Received: 22 July 2025 | Revised: 22 July 2025 | Accepted: 15 August 2025

DOI: 10.1002/aaai.70023

ARTICLE

Semi-Markovian planning to coordinate aerial and maritime medical evacuation platforms

Mahdi Al-Husseini[1,2] | **Kyle H. Wray**[3] | **Mykel J. Kochenderfer**[2]

[1]25th Infantry Division, United States Army, Schofield Barracks, USA

[2]Department of Aeronautics and Astronautics, Stanford University, Stanford, USA

[3]Khoury College of Computer Sciences, Northeastern University, Boston, USA

Correspondence
Mahdi Al-Husseini, United States Army.
Email: mah9@stanford.edu

Abstract

The transfer of patients between two aircraft using an underway watercraft increases medical evacuation reach and flexibility in maritime environments. The selection of any one of multiple underway watercraft for patient exchange is complicated by participating aircraft utilization histories and participating watercraft positions and velocities. The selection problem is modeled as a semi-Markov decision process with an action space, including both fixed land and moving watercraft exchange points. Monte Carlo tree search with root parallelization is used to select optimal exchange points and determine aircraft dispatch times. Model parameters are varied in simulation to identify representative scenarios where watercraft exchange points reduce incident response times. We find that an optimal policy with watercraft exchange points outperforms an optimal policy without watercraft exchange points and a greedy policy by 35% and 40%, respectively. In partnership with the United States Army, we deploy for the first time the watercraft exchange point by executing a mock patient transfer with a manikin between two HH-60M medical evacuation helicopters and an underway Army Logistic Support Vessel south of the Hawaiian island of Oahu. Both helicopters were dispatched in accordance with our optimized decision strategy.

INTRODUCTION

Medical planners coordinate multiple aerial and maritime evacuation platforms to facilitate the transfer of patients across expansive, non-contiguous maritime environments. To do so effectively involves understanding each platform's patient carrying capacity, position, forward speed, and other transportation and usage constraints, as well as the operational environment in time and space, for example, casualty estimate data. The trajectory of watercraft, often underway in support of non-medical mission requirements, further complicates the exchange point selection process. Due to the challenging nature of the multi-agent coordination problem, patients in maritime environments are currently transferred via evacuation aircraft or hospital ships, or occasionally, from evacuation aircraft to hospital ships. Transportation exclusively by evacuation aircraft is fast, but transport distance is limited by aircraft range. Transportation exclusively by hospital ship can support any transport distance, but is slow. This paper considers, for the first time, the use of underway watercraft as intermediary exchange points for two medical evacuation aircraft arriving from different islands (Figure 1). Doing so combines the strengths of both aerial and maritime

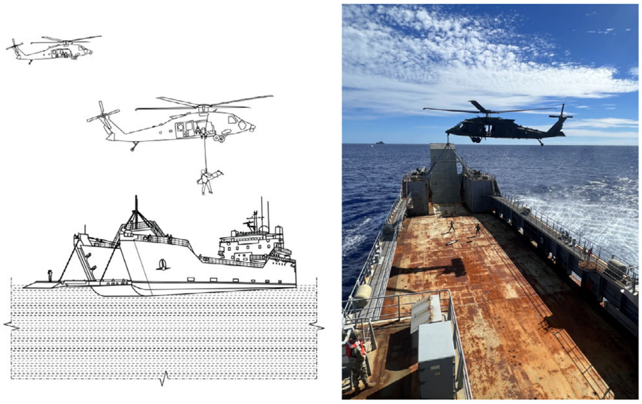

(A) (B) © Mahdi Al-Husseini 2023

FIGURE 1 An evacuation aircraft lowers a patient onto a watercraft while a second aircraft circles nearby. Using watercraft as exchange points between aircraft expedites patient movement and enables evacuation across vast distances.

evacuation platforms, but requires careful selection and dispatching to minimize transport delays. This novel capability is inspired by the ambulance exchange point, which uses fixed, land-based points for patient exchange between ambulances (United States Army 2019). Unlike traditional ambulance exchange points, however, watercraft exchange locations are neither preidentified nor fixed in place.

The use of watercraft as exchange points raises questions:

- Which watercraft should be selected as an exchange point to minimize incident response time for a given evacuation request?
- How does a given watercraft exchange point affect a participating aircraft's ability to support future missions?
- How does aircraft airspeed, total casualty magnitude, distribution of patients between islands, expected number of patients per evacuation request, and severity of patient injury affect exchange point selections?

To enable this capability, we first model a representative dynamic environment with multiple aircraft and watercraft operating across and between two islands, then introduce a sequential decision making process and online solver to select optimal watercraft exchange points and determine aircraft dispatching. Our paper's main contributions are:

- Develop a model environment and semi-Markov decision process (SMDP) (Baykal-Gürsoy and Gürsoy 2007; Hu and Yue 2008) for the coupled watercraft exchange point selection and aircraft dispatch problems.
- Introduce an online planner that applies Monte Carlo Tree Search (MCTS) with root parallelization (Chaslot, Winands, and van Den Herik 2008) using patient evacuation threads from a casualty generation model.
- Adjust five model parameters to identify representative scenarios where watercraft exchange points would expedite transfers. Cruise speeds of aircraft with historic, current, or proposed evacuation roles are considered.
- Deploy the watercraft exchange point in Hawaii using two HH-60M Black Hawk helicopters and an Army Logistics Support Vessel (LSV). The dispatch of helicopters to transfer a mock patient was informed by our model, resulting in minimal delay during patient exchange.

RELATED WORK

This journal extension applies semi-Markovian planning to the multi-platform maritime evacuation coordination problem (Al-Husseini, Wray, Kochenderfer 2025). We develop a planning framework for the proposed watercraft exchange point capability, the underlying mechanisms of which are explored in a prior companion paper (Al-Husseini, Wray, and Kochenderfer 2024). There, we provide the problem description and operational considerations for air medical experts. In contrast, our research presented here details the artificial intelligence, computation, and modeling contributions as well as experimental results. The literature introduces adjacent efforts for medical evacuation resource allocation, characteristics of which feature prominently in our work.

McLay and Mayorga (2013) developed a linear programming method for dispatch, and consider several equity measures to ensure an equalized distribution of resources across various demand signals. Keneally, Robbins, and Lunday (2016) developed an MDP that dispatches medical evacuation aircraft in a land-based combat environment over an extended period of time to maximize overall system utility. Jenkins, Robbins, and Lunday (2018) build on efforts by Keneally, Robbins, and Lunday, introducing a monotonically decreasing function over time to more accurately reflect patient survival probability. Pettet et al. (2021) presented a hierarchical framework for partitioning large evacuation dispatch and allocation problems into tractable intertwined sub-problems. To overcome the sparsity of evacuation incidents available for consideration, they generate several casualty threads using an internal prediction model and instantiate individual MCTS trees for each. Equity considerations, a monotonically decreasing reward function, and generated casualty threads that align with parallel search trees all feature in this paper.

