar{o} | b, \bar{a})$ and posteriors

$(b'_{\bar{o}, \bar{a}})$ are obtained via a belief update over the model’s transition and observation tensors, then cached for subsequent calls.

Ordering Observation Histories Any $\bar{\varphi}$ may contain both decentralized and centralized mappings conditioned on the underlying state and actions taken. We therefore explore joint observation histories (JOH) and local observation histories (LOH) in a predetermined sequence: stage, JOH then LOH, by agent (for LOHs), then lexicographically. Observe that, for length t , all $\bar{o}^{0:t} \preceq o^{0:t}$ and all $o^{0:t} \preceq \bar{o}^{0:t}$. Additionally, $\bar{o}^{0:t} \preceq_{\text{lex}} \bar{o}^{0:t}$, and $o_i^{0:t} \preceq o_j^{0:t}$ if $(i < j)$ or $(i = j \wedge o_i^{0:t} \preceq_{\text{lex}} o_j^{0:t})$.

Clustering We implement lossless incremental clustering in decentralized policy components based on a probabilistic equivalence criterion [23]. We similarly cluster centralized policy components based on the resulting joint belief, or $\forall s, P(s | \bar{o}_1^{0:t}) = P(s | \bar{o}_2^{0:t})$.

Admissible Heuristic As with multiagent A^* (MAA*) [30], an *admissible heuristic* Q guides a path through a search tree with partial policies as nodes $\bar{\varphi}$. A heuristic is admissible if it equals or over-approximates the true value of the policy node. An *open list* is maintained with nodes under consideration. The node in the open list with the highest heuristic value is expanded and replaced in the open list by its children. The tree search terminates once a fully-specified policy with the highest heuristic value is identified. For each parent node and candidate action, we split each posterior belief using $\mathcal{S}_{\text{sync}}$ and $\mathcal{A}_{\text{sync}}$ (more generally expressed using $F(s, a)$) and take a probability-weighted sum of the exact centralized value on the communication-dependent posterior and the exact decentralized value on its complement. Because every constituent (centralized, decentralized, and their mixture conditioned on communication) is an exact optimum of a relaxation of the remaining subproblem, the resulting heuristic **never underestimates** the achievable return and is therefore admissible.

8 EXPERIMENTS

Semi-Decentralized Benchmarks We evaluate RS-SDA*¹ in semi-decentralized versions of four common Dec-POMDP benchmarks: DEC-TIGER [18], FIREFIGHTING [22], BOXPUSHING [25], and MARS [2], and in a new MARITIMEMEDEVAC benchmark. Problem descriptions with illustrations are disclosed in the technical appendix. All experiments were conducted using an AMD Ryzen 9 9900X3D 12-Core Processor (4400 MHz), with timeout occurring at 20 minutes and memory out at 16 GB. We adopt hyper-parameters $M = 200$, $d = 3$, and $\alpha = 0.2$ for all experiments.




Figure 4: MARITIMEMEDEVAC exact optimal policy values across communication regimes for h one through eight.

¹Code repository and technical appendix: <https://github.com/mahdial-husseini/RSSDA>

Algorithm 1: Recursive Small-Step Semi-Decentralized A*

Input : $\Phi \triangleq (\mathcal{S}_{\text{sync}}, \mathcal{A}_{\text{sync}})$; h horizon; b initial belief; φ initial policy; d heuristic depth; M iterations; u upper bound

