

H-TD²: Hybrid Temporal Difference Learning for Adaptive Urban Taxi Dispatch

Benjamin Rivière¹, *Graduate Student Member, IEEE*, and Soon-Jo Chung¹, *Senior Member, IEEE*

Abstract—We present H-TD²: Hybrid Temporal Difference Learning for Taxi Dispatch, a model-free, adaptive decision-making algorithm to coordinate a large fleet of automated taxis in a dynamic urban environment to minimize expected customer waiting times. Our scalable algorithm exploits the natural transportation network company topology by switching between two behaviors: distributed temporal-difference learning computed locally at each taxi and infrequent centralized Bellman updates computed at the dispatch center. We derive a regret bound and design the trigger condition between the two behaviors to explicitly control the trade-off between computational complexity and the individual taxi policy’s bounded sub-optimality; this advances the state of the art by enabling distributed operation with bounded-suboptimality. Additionally, unlike recent reinforcement learning dispatch methods, this policy estimation is adaptive and robust to out-of-training domain events. This result is enabled by a two-step modelling approach: the policy is learned on an agent-agnostic, cell-based Markov Decision Process and individual taxis are coordinated using the learned policy in a distributed game-theoretic task assignment. We validate our algorithm against a receding horizon control baseline in a Gridworld environment with a simulated customer dataset, where the proposed solution decreases average customer waiting time by 50% over a wide range of parameters. We also validate in a Chicago city environment with real customer requests from the Chicago taxi public dataset where the proposed solution decreases average customer waiting time by 26% over irregular customer distributions during a 2016 Major League Baseball World Series game.

Index Terms—Real-time taxi dispatch, adaptive systems, multi-agent systems, distributed decision-making, autonomous vehicles.

I. INTRODUCTION

COORDINATING a large fleet of automated taxis in complex and dynamic urban environments is an anticipated challenge for transportation network companies such as Uber, Lyft, Waymo, and Tesla. A typical urban mobility problem for these companies is taxi dispatch, where a fleet of taxis service customers and the remaining, idle taxis are coordinated with a dispatch algorithm to minimize the customer waiting time of future requests. In practice, a transportation network

Manuscript received September 4, 2019; revised April 20, 2020 and April 22, 2021; accepted June 15, 2021. This work was supported in part by Raytheon Company and in part by the Jet Propulsion Laboratory. The Associate Editor for this article was A. Nunez. (*Corresponding author: Benjamin Rivière.*)

The authors are with the Graduate Aerospace Laboratories, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: briviere@caltech.edu; sjchung@caltech.edu).

Digital Object Identifier 10.1109/TITS.2021.3097297



Fig. 1. Concept graphic of an intelligent transportation network. Autonomous taxis, that can include both ground and air vehicles, estimate in real-time the customer demand and coordinate locally to behave with bounded sub-optimality.

company might be composed of a dispatch center equipped with complete information and a large computational budget and a fleet of taxis, each operating with local information and a limited amount of processing power and communication bandwidth (see Fig. 1). In this manner, the transportation network company network can be decomposed into an underlying “star-topology” network between taxis and the dispatch center, and an arbitrary peer-to-peer network between taxis. The proposed algorithm, H-TD², exploits this topology explicitly by proposing a hybrid algorithm with two distinct behaviors: the central node computes exact, large-batch policies infrequently, and each taxi computes approximate, online updates with local information.

The overview of H-TD² is shown in Fig. 2. At a given timestep, the closest taxis service the new customer requests, and the rest of the free taxis are dispatched to reduce expected waiting time of future requests. The free taxis coordinate with a distributed game theoretic scheme to optimize their policy estimate, where the policy is estimated as follows: the servicing taxis communicate the customer data to their neighboring taxis, where the expected reward (e.g. customer waiting time) is estimated with a distributed estimation algorithm. Then, all taxis use the estimated reward to update their policy estimate with temporal difference learning. If any of the taxis determine that their policy estimate error is larger than the user specified threshold, the taxi signals to the dispatch center for a centralized policy update.