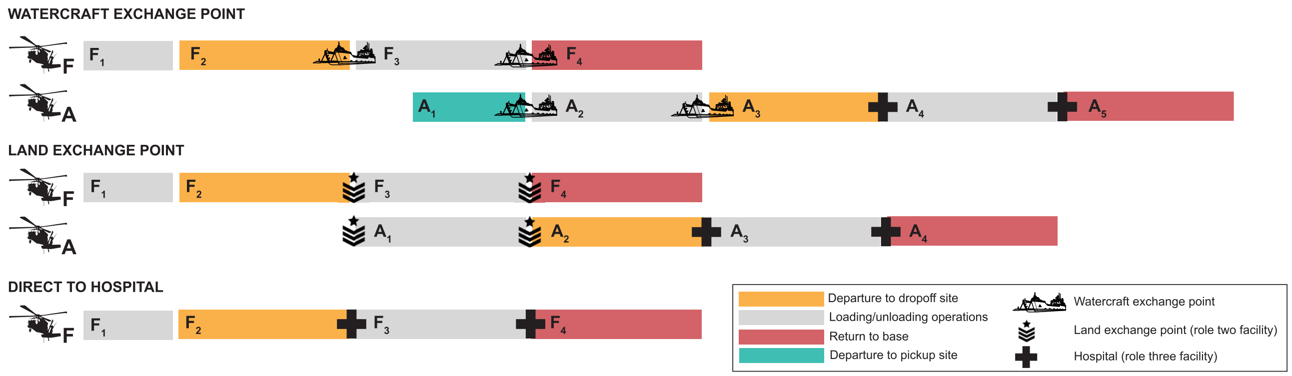

FIGURE 2 Action sequences for three categories of actions: watercraft exchange points, land exchange point, and direct to role three transfers. Time spans F_i are associated with the action sequence for the Kauai aircraft, while time spans A_i are associated with the action sequence for Oahu aircraft. F_1 includes the time required to pickup the patient from point of injury.

PROBLEM FORMULATION

Model environment

The model scenario involves a large-scale combat operation across the Hawaiian Islands of Oahu and Kauai. Patients are induced sequentially in accordance with a Poisson point process. Each patient is evacuated to the next higher role of care in accordance with Army doctrine (United States Army 2020). Evacuation requests flow from a role one (first responder) to a role two (forward resuscitative care) to a role three (theater hospitalization) (Cunningham et al. 2019). Patients begin at one of eight spatially distributed role one or role two facilities. Medical evacuation platoons with HH-60M Black Hawk helicopters are staged on both islands at role twos. We refer to the forward support medical evacuation platoon (FSMP) aircraft as the forward or Kauai aircraft, and the area support medical evacuation platoon (ASMP) aircraft as the rear or Oahu aircraft. FSMP aircraft are postured to coincide with demand, typically near troops engaged in combat or in high-population density areas. ASMP aircraft move patients between treatment facilities on an area basis. Three military watercraft, an LSV, a Landing Craft Utility (LCU), and an Expeditionary Fast Transport (EPF) travel between both islands at 10, 8, and 43 knots, respectively. As described in Figure 2, transfers may employ watercraft exchange points, land exchange points, or may fly directly to the role three. The distance between Oahu and Kauai is such that a typical rotary-wing aircraft can traverse it one way without requiring an intermediary transfer or refueling point, allowing fixed land and moving watercraft exchange points to be considered simultaneously. Each evacuation request contains the origin location, the number of patients to be transported, and their transport destination. On evacuation request receipt, the appropriate aircraft is dispatched to the patient pickup site. After collection, the aircraft delivers its patients to either an exchange point or the transport destination. The aircraft then returns to base where it refuels and subsequently enters an idle state. Please see Al-Husseini, Wray, and Kochenderfer (2024) for additional problem description details.

Semi-Markov decision process formulation

We model the interisland patient transfer process, which includes exchange point selection, as a fully-observable, discrete-time SMDP. The sojourn time τ is defined as the time between subsequent actions, may be thought of as the duration during which control of an agent is relinquished, and can be represented by a discrete random variable in discrete-time SMDPs. The SMDP can be formalized as $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, T, R, \rho, \tau \rangle$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, T is the state transition model, R is the immediate reward function, and ρ is the sojourn time reward rate (Puterman 2014). The SMDP dynamic influence diagram is shown in Figure 3. The transition model T assigns a probability of transitioning from an existing state s by action a to state s', with $\sum_{s \in S} P(s, a, s') = 1$. We can write the optimality criterion for the discrete-time SMDP as:

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^{\eta^t} \left(R(s^t, \pi(s^t)) + \sum_{\tau=0}^{\tau^t - 1} \gamma^{\tau} \rho(s^t, \pi(s^t), \tau) \right) \mid \pi, s^0 \right].$$

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. By consequence, natural process time η is the sum of sojourn times $\sum \tau^t + \epsilon$, where ϵ is the elapsed time since η^t. Similarly, sojourn time τ^t may be expressed as $\eta^{t+1} - \eta^t$, or the difference in natural process times for subsequent decision epochs. The discount factor γ is a hyperparameter

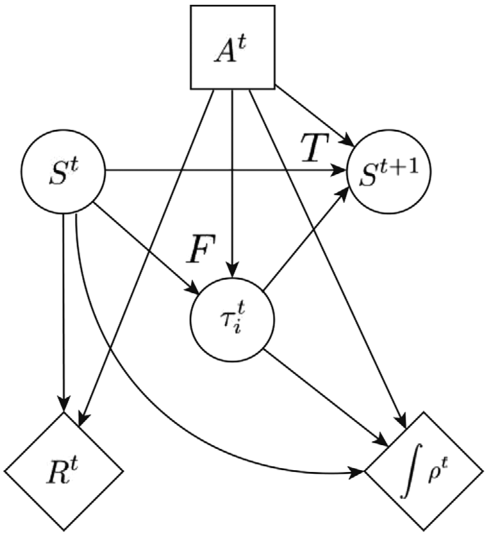

FIGURE 3 Dynamic influence diagram for the semi-Markov decision process, used to model the interisland patient transfer process, to include exchange point selection.

that balances short and long-term rewards. A larger γ prioritizes long-term rewards, whereas a smaller γ prioritizes short-term rewards. The resulting value function is:

$$V^{\pi}(s) = \mathcal{R}(s,\pi(s)) + \sum_{\tau^*=1}^{\infty} \gamma^{\tau} F(d\tau \mid s,\pi(s)) \sum_{s'\in S} T(s', \mid s,\pi(s),d\tau) V^{\pi}(s')$$

where expected reward $\mathcal{R}$ is the sum of immediate reward R and sojourn time reward rate ρ:

$$\mathcal{R}(s,a) = R(s,a) + \sum_{\tau=1}^{\infty} F(d\tau \mid s,a) \sum_{\tau^*=1}^{\tau-1} \gamma^{\tau^*} \rho(s,a,\tau^*).$$

We assume the presence of a centralized controller that selects from established exchange point pairing sequences in Figure 2 to transform the multiagent system into one that can be considered using an SMDP. Each state s in $\mathcal{S}$ includes the mission delay for participating aircraft, the likelihood a participating aircraft will have a maintenance fault and require replacement, the moving average utilization rate of participating aircraft, current and estimated future locations of participating watercraft and medical ships, and the considered evacuation request. The action space $\mathcal{A}$ includes the LSV exchange point, LCU exchange point, EPF exchange point, role two land exchange point, and direct to role three hospital. Sojourn time τ represents the time to resolve an interisland transfer using platforms assigned in a pairing sequence, and ρ is the sojourn reward rate generated by platforms not assigned to a pairing sequence instead tending to queued on-island evacuation requests.

Exchange points are selected at decision epochs coinciding with the receipt of evacuation requests. In environments with non-stationary behavior, exchange point selections may be revisited at the time of execution or when aircraft become available to service the request. This aligns decision-making with refreshed state information but forces the reevaluation of subsequent decisions at additional computational cost. In other words, the evolving status and behavior of watercraft and aircraft in response to battlefield conditions may require decision-making at the point of execution and not on request arrival. All evacuation requests are urgent and should therefore immediately be serviced or queued to an evacuation asset. This queuing can be later adjusted as needed.