Output : optimal policy φ^* , optimal value v^*

function $\text{RS-SDA}^*(h, b, \varphi, d, M, u, \Phi)$

```

 $Q \leftarrow \text{PRIORITYQUEUE}(\uparrow)$ 
 $Q.\text{push}(\min(\varphi.\text{heuristics}), \varphi)$ 
 $i \leftarrow 0$ 
while  $Q$  is not empty do
   $(v, \varphi) \leftarrow Q.\text{pop}()$ 
  if  $v < \infty$  then  $i \leftarrow i + 1$ 
  if  $i \geq M$  or  $v \leq u$  then return  $(\min(v, u), \text{None})$ 
   $\sigma \leftarrow$  current stage of  $\varphi$ ;  $j \leftarrow$  current agent index of  $\varphi$ 
  // Phase 1: Transition to Subsequent Stage
  if stage  $\sigma$  is complete ( $j = |A|$ ) then
    if  $\sigma = h$  then return  $(v, \varphi)$ 
     $(D_{\text{dec}}, D_{\text{cen}}, P_{\text{dec}}, P_{\text{cen}}, \rho) \leftarrow \text{TERMINALPROBS}(\varphi, \Phi)$ 
    if  $\rho > 0$  then
       $\varphi \leftarrow \text{CLUSTERPOLICYDEC}(\varphi, D_{\text{dec}}, P_{\text{dec}})$ 
    if  $\rho < 1$  then
       $\varphi \leftarrow \text{CLUSTERPOLICYCEN}(\varphi, D_{\text{cen}}, P_{\text{cen}})$ 
       $w \leftarrow \text{EVALUATEPOLICY}(\varphi, \rho)$ 
       $\varphi.\text{heuristics}.\text{append}(w)$ 
       $\varphi.\text{depth} \leftarrow \min(\sigma, d)$ 
       $\varphi.\text{extend\_empty\_stage}(); \sigma \leftarrow \sigma + 1; j \leftarrow 0$ 
    end
  // Phase 2: Policy Expansion
  if  $\sigma = h$  (Final Stage Optimization) then
    if  $\rho < 1$  and centralized unfilled then
      determine  $\bar{a}^*$  for centralized component
       $\varphi' \leftarrow \varphi$  with  $\bar{a}^*$  set;
       $w \leftarrow \text{EVALUATEPOLICY}(\varphi', \rho)$ 
       $Q.\text{push}(w, \varphi')$ 
    else
      determine  $a_j^*$  for decentralized component
       $\varphi' \leftarrow \varphi$  with  $a_j^*$  set;
       $w \leftarrow \text{EVALUATEPOLICY}(\varphi', \rho)$ 
      if  $w = v$  then return  $(w, \varphi')$ 
       $Q.\text{push}(w, \varphi')$ 
    else
      if  $\rho < 1$  and centralized unfilled then
        foreach joint action  $\bar{a} \in \bar{A}$  do
           $\varphi' \leftarrow \varphi$  with  $\bar{a}$  set
           $w \leftarrow \text{EXACTCENTRALQ}(\varphi', \bar{a})$ 
           $Q.\text{push}(w, \varphi')$ 
        end
      else if agent  $j$  unfilled then
        foreach  $a_j \in A_j$  do
           $\varphi' \leftarrow \varphi$  with  $a_j$  set
           $w \leftarrow \text{EVALUATEPOLICY}(\varphi', \rho)$ 
           $Q.\text{push}(w, \varphi')$ 
        end
      end
    end
  end

```

Results As shown in Table 2 and Figure 4, RS-SDA* is competitive with the centralized upper bound across most semi-decentralized benchmarks and MARITIMEDEVAC. The modified benchmarks demonstrate how model dynamics influence the value of information in multi-agent systems. SDEC-FIREFIGHTING exemplifies problems where centralization benefits are negligible, and the optimal RS-SDA* solution equals the optimal RS-MAA* solution for all considered h . By contrast, SDEC-BOX exemplifies problems where partial centralization results in complete information sharing, and the optimal RS-SDA* solution equals the fully centralized optimum for all considered h . In MARITIMEDEVAC, the three regimes are nearly indistinguishable at moderate horizons ($H = 4, 5, 6$), but at

$H = 7$ centralized reaches 6.62 while semi-decentralized attains 6.36, clearly outperforming full decentralization (3.27). At $H = 7$, the semi-decentralized policy recovers about 96% of the centralized value. These results indicate that semi-decentralization and RS-SDA* preserve much of the benefit of centralized coordination while staying tractable, with occasional slowdowns or timeouts on problem instances where lossless clustering is largely ineffective.