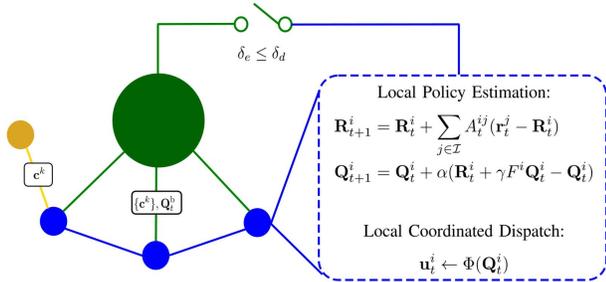


Fig. 2. Overview of H-TD², where blue represents the taxi network, yellow represents the customers, and green represents the dispatch center. The i th taxi estimates the dispatch policy with local operations: distributed estimation of reward, \mathbf{R}_t^i computed in (9), and temporal difference learning to update the policy, \mathbf{Q}_t^i , (13). If any of the taxis determines that its policy estimate error, δ_e , is larger than the user specified threshold, δ_d , the taxi signals to the dispatch center to receive a centralized policy update, \mathbf{Q}_t^p (6). Finally, each free taxi uses the policy in a game theoretic formulation, Φ (20), to find its dispatch position vector, \mathbf{u}_t^i .

The contributions of the paper are stated as follows:

- We derive a novel regret bound by leveraging distributed estimation methods in local online policy estimation and introduce a trigger condition to the batch update, permitting the user to explicitly specify the policy's computational and communication expense vs. bounded sub-optimality trade-off. This advances the state of the art by enabling distributed operation with bounded sub-optimality.
- We propose a taxi-dispatch solution that is adaptive, model-free, and coordinated. Unlike state-of-the-art reinforcement-learning dispatch methods, our method directly adapts the policy based on real-time data, thereby providing a property of robustness to irregular urban mobility events such as traffic, weather, and major public events. This advancement is enabled by two step approach: first we propose a hybrid policy estimation in a finite-dimensional, agent-agnostic cell abstraction, and then we interface the resulting policy estimation for agent-based coordination with a local prescriptive game-theoretic task assignment.

We demonstrate the performance and computational properties of our method with numerical experiments: our algorithm reduces customer waiting time compared to a receding horizon control baseline and the simulation runtime is linear with the number of agents. We also validate our claim that adaptive algorithms are robust to general irregular events with a case study of the Chicago City taxis during the 2016 Major League Baseball World Series.

The remainder of the paper is organized as follows: in Sec. II, we review the related literature and compare our method with the state of the art. In Sec. III, we present the taxi dispatch problem description and a motivating example. In Sec. IV, we present the cell-based Markov Decision Process (MDP). In Sec. V, we discuss the exact and approximate solutions to the MDP and the integration of the learned policy into a game theoretic method. In Sec. VI, we present numerical experiments demonstrating the advantages of our algorithm compared to a receding horizon control baseline in simulated

and real customer datasets. The details of the fleet simulator implementation are presented in Appendix.

II. RELATED WORK

Recent urban mobility research has developed dynamic and scalable methods. A well-studied example is the vehicle routing and dial-a-ride problems [1], [2] where taxis find a minimum cost path through a routing graph, and its dynamic extension, where part or all of the customer information is unknown and revealed dynamically. Recent dynamic routing research proposes scalable solutions to dynamic routing problems with bio-inspired methods [3], data-driven methods [4]–[6], and model-based methods [7], [8]. In this paper, we study a variant of the dynamic routing problem, taxi dispatch, where we propose a novel two-stage approach: distributed estimation with temporal difference learning, and game-theoretic coordination. This advances the state of the art by permitting adaptive distributed operation with bounded sub-optimality with respect to the optimal centralized policy.

Taxi dispatch is an emerging urban mobility problem where free taxis are dispatched to locations in the map to minimize customer waiting time of future requests. Recent approaches have adopted model-based [9], [10] and model-free [11] methods. An online model-based method like [9] uses real-time data to fit a system prediction model (for example, customer demand and taxi supply) and then compute a receding horizon control solution in response to that model. Pre-specified system models can be over-restrictive, and recent reinforcement-learning model-free methods [11]–[13] have been used successfully to overcome this limitation. In a model-free method, events are not explicitly modelled, they are captured by the arbitrary dynamics of the underlying reward. In regular operation, this reward is periodic, and can be accurately predicted (either explicitly or implicitly) and used for fleet control. However, it is possible that an irregular event occurs out of training domain and causes the reward dynamics to be unpredictable. In this case, we argue that it is better to adapt in real-time than predict with irrelevant data. Our model-free approach adapts the policy directly in response to real-time data, achieving performance that is robust to unpredictable, irregular events such as weather, accidents, and major public events.