We develop separate reward functions for greedy and optimized policies, the latter, which considers ρ for on-island evacuation requests. Both greedy and optimized policies employ a monotonically decreasing fusion of a Weibull survival function with a linear reward function, as shown in Figure 4. A survival function is the complimentary cumulative distribution over a lifetime, and it represents the probability a casualty will survive after a certain amount of time has elapsed following injury. The Weibull survival function extends an exponential distribution to the hazard rate (Zhang et al. 2011). This represents the non-linear nature of the physiological "golden hour" employed in military medical evacuation asset allocation policy. The golden hour reflects the belief that trauma patients are far more likely to survive if definitive care is provided within 60 min of sustaining an injury (Kotwal et al. 2016). Evacuation operations during large-scale combat must also consider battlefield clearance requirements (Hamilton 2020), here reflected in the linear component of our reward function. The generic per patient fused reward function R_f formulation follows, where t is time since evacuation request, m is the battlefield clearance slope, and λ and k are Weibull cumulative distribution function parameters:

$$R_f = \max\left(e^{-(\frac{t}{\lambda})^k}, -mt+1\right).$$

The greedy policy reward function is therefore the per patient fused reward function times the number of patients transported P, over the time differential for τ:

$$\mathcal{R}_{\text{greedy}} = \sum_{\tau=1}^{\infty} F(d\tau \mid s,a) \cdot R_f(\lambda=125, k=7, m=0.0042, t=\tau) \cdot P.$$

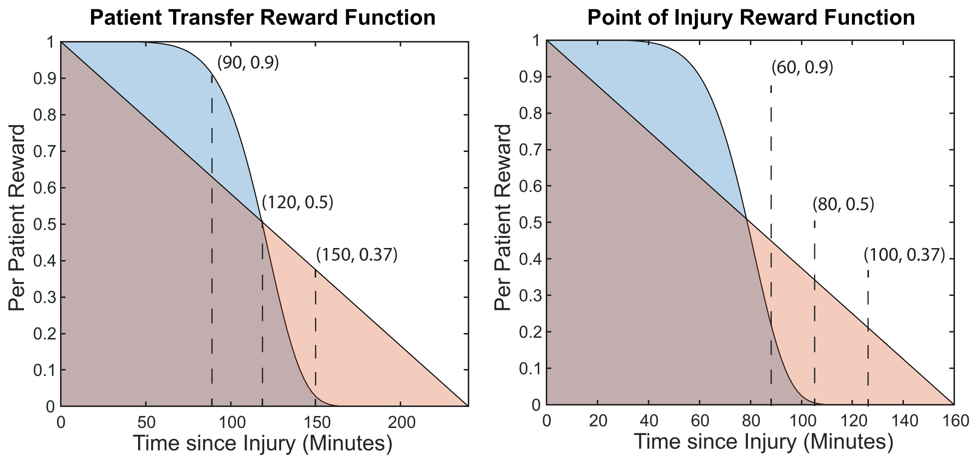

FIGURE 4 Fused reward functions capturing non-linear survivability estimates in blue and linear battlefield clearance requirements in red.

R_f applied with the above values for λ, k, and m indicates that a patient transfer completed within 90 min of lifesaving stabilization at a role two has a 90% probability of survival, whereas a patient transfer completed 120 min of lifesaving stabilization has a 50% probability of survival. The fusion reward function may be adjusted for injury severity, injury type, and means of stabilization, and per command-driven requirements for battlefield clearance during different phases of a military operation. Sojourn time τ and time differential $d\tau$ are dependent on the sum of incident response time T and the time Z between receiving the interisland evacuation request and dispatch of the assigned evacuation platform. T is determined per action sequences in Figure 2, and Z depends on queuing delays induced by aligning platforms on active missions to evacuation requests.

$$F(d\tau \mid s, a) = F(d\tau \mid T, Z)$$

the T formulation follows, where F is time for the forward aircraft, A is a time for the area aircraft, and exchange point E may be in the set of watercraft exchange points W or land exchange points L:

$$T = \begin{cases} [F_1 + F_2 + F_3 + A_2 + A_3 + A_4] & \text{if } E \in W \\ [F_1 + F_2 + F_3 + A_2 + A_3] & \text{if } E \in L \\ [F_1 + F_2 + F_3] & \text{otherwise.} \end{cases}$$

The reward function for optimized policies $\mathcal{R}_{\text{optimal}}$ considers accumulated rewards generated from on-island evacuation requests. Accumulated rewards capture the negative impact interisland patient transfer delays have on on-island evacuations and subsequent interisland patient transfers. The optimal reward $\mathcal{R}_{\text{optimal}}$ formulation follows, where $x_{\tau'}$ is an Oahu or Kauai on-island evacuation request in the set of requests $\mathcal{X}_{\tau'}$ instantiated at incrementing discrete time step τ':

$$\mathcal{R}_{\text{optimal}} = \mathcal{R}_{\text{greedy}} + \sum_{\tau=1}^{\infty} F(d\tau \mid s, a) \sum_{\tau'=1}^{\tau-1} \gamma^{\tau'} \sum_{x_{\tau'} \in \mathcal{X}\tau'} R_f(\lambda = 63, k = 7, m = 0.0063, t = t_{x_{\tau'}}) \cdot P_x$$

where $t_{x_{\tau'}}$ is the time between the on-island injury occurring and the patient arriving at the on-island role two.

Accumulated rewards may be further supplemented with a request-agnostic sojourn penalty rate Y. Y penalizes each aircraft h in the set of participating aircraft H for time of employment T_h relative to historic aircraft utilization U_h. T_h, like T, is a consequence of exchange point selection. The penalty formulation follows, where ω_1 and ω_2 are tunable parameters:

$$T_1 = [F_1 + F_2 + F_3 + F_4]$$

$$T_2 = \begin{cases} [A_1 + A_2 + A_3 + A_4 + A_5] & \text{if } E \in W \\ [A_1 + A_2 + A_3 + A_4] & \text{if } E \in L \\ [0] & \text{otherwise} \end{cases}$$

$$Y = \omega_1 T_1 U_1 + \omega_2 T_2 U_2.$$

Monte Carlo Tree Search

An expansive military campaign across a non-contiguous maritime environment with multiple heterogeneous evacuation platforms results in complex patient evacuation dynamics that prove difficult to model in closed form.

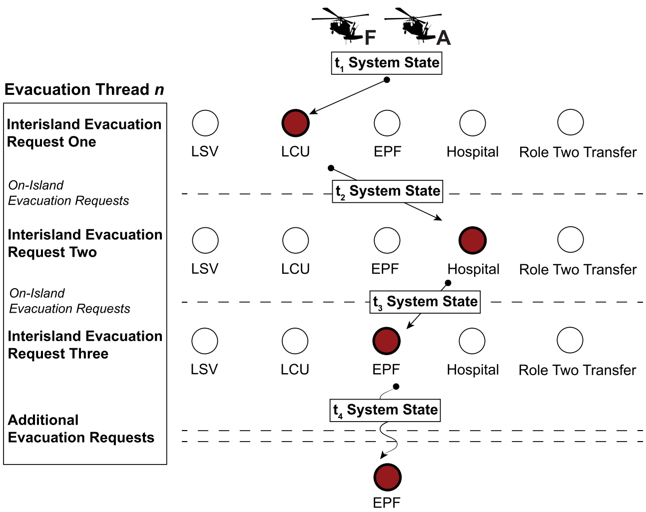

FIGURE 5 Monte Carlo Tree Search with root parallelization to solve the interisland patient transfer dispatch problem.