	lower bound		our approach		upper bound	
	decentralized		semi-decentralized		centralized	
	RS-MAA*		RS-SDA*		RS-SDA*	
h	value	time	value	time	value	time
SDEC-TIGER						
8	12.21726	1	27.21518	<1	47.71696	<1
9	15.57244	6	30.90457	<1	53.47353	<1
10	15.18438	271	34.72370	<1	60.50988	<1
SDEC-FIREFIGHTING ($n_h = 3, n_f = 3$)						
3	-5.73697	<1	-5.73697	<1	-5.72285	<1
4	-6.57883	2	-6.57883	3	-6.51859	3
5	-7.06987	29	-7.06987	49	-6.98069	37
SDEC-BOXPUSHING						
3	66.08100	<1	66.81000	<1	66.81000	<1
4	98.59361	18	99.55630	<1	99.55630	<1
5	107.72985	MO	109.09251	1	109.09251	1
SDEC-MARS (Right Band Rendezvous)						
4	10.18080	1	10.18080	<1	10.87020	<1
5	13.26654	2	14.29038	<1	14.38556	1
6	18.62317	4	20.06430	2	20.06706	3
SDEC-MARS (Survey Site Beacons)						
4	10.18080	1	10.54620	<1	10.87020	<1
5	13.26654	2	13.26654	<1	14.38556	1
6	18.62317	4	18.62317	143	20.06706	3
SDEC-MARS (Drill Site Beacons)						
4	10.18080	1	10.87020	<1	10.87020	<1
5	13.26654	2	14.38556	<1	14.38556	1
6	18.62317	4	20.06168	2	20.06706	3
MARITIMEDEVAC						
6	3.18348	<1	3.19807	28	3.19945	<1
7	3.26710	2	6.36301	33	6.61819	1
8	8.03228	156	10.61275	MO	10.88244	1

Table 2: RS-SDA* offline planning performance on semi-decentralized benchmarks and MARITIMEDEVAC. TO denotes timeout (>1200s) and MO denotes memout (>16GB).

9 CONCLUSION

We present a foundational framework for multiagent decision making under probabilistic communication. We introduce the SDec-POMDP, which unifies the Dec-POMDP, MPOMDP, and several communication mechanisms with delay, loss, or cost. A secondary contribution is RS-SDA*, an admissible heuristic search algorithm for semi-decentralized systems with excellent performance, and semi-decentralized versions of four standard benchmarks and a new medical evacuation scenario. Taken together, SDec-POMDP and RS-SDA* provide a principled basis for studying and exploiting probabilistic communication in cooperative teams. Future work includes exploiting interleaving offline planning and online search to improve approximate RS-SDA* performance and investigating systems with non-stationary sojourn time distributions.