Our method leverages results from reinforcement learning in convergence of temporal difference iteration [14]–[16] in a dynamic environment, i.e. when the reward or transition probabilities are changing over time. An alternative online model-free approach is online actor-critic [17], where our work differs from this result in two ways: we consider a general non-quadratic reward function and we consider a multi-agent setting by analyzing a hierarchical system of a temporal difference iteration with distributed estimation of the reward model. To the best of the authors knowledge, the only other work to propose a distributed temporal difference algorithm is recent work [18] that addresses the convergence properties of consensus on model parameters in the case of linear function approximations.

In general, multi-agent reinforcement learning research is challenging because the MDP's state and action space

dimensionality is coupled to the number of agents, which is typically handled by using either (i) function approximation methods such as deep neural networks or (ii) decoupled, decentralized solutions. A survey paper on multi-agent reinforcement learning discusses additional methods [19]. In contrast to an agent-based (or Lagrangian) approach, our method uses a naturally scalable cell-based (or Eulerian) model that decouples the problem dimensionality from the number of agents, inspired by a method used in probabilistic swarm guidance [20].

Because of the cell-based abstraction, our algorithm requires an additional task assignment component to coordinate taxis. Task assignment is a canonical operations research problem and there exists many available centralized [21]–[24] and decentralized [25]–[27]. Among these options, we use a distributed prescriptive game theory [28] approach that leverages existing asymptotic game theoretic optimality and convergence results. In contrast to conventional *descriptive* game theory, *prescriptive* game theory designs multi-agent local interactions to achieve desirable global behavior. Using one such method, binary log-linear learning [28], the taxis achieve global cooperative behavior with only local information.

III. PROBLEM DESCRIPTION

A. Notation

We denote vectors with a bold symbol, matrices with plain uppercase, scalars parameters with plain lowercase, functions with italics, and we use calligraphic symbols for operators and sets. We denote a taxi index with an i or j superscript, a customer index with a k superscript, and the time index with a subscript t . Also, I_n denotes the n -dimensional identity matrix.

B. Problem Statement

We consider the urban taxi dispatch problem, where we control a fleet of taxis to minimize customer waiting time. At each timestep, each customer requests is serviced by the nearest taxi. These *servicing* taxis use customer information to update their reward model and exchange information with neighboring taxis. The remaining *free* taxis are dispatched to locations in the map according to the proposed dispatch algorithm. The overall fleet control is summarized in Algorithm 1 and implementation details are given in Appendix.

Algorithm 1 Fleet Control Problem

```

1 initialize taxi fleet;
2 for  $t \in [t_0 : t_f]$  do
3   broadcast local customer requests to taxis;
4   assign closest free taxis to service customers;
5   dispatch free taxis to locations in the map;
6 end

```

The system, as shown in Fig. 3, is composed of customers and taxis. The k^{th} -customer state, \mathbf{c}^k , is composed of the time of request, trip duration, pickup location, and dropoff location, i.e. $\mathbf{c}^k = [t_r^k, t_d^k, \mathbf{p}^{k,p}, \mathbf{p}^{k,d}]$ and its pickup location is shown in

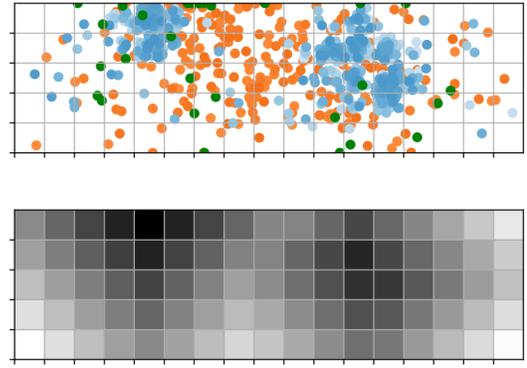


Fig. 3. State space representation of a Gridworld simulation with 1000 taxis, with corresponding value function estimation. In the top subplot, the blue dots are free taxi positions, the orange dots are positions of taxis currently servicing customers, and the green dots are the new customers requests pickup positions. In the bottom subplot, the approximate value function distribution is shown over the state space.