These evacuation models motivate the use of probabilistic search algorithms, such as MCTS, in tandem with a generative environment. MCTS is an anytime algorithm, meaning it can return a valid, if sub-optimal, solution at any point in its run-time while also adapting to changes in the model environment (Browne et al. 2012). MCTS, visualized in Figure 5, represents the set of possible actions as edges in a search tree. The nodes in the search tree are then the ensuing states. MCTS expands the search tree using random sampling or a predefined heuristic informed by the problem domain. MCTS explores the search tree asymmetrically such that the most promising actions are prioritized. State value estimation occurs by simulating a play-out from a node to the end of a predefined planning horizon. We apply the standard Upper Confidence bound for Trees (UCT) algorithm (Kocsis, Szepesvári, and Willemson 2006) to balance exploration versus exploitation in the tree search policy and determine the value of each node n, as follows:

$$\mathrm{UCT}(n) = \bar{u}(n) + c\sqrt{\frac{\log(\mathrm{visits}(n'))}{\mathrm{visits}(n)}}$$

such that $\bar{u}(n)$ is the value of the state at n, c is an exploration constant that adjusts the balance between exploration of new nodes and exploitation of previously visited nodes, $\mathrm{visits}(n)$ is the visitation count for n, and $\mathrm{visits}(n')$ is the visitation count for the parent node of n, n'. A default, heuristic policy is typically applied to estimate action value during search tree roll-outs. We introduce a greedy policy for roll-outs that selects the exchange point which, based on proximity alone and without accurate state information, minimizes transfer time for the evacuation request under consideration.

TABLE 1 Environmental parameters.

Parameter	Values
Action space	$[\mathcal{A}_1, \mathcal{A}_2]$
Dispatch policy	[Greedy, Optimal]
Casualty magnitude multiplier	[0.6, 0.7, …, 1.4]
Platoon casualty ratio	[0.6, 0.7, …, 1.4]
Proportion of patient transfers	[0.1, 0.2, 0.3, 0.4, 0.5]
Aircraft airspeed [kn]	[120, 130, 150, 250, 280]
Patients per evacuation request	[2, 3, 4, 5, 6]

Root parallelization

Root parallelization is a process whereby multiple independent search trees are constructed and solved using MCTS (Chaslot, Winands, and van Den Herik 2008). The search trees differ in composition due to stochastic sampling during construction. Results are then aggregated across all search trees to select an optimal course of action. In this paper, search trees are aligned to evacuation request "threads" generated from simulated casualty data. Search tree "scores" may be combined across search trees that begin with a particular action, in which case the action with the greatest summed score is selected. Sampling a single evacuation request thread proves insufficient, due to significant environmental uncertainty. Root parallelization paired with casualty generation data can provide robust results even when actual evacuation requests in theater are sparse.

Parameters

Environmental parameters can be grouped into evacuation support and evacuation request parameters, shown in Table 1. We introduce two support parameters: exchange point selection action space and dispatch policy. SMDP effectiveness is evaluated for action spaces $\mathcal{A}_1$ and $\mathcal{A}_2$. $\mathcal{A}_1$ includes all moving watercraft exchange points, fixed land exchange points, and the direct-to-role three transfer. $\mathcal{A}_2$ features only fixed land exchange points and the direct-to-role three transfer. We compare a greedy policy against a root-parallelized MCTS solver with accumulated rewards. The main source of environmental uncertainty is the ebb and flow of patients across the battlefield, which we characterize using evacuation request parameters. These include casualty magnitude multiplier,

TABLE 2 MCTS with root parallelization hyperparameters.

Hyperparameter	Values
Thread count	[4, 6, 8, 10, 12, 14, 16]
Thread duration	[4 h, 6 h, 8 h, 10 h, 12 h, 14 h, 16 h]
Discount factor	[0.90]
Exploration constant	[1.0]

platoon casualty ratio, proportion of patient transfers, and patients per evacuation request. Also considered the cruise airspeed of the aircraft servicing the evacuation requests.

We experiment with four hyperparameters for MCTS with root parallelization: thread count, thread duration, discount factor, and exploration constant. Considered values for each are shown in Table 2. A two-dimensional grid search compares thread count and duration against total reward, for a given optimal policy. Ten threads with a duration of 10 h each was selected for all simulations.

EXPERIMENTS

The introduced watercraft exchange point model and planner are first evaluated in simulation, then deployed in a military exercise with an actual LSV and two HH-60M Black Hawks. This section covers the simulation results, deployment and deployment insights, and capability discussion.

The model environment was simulated 20 times for each configuration of Table 1 parameters to determine the impact watercraft exchange points have on total rewards and incident response times by platoon, as shown in Figure 6. Considered is how each parameter affects the ratio of patient transfers conducted via watercraft versus either land exchange points or direct transfer. Unless otherwise stated, the casualty magnitude multiplier is 1.0, the platoon casualty ratio is 1.4, proportion of patient transfers is 0.25, aircraft speed is 150 knots, and patients per evacuation request is three. All ranges reflect a 95% confidence interval.

Total casualty magnitude

Increasing the number of casualties increases total rewards. This increase is greater for an optimal policy with watercraft exchange points than for an optimal policy without watercraft exchange points. However, the evacuation system reaches a saturation point, and total rewards begin to decrease across all three considered policies above a magnitude multiplier of 1.2. A casualty magnitude multiplier of 1.0 indicates that 96 patients were transferred through 32 evacuation requests distributed over 24 h. Incidence response times increase dramatically with an increase in casualty magnitude for both an optimal policy without watercraft and a greedy policy. Incidence response times increased slightly with an increase in casualty magnitude for an optimal policy with watercraft. The ratio of patient transfers moved via watercraft does not change with an increase in casualty magnitude.

Distribution of patients across islands

A geographic shift in casualty distribution significantly affects exchange point selection. Total reward and platoon incident response times remain steady if the majority of casualties are on Oahu and serviced by the ASMP aircraft. Watercraft exchange point availability increases total reward and reduces response times when the casualty magnitude proportion multiplier increases above 1.1. Watercraft exchange points are regularly employed by an optimal policy when the forward aircraft is sufficiently burdened relative to the rear aircraft. Between a 0.6 and 1.0 proportion multiplier, 0% of patient transfers employed watercraft exchange points. By 1.4, more than 95% of patient transfers employed watercraft exchange points. An optimal policy with watercraft exchange points outperforms both an optimal policy without watercraft exchange points and a greedy policy when the number of casualties on the forward island of Kauai exceeds that on the rear island of Oahu by 25% or more.

Number of patients per evacuation request

An increase in the number of patients transported per evacuation request, and a corresponding reduction in the number of total evacuation requests, increases total reward across optimal policies with and without watercraft exchange points. This is true for up to four patients, or two-thirds of an HH-60M's cabin capacity. The total reward for an optimal policy then begins to decrease for five or more patients per request. Watercraft exchange point availability results in a significant increase in total rewards when the number of patients per evacuation request averages three or four. The number of patients per evacuation request has minimal impact on the ratio of patient transfers moved via watercraft for an optimal policy with watercraft exchange points.