REFERENCES

- [1] Christopher Amato, Girish Chowdhary, Alborz Geramifard, N Kemal Üre, and Mykel J Kochenderfer. 2013. Decentralized control of partially observable Markov decision processes. In *52nd IEEE Conference on Decision and Control*. IEEE, IEEE, 2398–2405.
- [2] Christopher Amato and Shlomo Zilberstein. 2009. Achieving goals in decentralized POMDPs. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1 (Budapest, Hungary) (AAMAS '09)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 593–600.
- [3] Md Ali Azam, Hans D Mittelmann, and Shankarachary Ragi. 2021. UAV formation shape control via decentralized markov decision processes. *Algorithms* 14, 3 (2021), 91.
- [4] Raphen Becker, Shlomo Zilberstein, Victor Lesser, and Claudia V Goldman. 2004. Solving transition independent decentralized Markov decision processes. *Journal of Artificial Intelligence Research* 22 (2004), 423–455.
- [5] Daniel S Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research* 27, 4 (2002), 819–840.
- [6] Tristan Cazenave. 2010. Partial Move A*. In *2010 22nd IEEE International Conference on Tools with Artificial Intelligence*, Vol. 2. IEEE, 25–31.
- [7] Iain D Craig. 1988. Blackboard systems. *Artificial Intelligence Review* 2, 2 (1988), 103–118.
- [8] Lee D Erman, Frederick Hayes-Roth, Victor R Lesser, and D Raj Reddy. 1980. The Hearsay-II speech-understanding system: Integrating knowledge to resolve uncertainty. *ACM Computing Surveys (CSUR)* 12, 2 (1980), 213–253.
- [9] Claudia V Goldman and Shlomo Zilberstein. 2003. Optimizing information exchange in cooperative multi-agent systems. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*. 137–144.
- [10] Claudia V Goldman and Shlomo Zilberstein. 2008. Communication-based decomposition mechanisms for decentralized MDPs. *Journal of Artificial Intelligence Research* 32 (2008), 169–202.
- [11] Alan Grant, Paul Williams, Nick Ward, and Sally Basker. 2009. GPS jamming and the impact on maritime navigation. *The Journal of Navigation* 62, 2 (2009), 173–187.
- [12] Pulkit Grover and Anant Sahai. 2010. Implicit and explicit communication in decentralized control. In *2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 278–285.
- [13] Jiechuan Jiang and Zongqing Lu. 2018. Learning Attentional Communication for Multi-Agent Cooperation. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc.
- [14] Kyle D Julian and Mykel J Kochenderfer. 2019. Distributed wildfire surveillance with autonomous aircraft using deep reinforcement learning. *Journal of Guidance, Control, and Dynamics* 42, 8 (2019), 1768–1778.
- [15] Mykel J Kochenderfer. 2015. *Decision making under uncertainty: theory and application*. MIT press.
- [16] Wietze Koops, Nils Jansen, Sebastian Junges, and Thiago D. Simão. 2023. Recursive small-step multi-agent A* for dec-POMDPs. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (Macao, P.R.China) (IJCAI '23)*. Article 600, 9 pages.
- [17] Thomas M Moerland, Joost Broekens, Aske Plaat, and Catholijn M Jonker. 2022. A unifying framework for reinforcement learning and planning. *Frontiers in Artificial Intelligence* 5 (2022).
- [18] Ranjit Nair, Milind Tambe, Makoto Yokoo, David Pynadath, and Stacy Marsella. 2003. Taming decentralized POMDPs: Towards efficient policy computation for multiagent settings. In *IJCAI*, Vol. 3. 705–711.
- [19] Frans Oliehoek and Matthijs Spaan. 2012. Tree-based solution methods for multiagent POMDPs with delayed communication. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 26. 1415–1421.
- [20] Frans Adriaan Oliehoek. 2013. Sufficient Plan-Time Statistics for Decentralized POMDPs. In *IJCAI*. 302–308.
- [21] Frans A. Oliehoek and Christopher Amato. 2016. *A Concise Introduction to Decentralized POMDPs* (1st ed.). Springer Publishing Company, Incorporated.
- [22] Frans A Oliehoek, Matthijs TJ Spaan, and Nikos Vlassis. 2008. Optimal and approximate Q-value functions for decentralized POMDPs. *Journal of Artificial Intelligence Research* 32 (2008), 289–353.
- [23] Frans A. Oliehoek, Shimon Whiteson, and Matthijs T. J. Spaan. 2009. Lossless clustering of histories in decentralized POMDPs. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1 (Budapest, Hungary) (AAMAS '09)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 577–584.
- [24] Martin L Puterman. 1994. Markov decision processes: Discrete stochastic dynamic programming. John Wiley & Sons. Inc., New York, NY (1994).
- [25] Sven Seuken and Shlomo Zilberstein. 2007. Improved memory-bounded dynamic programming for decentralized POMDPs. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*. 344–351.
- [26] Junjie Sheng, Xiangfeng Wang, Bo Jin, Junchi Yan, Wenhao Li, Tsung-Hui Chang, Jun Wang, and Hongyuan Zha. 2022. Learning structured communication for multi-agent reinforcement learning. *Autonomous Agents and Multi-Agent Systems* 36, 2 (2022), 50.
- [27] Amanpreet Singh, Tushar Jain, and Sainbayar Sukhbaatar. 2018. Learning when to communicate at scale in multiagent cooperative and competitive tasks. *arXiv preprint arXiv:1812.09755* (2018).
- [28] Arambam James Singh, Duc Thien Nguyen, Akshat Kumar, and Hoong Chuin Lau. 2019. Multiagent decision making for maritime traffic management. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 6171–6178.
- [29] Sagar Sudhakara, Dhruva Kartik, Rahul Jain, and Ashutosh Nayyar. 2024. Optimal Communication and Control Strategies in a Cooperative Multiagent MDP Problem. *IEEE Trans. Automat. Control* 69, 10 (2024), 6959–6966.
- [30] Daniel Szer, François Charpillet, and Shlomo Zilberstein. 2005. MAA*: a heuristic search algorithm for solving decentralized POMDPs. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (Edinburgh, Scotland) (UAI'05)*. AUAI Press, Arlington, Virginia, USA, 576–583.
- [31] Tze-Yang Tung, Szymon Kobus, Joan Pujol Roig, and Deniz Gündüz. 2021. Effective communications: A joint learning and communication framework for multi-agent reinforcement learning over noisy channels. *IEEE Journal on Selected Areas in Communications* 39, 8 (2021), 2590–2603.
- [32] Ruixue Zhang, Jiao Wang, Jun Ge, and Qiyuan Huang. 2024. Multiagent cooperative search learning with intermittent communication. *IEEE Intelligent Systems* 39, 2 (2024), 11–20.
- [33] Changxi Zhu, Mehdi Dastani, and Shihan Wang. 2024. A survey of multi-agent deep reinforcement learning with communication. *Autonomous Agents and Multi-Agent Systems* 38, 1 (2024), 4.