green in the top subplot. The i^{th} -taxi is defined by a position vector, \mathbf{p}_t^i and an operation mode: free (shown in blue) or servicing (shown in orange). The dispatch solution is a desired position vector for each of the free taxis, \mathbf{u}_t^i , that results in minimizing customer waiting time over a time horizon. Our method estimates the optimal policy, visualized with the value function over the state-space in the bottom subplot, that maximizes the expected reward over time.

IV. CELL-BASED MARKOV DECISION PROCESS

The previous section described the dispatch problem with an agent-based perspective, i.e. in terms of positions and actions of individual taxis and customers. Next, we will introduce the cell-based Markov Decision Process (MDP), where cell-based refers to an Eulerian perspective in which we analyze values like location and reward with respect to cells of a discretized map as shown in Fig. 3. The cell-based formulation decouples the decision making problem from the number of agents, permitting a finite dimensional policy representation. The decision making problem is formalized with a MDP, \mathcal{M} , defined as a tuple of state space, action space, transition model, reward model, and discount factor [29]:

$$\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle. \quad (1)$$

- The state space, \mathcal{S} , is defined as the set of map cells shown in Fig. 3, where the state of the i^{th} -taxi, s_t^i is the cell index that contains that taxi's position. The number of cells in the environment is denoted by the cardinality of the set, $|\mathcal{S}|$.
- The action space, \mathcal{A} , is defined as a movement between map cells for taxi i . The taxi on dispatch has 5 actions: $a_t^i \in \mathcal{A} = \{\text{stay, right, up, left, down}\}$, defined with respect to its current map cell in \mathcal{S} .
- The transition function, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{P}$ is defined as $P(s_t^i, a_t^i, s_{t+1}^i) = \mathbb{P}(s_{t+1}^i | s_t^i, a_t^i)$. We define P with a deterministic, cell-based dynamical model, f :

$$s_{t+1}^i = f(s_t^i, a_t^i) \quad \text{where } P(s_t^i, a_t^i, s_{t+1}^i) = 1. \quad (2)$$

If the next state is valid, the cell-based dynamical model moves the taxi from the initial state, s_t^i to the neighboring

state s_{t+1}^i according to its action. If the next state is not valid, the dynamical model returns the initial state.

- The reward function $R_t : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is defined to be the negative of the expected customer waiting time and is estimated from reward samples; reward is the negative of the time it takes the taxi to go from its current position to the dispatch position defined by the cell-action, and then from that position to the customer request. The reward has a subscript t because we assume the reward changes over time due to the changing customer distribution. We equivalently write the R_t function as a vector of state-action pairs, $R_t(s, a) = \mathbf{R}_t[s|\mathcal{A}| + i(a)]$ and $\mathbf{R}_t \in \mathbb{R}^{n_q}$, where $n_q = |\mathcal{S}||\mathcal{A}|$ and $i(a)$ denotes the action's index. Given a customer request, \mathbf{c}^k , we compute a sample of the reward function, r_t^i that can be used to estimate the underlying reward, \mathbf{R}_t according to an observation model:

$$r_t^i(s_t^i, a_t^i) = -(\eta(\mathbf{p}_t^i, \mathbf{u}_t^i) + \eta(\mathbf{u}_t^i, \mathbf{p}^{k,p})) \quad (3)$$

$$\mathbf{r}_t^i = H_t^i \mathbf{R}_t + \mathbf{v}_t^i \quad (4)$$

Recall that \mathbf{p}_t^i is the position of the i^{th} -taxi, \mathbf{u}_t^i is the dispatch desired position, and $\mathbf{p}^{k,p}$ is the k^{th} -customer pickup position. Also, η is the estimated time-of-arrival function that accepts position vectors, returns a scalar time value, and is specified in Appendix. It is parameterized by the average taxi velocity \bar{v}_{taxi} . The measurement noise is sampled from a normal distribution with variance ς , $\mathbf{v}_t^i \sim \mathcal{N}(0, \varsigma I)$. Finally, the cell-based observation model, $H_t^i \in \mathbb{R}^{n_q \times n_q}$, is a binary diagonal matrix with unity elements at the state-actions pairs where the customer request \mathbf{c}^k contains information of the corresponding state-action pair, and 0 otherwise.