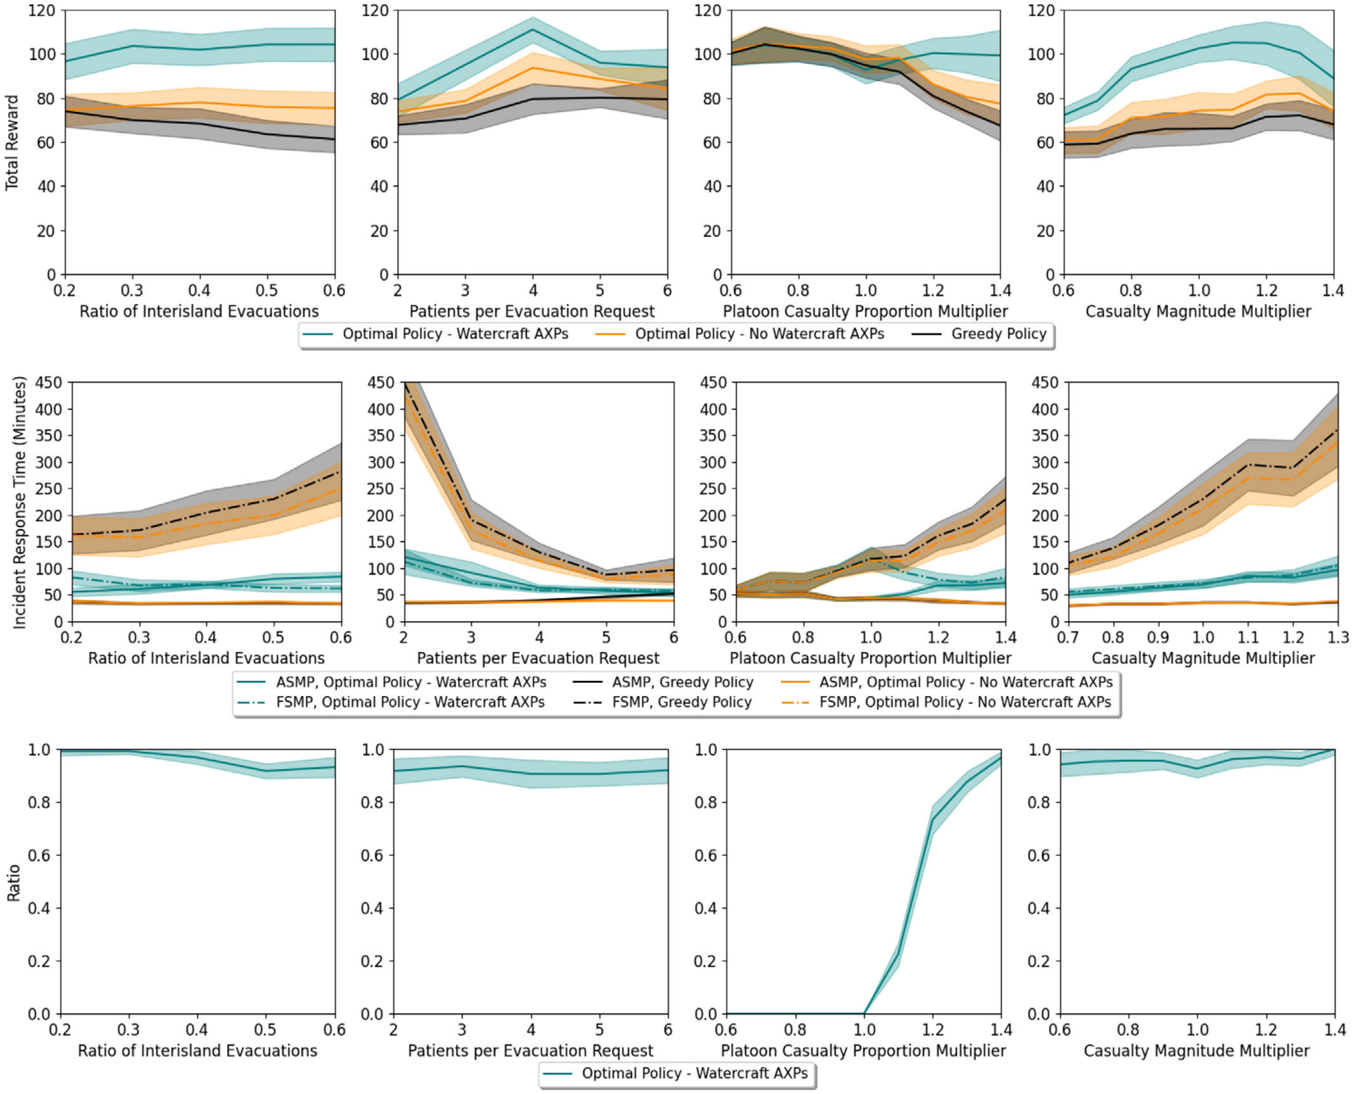

FIGURE 6 Casualty magnitude, distribution of patients, patients per evacuation request, and proportion of interisland transfer patients impact on total rewards and incident response time by platoon for various exchange point policies, and ratio of patient transfers moved via watercraft given an optimal watercraft exchange point policy.

Proportion of interisland transfer patients

A change in the proportion of interisland transfer patients has a negligible impact on total rewards for both optimal policies with and without watercraft exchange points. There is a weak inverse relationship between the proportion of patients requiring interisland transfer and total rewards for a greedy policy. For an optimal policy with watercraft exchange points, an increase in the percentage of interisland transfers results in a decrease in FSMP response times and an increase in ASMP response times. The reverse is true for both an optimal policy without watercraft exchange points and a greedy policy. Changing the proportion of interisland transfer patients has little effect on the use of watercraft exchange points. An optimal policy with watercraft exchange points significantly outperforms both an optimal policy without watercraft exchange points and a greedy policy across the full range of values considered.

Aircraft airspeed

Figure 7 shows how watercraft exchange points affect incident response times across combinations of total casualty magnitude and aircraft cruise airspeeds. FSMP and ASMP aircraft are assumed to be homogeneous assets with the same cruise airspeeds. As discussed, watercraft exchange points reduce the dramatic spike in response times resulting from an increased number of casualties. The faster the aircraft cruise airspeed, however, the less impact watercraft exchange points have on incident response times.

	Fewer Total Casualties			More Total Casualties		
Medical Aircraft	-30%	-20%	-10%	+10%	+20%	+30%
UH-1 Iroquois 120 knots	0.2	0.3	0.34	0.43	0.45	0.5
MH-65 Dolphin 130 knots	0.19	0.2	0.26	0.36	0.45	0.44
UH-60 Black Hawk 150 knots	0.11	0.15	0.26	0.32	0.36	0.49
SB-1 Defiant 250 knots	0.033	0.041	0.065	0.083	0.13	0.16
V-280 Valor 280 knots	0.013	0.022	0.031	0.076	0.069	0.14

FIGURE 7 The impact using watercraft exchange points has on incident response times across aircraft cruise speeds and casualty loads. A cell value of 0.5 indicates a 50% improvement in incident response times when watercraft exchange points are employed relative to when they are not. Five aircraft are considered, each with a historic, current, or proposed role in medical evacuation.

Watercraft exchange points reduce average incident response times by 50% for two Bell UH-1 Iroquois and a 30% increase in total casualties. In comparison, the SB-1 Defiant and Bell V-280 Valor minimally reduce incident response times when there is a 30% increase in casualties and watercraft exchange points are available (16% and 14%, respectively).

DEPLOYMENT

We move from simulation to real-world deployment by facilitating a patient transfer with a manikin in Hawaii on October 2023 using two HH-60M Black Hawk helicopters and Army LSV-3 *General B. Somervell* as a watercraft exchange point. This deployment occurred as part of a larger Army exercise to evaluate novel medical evacuation capabilities. Both helicopters were launched from Wheeler Army Airfield, located in central Oahu. The mock patient was picked up from the ground force three miles east of Wheeler Army Airfield, transported to and from LSV-3 via helicopter as shown in Figure 8, and ultimately delivered to Tripler Army Hospital. An execution checklist with helicopter dispatch times, informed by our decision process and solver, was provided to command posts located on Wheeler Army Airfield and LSV-3. Figure 9 depicts the real-world inputs to our decision process and solver. These include the number of transportable patients in the medical evacuation request, participating evacuation platform characteristics, and real-time updates on platform locations and delays induced by non-participating air and maritime traffic. LSV-3 was positioned ten miles south of Honolulu and moved in a south-westerly direction at 5 knots during patient drop-off and pickup.

Figure 10 compares real-world deployment outcomes to the simulation results, which informed the real-world helicopter dispatch times. Boxes one and three annotate helicopter airfield departures. Box two annotates the first

(A) © Tristan Moore 2023 (B) © Tristan Moore 2023

FIGURE 8 An HH-60M Medical Evacuation Black Hawk helicopter executes hoist iterations to Army Logistics Support Vessel 3, which is traversing the open ocean south of Honolulu, Hawaii at 5 knots. These operations were part of a two-week patient transfer exercise demonstrating the use of Army watercraft as ambulance exchange points.

helicopter's arrival at LSV-3. Box four annotates when the patient transfer occurs. Only 3 min separated the first helicopter departing LSV-3 after dropping off the patient from the second helicopter arriving at LSV-3 to pick up the patient. Minimizing this difference ensures medical continuity of care and is a key benefit of the introduced decision strategy. The incident response time from evacuation request notification to role three patient delivery was 76 min, 16 min of which were caused by maritime traffic delays, the impact of which is annotated by the asterisk box with bars. Box six annotates the first helicopter's return to the airfield, box five annotates the second helicopter's departure from the watercraft, and box seven annotates the second helicopter's arrival at the role three.

Air and maritime traffic delays present a serious challenge to deployment dispatch timing. Air traffic delays were deemed likely due to the high volume of aircraft in the Honolulu Class B airspace, which both helicopters traversed while transitioning through Pearl Harbor enroute to LSV-3. Similarly, the proximity of LSV-3 to Pearl Harbor suggested the presence of obstructing maritime traffic hindering patient transfer operations. We design the aircraft dispatching system to reevaluate past decision-making when critical state information is updated, thereby minimizing impact to operations. Air and maritime traffic delays are coded to increase select time spans in the action sequences presented in Figure 2. These updates were programmed in by a command and control authority in the command post on LSV-3, who then relayed updated instructions to the helicopters through the command post on Wheeler Army Airfield.

During deployment, a non-affiliated maritime vessel approached LSV-3 about 8 min after evacuation request

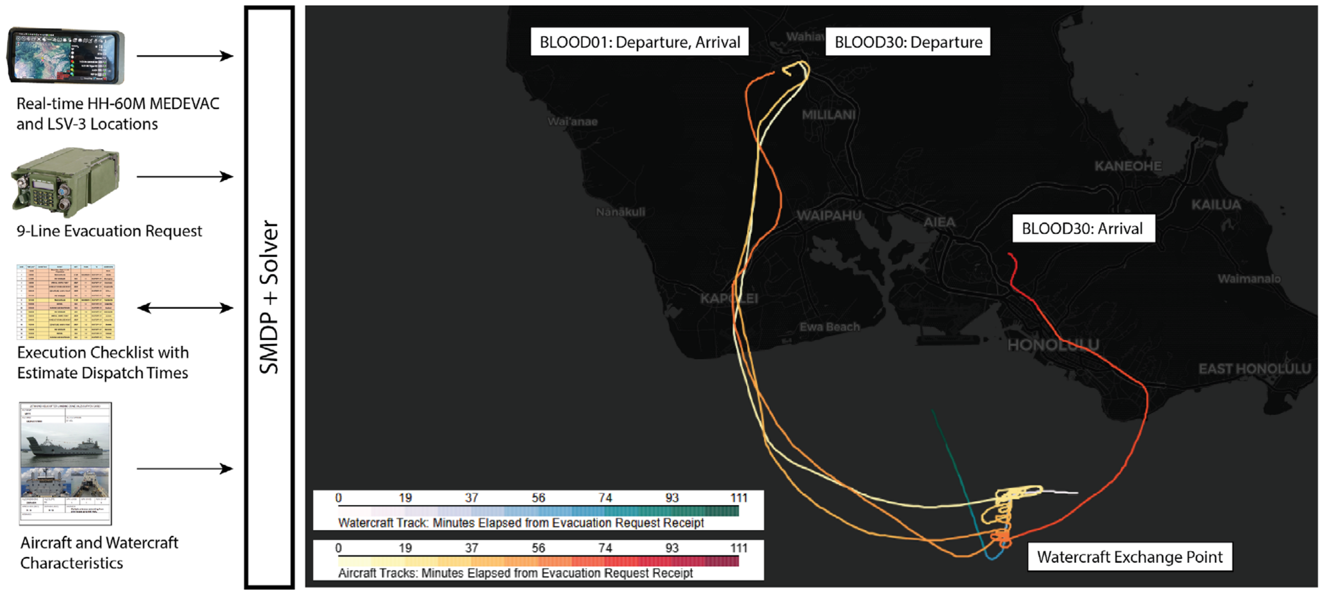

FIGURE 9 Schematic architecture of the decision process and solver, and resulting ground track of two HH-60M Black Hawks (call signs BLOOD01 and BLOOD30) and the participating Logistics Support Vessel during deployment on October 11, 2023.

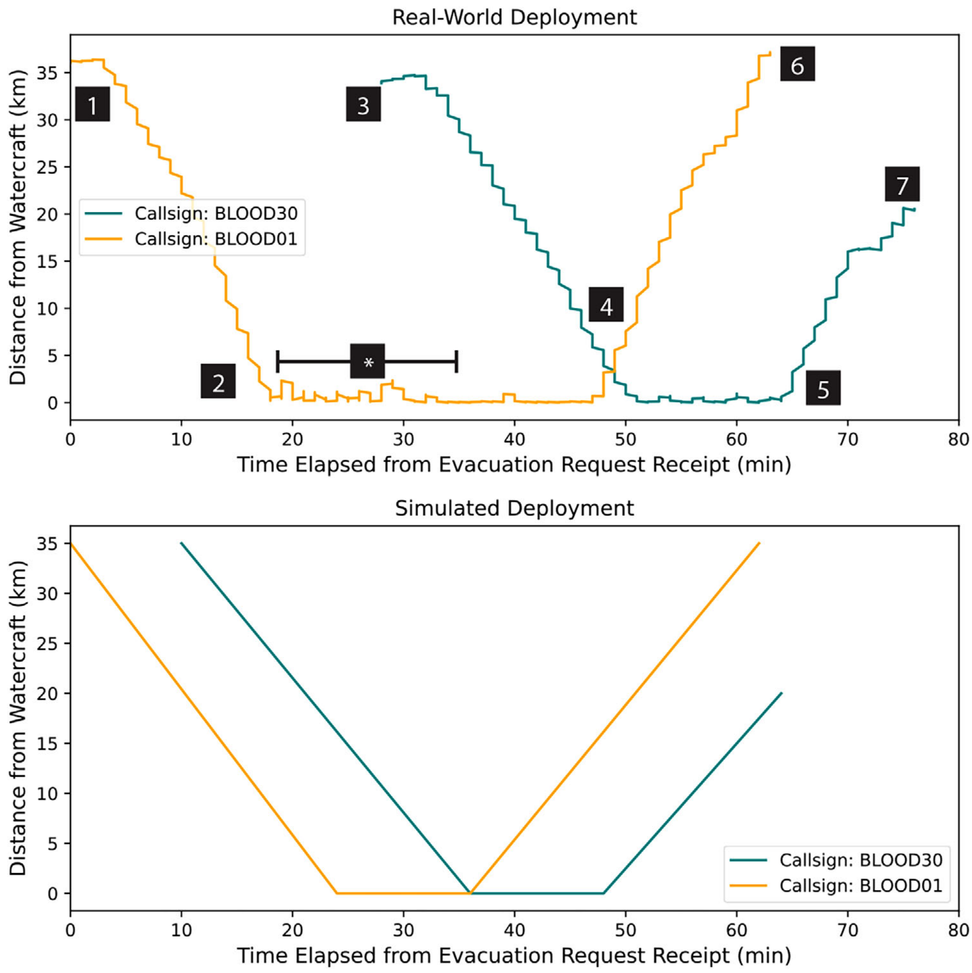

FIGURE 10 Annotated comparison of deployment and simulation evacuation times and distances from the underway watercraft exchange point. Key insight: the simulation model accurately captures the real-world dynamics.