- Note that γ is the discount rate of the system. This parameter determines the trade-off between greedy and long-term optimal behavior.

V. ALGORITHM DESCRIPTION AND ANALYSIS: H-TD²

We describe the details of H-TD² in this section, defining the exact and approximate policy estimation, the hybrid switching behavior, and the game theoretic task assignment. The overview of the method is given in Algorithm 2.

A. Centralized Q-Value Computation

We present the idealized Bellman solution to the cell-based decision making problem specified in (1). The solution is a policy function that maps states to an action that maximizes the discounted reward over time and can be represented as a value function, as shown in Fig. 3, or an action-value function known as Q_t^b -values. We use the latter and use the superscript b notation to denote the policy that is synthesized with a Bellman iteration method. We adopt the conventional optimal Q_t^b -value function as follows:

$$Q_t^b(s_t, a_t) = \mathbb{E}_{s \sim P(s_t|s_t, a_t)} [R_t + \gamma \mathbb{E}_{a_t' \sim \pi^b} Q^b(s_t', a_t')] \quad (5)$$

We specify this general formulation with some assumptions. First, we use a finite-dimensional tabular Q_t^b and write the Q_t^b function as a vector of state-action pairs, $\mathbf{Q}_t^b \in \mathbb{R}^{n_q}$,

Algorithm 2 H-TD² at Timestep t

```

1 input: set of total, free, and servicing taxis:  $\mathcal{I}, \mathcal{I}_f, \mathcal{I}_s$ 
2 output: action profile for free taxis,  $\mathbf{u}_t$ 
   /* Hybrid Temporal Difference */
3 for  $\forall i \in \mathcal{I}$  do
4   if  $\delta_e > \delta_d$  (15) then
5     slow update  $\mathbf{Q}_t^i$  with aggregated global
     information (6) at the central node;
6   else
7     fast update  $\mathbf{Q}_t^i$  with local information (13) at each
     taxi;
8   end
9 end
   /* Game Theoretic Task Assignment */
10 randomly initialize cell-based action profile  $\hat{\mathbf{A}}_t$ ;
11 while  $\hat{\mathbf{A}}_t$  not converged do
12   randomly pick  $i \in \mathcal{I}_f$  that has not converged;
13   consider current action,  $a_t^i$ ;
14   propose random action  $a_t^i$ ;
15   compute marginal utility,  $J$  with  $\mathbf{Q}_t^i$  (20);
16   stochastically assign action with  $J$  (21);
17   check  $i^{\text{th}}$ -taxi action convergence;
18 end
19 Convert cell-based actions  $a_t^i$  to position vectors  $\mathbf{u}_t^i$ ;

```

as done with the reward in Sec. IV. Next, we apply the deterministic transition function, $P(s_t^i, a_t^i, s_{t+1}^i)$, as specified in (2) to remove the outer expectation. We remove the inner expectation by specifying the policy π^b to be a transition kernel matrix F^b such that $F^b \mathbf{Q}_t^b(s_t^i) = \max_{a_t^i} Q_t^b(s_t^i, a_t^i)$. Combining these, we rewrite a simplified expression for \mathbf{Q}_t^b as the fixed point of the Bellman operator \mathcal{T} :

$$\mathbf{Q}_t^b = \mathbf{R}_t + \gamma F^b \mathbf{Q}_t^b = \mathcal{T} \mathbf{Q}_t^b \quad (6)$$

The Bellman iteration can be solved using conventional value iteration or policy iteration methods from batch customer data. For the purpose of this paper, we use a Modified Policy Iteration (MPI) method [30] to solve line 5 of Algorithm 2. However, there are complications with implementing a pure Bellman approach, which we address in the next subsection.