receipt. LSV-3 responded by adjusting its course south, which delayed the first helicopter from conducting patient unloading operations. The delay start and stop time was communicated by the LSV-3 crew and used to adjust time span F_3, which also triggered a new decision epoch and subsequent search. The second helicopter's airfield departure was then automatically delayed commensurate to delays experienced by the first helicopter. This built-in update infrastructure enables continuous evacuation operations despite the non-stationary and unpredictable nature of the operating environment. The window of opportunity for amending dispatch decisions can be exceedingly short; during deployment, the decision to delay the second helicopter came only 2 min prior to the intended taxi time. The update infrastructure presented resembles a SMDP, where some decision epochs align with probabilistic state updates rather than resolved actions, and are in a sense, interruptions. Our approach to updates is inspired by the τ_{any} synchronous termination scheme used in concurrent decision making for multi-agent SMDPs (Rohanimanesh 2006).

This deployment resulted in the watercraft exchange point capability being recommended for addition to Army doctrinal publication ATP 4-02.2, *Medical Evacuation*. This establishes the watercraft exchange point as a standard technique for 300 plus Army medical evacuation helicopters in more than 25 air ambulance companies and detachments (Figures 9, 10).

DISCUSSION

An optimal policy with watercraft exchange points significantly outperforms an optimal policy without watercraft exchange points across specified ranges for all five examined casualty flow parameters in terms of both total rewards and incident response times. The presence of watercraft exchange points generally results in a 30%–35% improvement in total rewards. Similarly, an optimal policy without watercraft exchange points significantly outperforms a greedy policy across specified ranges for three of the five parameters examined, demonstrating a more modest 10%–15% improvement. We find a trade-off between reducing FSMP response times and increasing ASMP response times, such that the time reduced for the FSMP greatly outweighs the time increased for the ASMP. This imbalance indicates a geographically derived inefficiency in the dynamic resource allocation problem. Different geographies and operational scenarios will result in unique demand imbalances that may fluctuate over the course of a battle. Watercraft exchange points then play a balancing role such that the time required to service a particular evacuation request can be partitioned to best support each island's unique utilization.

Impressively, in nearly all instances, an optimal policy with watercraft exchange points minimizes the disparity in incident response time between platoons to near zero. That this occurred without say over the location and movement of participating watercraft suggests that a certain watercraft density may be a sufficient substitute for watercraft control. Without watercraft exchange points, even an optimal policy cannot bridge imbalances in demand, and the disparity in response time between platoons balloons up to hundreds of minutes for certain parameter configurations. Partitioning the transfer between ambulances comes at a cost: the patient handover process at any exchange point increases the total time to deliver a patient to their destination. This cost is often outweighed by the benefit induced by partitioning.

Several insights were obtained during the deployment that support the future implementation of decision support tools for cooperative aerial and maritime evacuation planning.

- **Continuous monitoring and evaluation**: Evacuation operations may occur in high-density traffic areas that induce substantial air and maritime traffic delays. These delays increase incident response times and negatively affect patient outcomes. Delays often break the medical continuity of care by leaving patients awaiting pickup on watercraft for extended periods of time. Understanding the likelihood and impact of delays by type and over time can inform model updates, resulting in refreshed decision-making that lessens operational impacts. Future delays may be identified automatically by evaluating real-time ADS-B and AIS data instead of relying on a human-in-the-loop as done during deployment.
- **Flexibility, adaptability, and risk**: The coordination of multiple platforms with asynchronous actions is susceptible to sudden changes in the demand signal or the environment. Update infrastructure must gauge risk when determining whether to amend existing guidance. Interruptions, such as those employed during deployment, should be carefully timed and sparsely used so as to not overwhelm human operators. Trust in automated systems can quickly degrade with sudden changes in guidance, especially in domains as sensitive as aeromedical flight.
- **Communication and contingencies**: Assuming full observability requires platforms to communicate intentions and state information on a regular and predictable basis, whether via automated real-time data sharing through mobile applications or reported updates from crews. We successfully mitigated delays during

deployment because helicopter and watercraft crews communicated relevant observations over satellite internet and VHF radio to a trained command and control authority. Lossy (Tung et al. 2021), delayed (Oliehoek and Spaan 2012), and costly (Goldman and Zilberstein 2008) communication severely limit the strong centralization assumption held throughout this paper, and complicate decision-making. Decentralized multi-agent and partial-observability models should be considered for when communication fails. Communication failures are especially common in the Indo-Pacific due to area-wide and targeted jamming (Grant et al. 2009).

FUTURE WORK

We identify five opportunities for future research that will further validate the watercraft exchange point with accompanying decision strategy as an established maritime evacuation technique. The first is an improved casualty generator that draws from relevant historical casualty databases. The current scenario would benefit from casualty evacuation threads for root parallelization based on data from the Central Pacific Guadalcanal Campaign during World War II. This may be accomplished by integrating the present application with military casualty estimation software like the Joint Medical Planning Tool (JMPT) and Medical Planners' Toolkit (MPTk). An extension of this approach is applying the watercraft exchange point concept to models of historical conflicts and comparing historical and simulated casualty data. The second opportunity is extending the present work to encompass degraded communication and partial-observability models. The third involves understanding watercraft density requirements. Our simulation model assumes three different military watercraft operating between islands of interest, but this number could conceivably be lower given the vastness of the Indo-Pacific. Similarly, the types of watercraft available, and their cruising speeds and carrying capacities, will likely impact the magnitude of benefits cited in our discussion. The fourth is to explore the use of controllable watercraft exchange points. The watercraft considered here, the LSV, LCU, and EPF, are logistics transport vessels and are therefore routed independently of a medical command and control authority. The Expeditionary Medical Ship (EMS) is the medical-variant of the EPF and will be fielded for the first time in 2026. The EMS may be strategically routed to enhance exchange point operations between islands. This is essentially a combined multi-agent routing and allocation problem. The fifth is considering a hierarchical approach to planning to enable scaling. Large-scale combat operations are sprawling conflicts with large aerial and maritime fleets and are characterized by extremely high casualty numbers.

CONCLUSION

Watercraft exchange points increase evacuation flexibility and reach in maritime environments. We demonstrate how watercraft exchange points enable patient transfers that correct for imbalances in aircraft utilization across islands. This results in significantly reduced incident response times. Although advantageous, watercraft exchange points are complicated by their non-dedicated and underway nature. The location and movement of watercraft and use of participating aircraft must be considered. We formulate an SMDP according to a relevant operational scenario and simulate using MCTS with root parallelization to solve the exchange point selection and aircraft dispatching problems. Platoon response times are determined for five parameters corresponding to the casualty flow and aircraft characteristics. Evacuation planning with watercraft exchange points results in a 35%–40% improvement over existing methods. We deploy the watercraft exchange point model in the Hawaiian Islands using two HH-60M MEDEVAC helicopters and an Army LSV, which led to proposed revisions to Army doctrinal publications on medical evacuation.

The views and conclusions contained herein are those of the authors alone and should not be interpreted as necessarily representing the official policies or endorsements, expressed or implied, of the United States Department of Defense.

ACKNOWLEDGMENTS

This research was supported by the United States Army LSV-3 and LSV-7 watercraft crews and 545th Harbormasters. We are also grateful to C/3-25 Aviation Regiment "Lightning DUSTOFF" based at Wheeler Army Airfield, Hawaii. Their HH-60M aircrews trained for months in the lead-up to deployment, during which they executed the watercraft exchange point to great effect.

This research received no external funding. There are no conflicts of interest to disclose.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

ORCID

Mahdi Al-Husseini https://orcid.org/0009-0003-9087-814X
Kyle H. Wray https://orcid.org/0000-0001-6986-9941
Mykel J. Kochenderfer https://orcid.org/0000-0002-7238-9663

REFERENCES

Al-Husseini, M., K. H. Wray, and M. J. Kochenderfer. 2024. "Watercraft as Overwater Ambulance Exchange Points to Enhance Air Medical Evacuation." *Air Medical Journal* 44(1): 23–29.

Al-Husseini, M., K. H. Wray, and M. J. Kochenderfer. 2025. "Semi-Markovian Planning to Coordinate Aerial and Maritime Medical Evacuation Platforms." *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 28.

Baykal-Gürsoy, M., and K. Gürsoy. 2007. "Semi-Markov Decision Processes: Nonstandard Criteria." *Probability in the Engineering and Informational Sciences* 21(4): 635–57.

Browne, C. B., E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. 2012. "A Survey of Monte Carlo Tree Search Methods." *IEEE Transactions on Computational Intelligence and AI in Games* 4(1): 1–43.

Chaslot, G. M. B., M. H. Winands, and H. J. van Den Herik. 2008. "Parallel Monte-Carlo Tree Search." in *International Conference on Computers and Games*, 60–71. Berlin, Springer.

Cunningham, C. W., D. E. Keen, S. G. Schauer, C. U. Kharod, and R. A. De Lorenzo. 2019. "Military Casualty Evacuation: MEDEVAC." *In Aeromedical Evacuation: Management of Acute and Stabilized Patients*, 21–40. Berlin, Springer.

Goldman, C. V., and S. Zilberstein. 2008. "Communication-Based Decomposition Mechanisms for Decentralized MDPs." *Journal of Artificial Intelligence Research* 32: 169–202.

Grant, A., P. Williams, N. Ward, and S. Basker. 2009. "GPS Jamming and the Impact on Maritime Navigation." *The Journal of Navigation* 62(2): 173–87.

Hamilton, M. C. T. 2020. "Clearing the Battlefield of Wounded in Large-Scale Combat Operations." *Technical report, U.S. Army Command and General Staff College*.

Hu, Q., and W. Yue. 2008. "Semi-Markov Decision Processes." In *Markov Decision Processes with Their Applications*, 105–20. Berlin, Springer.

Jenkins, P. R., M. J. Robbins, and B. J. Lunday. 2018. "Examining Military Medical Evacuation Dispatching Policies Utilizing a Markov Decision Process Model of a Controlled Queueing System." *Annals of Operations Research* 271: 641–78.

Keneally, S. K., M. J. Robbins, and B. J. Lunday. 2016. "A Markov Decision Process Model for the Optimal Dispatch of Military Medical Evacuation Assets." *Health Care Management Science* 19: 111–29.

Kocsis, L., C. Szepesvári, and J. Willemson. 2006. "Improved Monte-Carlo Search." *Technical report, University of Tartu, Institute of Computer Science*.

Kotwal, R. S., J. T. Howard, J. A. Orman, B. W. Tarpey, J. A. Bailey, H. R. Champion, R. L. Mabry, J. B. Holcomb, and K. R. Gross. 2016. "The Effect of a Golden Hour Policy on the Morbidity and Mortality of Combat Casualties." *JAMA Surgery* 151(1): 15–24.

McLay, L. A., and M. E. Mayorga. 2013. "A Dispatching Model for Server-to-Customer Systems That Balances Efficiency and Equity." *Manufacturing & Service Operations Management* 15(2): 205–20.

Oliehoek, F., and M. Spaan. 2012. "Tree-based solution methods for multiagent POMDPs with delayed communication." In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 26, 1415–21. Toronto, ON, Canada: AAAI Technical Track: Multiagent Systems.

Pettet, G., A. Mukhopadhyay, M. J. Kochenderfer, and A. Dubey. 2021. "Hierarchical Planning for Resource Allocation in Emergency Response Systems." In *ACM/IEEE International Conference on Cyber-Physical Systems*, 155–66. New York, NY: Association for Computing Machinery.

L Puterman, M. 2014. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Hoboken, NJ: John Wiley & Sons.

Rohanimanesh, K. 2006. *Concurrent Decision Making in Markov Decision Processes*. Amherst, MA: University of Massachusetts Amherst.

Tung, T.-Y., S. Kobus, J. P. Roig, and D. 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): 2590–603.

United States Department of the Army. 2019. "4-02.2 Medical Evacuation." *Army Techniques Publication*, United States Army.

United States Department of the Army. 2020. "Army Health System." *Field Manual*, United States Army.

Zhang, K., R. Wu, H. Du, X. Zhang, and X. Tao. 2011. "Research on Survival Model of Casualty in Wartime." In *2011 IEEE International Symposium on IT in Medicine and Education*, vol. 1, 473–75. New York, NY: IEEE.

How to cite this article: Al-Husseini, M., K. H. Wray, and M. J. Kochenderfer. 2025. "Semi-Markovian planning to coordinate aerial and maritime medical evacuation platforms." *AI Magazine* 46: e70023. https://doi.org/10.1002/aaai.70023

AUTHOR BIOGRAPHIES

Mahdi Al-Husseini is a Ph.D. student in the Aeronautics and Astronautics department at Stanford University and an active-duty HH-60M Black Hawk helicopter pilot-in-command in the United States Army. He develops models and algorithms for semi-decentralized multi-agent systems in communication constrained environments. He is the inventor on more than 30 patents and patent applications related to military aviation. He earned masters' degrees in aeronautics and astronautics and computer science at Stanford and the Georgia Institute of Technology.

Kyle Wray is a Research Assistant Professor at Northeastern University in the Khoury College of Computer Sciences and the Institute for Experiential Robotics. He earned his doctorate in computer science from the University of Massachusetts Amherst, advised by Shlomo

Zilberstein, specializing in artificial intelligence and autonomous robots. Kyle earned his bachelors' and masters' degrees in computer science, engineering, and mathematics from the University of Massachusetts Amherst and The Pennsylvania State University.

Mykel Kochenderfer is an Associate Professor of Aeronautics and Astronautics and Associate Professor, by courtesy, of Computer Science at Stanford University and senior fellow at the Institute for Human-Centered AI. He is the director of the Stanford Intelligent Systems Laboratory (SISL), conducting research on advanced algorithms and analytical methods for the design of robust decision making systems. Of particular interest are systems for air traffic control, unmanned aircraft, and automated driving where decisions must be made in uncertain, dynamic environments while maintaining safety and efficiency.