B. Distributed Reward Estimation and Q-Value Iteration

We present a policy approximation that overcomes the Bellman solution's practical limitations. The first issue with a pure Bellman solution is that in an online setting, the reward information is hidden a priori and received incrementally in samples, \mathbf{r}_t^i . So the expectation of the reward is not immediately available. Furthermore, each taxi only has access to local information. To overcome these problems, we assume the hidden reward evolves as a random walk process and synthesize linear estimators for the hidden reward \mathbf{R}_t :

$$\mathbf{R}_{t+1} = \mathbf{R}_t + \mathbf{w}_t \quad (7)$$

$$\mathbf{R}_{t+1}^c = \mathbf{R}_t^c + \sum_{j \in \mathcal{I}} K_t^j (\mathbf{r}_t^j - H_t^j \mathbf{R}_t^c) \quad (8)$$

$$\mathbf{R}_{t+1}^i = \mathbf{R}_t^i + \sum_{j \in \mathcal{I}} A_t^{ij} (\mathbf{r}_t^j - \mathbf{R}_t^i) \quad (9)$$

Recall \mathbf{r}_t^i is the reward sample defined in (4). Also $\mathbf{R}_t^c, \mathbf{R}_t^i \in \mathbb{R}^{n_q}$ are centralized and distributed reward estimators that will be used in the upcoming temporal difference learning, where the superscript c denotes a centralized quantity. The H_t^i matrix is the i^{th} -taxi's measurement model defined in (4) and the K_t^i matrix are the corresponding estimator gains. The process noise, $\mathbf{w}_t \sim \mathcal{N}(0, \varepsilon I)$ is sampled from a normal distribution where the parameter ε is computed offline from training data. The row-stochastic adjacency block matrix, A_t , specifies the local information available to each agent and is defined as:

$$A_t^{ij} = B_t^{ij} / \sum_{j=1}^{n_i} B_t^{ij}, \quad \text{where } B_t^{ij} = \begin{cases} K_t^j H_t^j & j \in \mathcal{I}_t^i \\ \varnothing_{n_q \times n_q} & \text{else} \end{cases} \quad (10)$$

$$\mathcal{I}_t^i = \{j \in \mathcal{I} \mid \|\mathbf{p}_t^j - \mathbf{p}_t^i\| < R_{\text{comm}}\}. \quad (11)$$

with diagonal block matrices $A_t^{ij}, B_t^{ij} \in \mathbb{R}^{n_q \times n_q}$ and full matrices $A_t, B_t \in \mathbb{R}^{n_{iq} \times n_{iq}}$. The i^{th} agent constructs the B_t^{ij} matrices from the gain and measurement matrices shared by its neighbors $j \in \mathcal{I}_t^i$. The local observation is parameterized by the radius of communication, R_{comm} .

The second issue with the Bellman approach is an intrinsic drawback that at higher state/action dimensions the Bellman-iteration calculation becomes computationally-expensive and cannot be quickly evaluated online. Instead, temporal difference learning [31] can be used as an approximate method to estimate \mathbf{Q} -values online using \mathbf{R}_t^c :

$$\mathbf{Q}_{t+1}^c = \mathbf{Q}_t^c + \alpha (\mathbf{R}_t^c + \gamma F^c \mathbf{Q}_t^c - \mathbf{Q}_t^c) \quad (12)$$

where $\alpha \in \mathbb{R}$ is the system learning rate and F^c is the transition kernel for this policy.

This formulation requires a central node to collect the data, compute a policy, and broadcast the new information at every timestep, scaling the computation complexity, bandwidth, and network delay with the number of taxis. To address this limitation, we introduce a distributed algorithm using communication between the taxis, and propose a policy update computed at each taxi using only local information:

$$\mathbf{Q}_{t+1}^i = \mathbf{Q}_t^i + \alpha (\mathbf{R}_t^i + \gamma F^i \mathbf{Q}_t^i - \mathbf{Q}_t^i) \quad (13)$$

where this temporal difference (13), with the distributed reward estimation (9) defines line 7 of Algorithm 2.

Using the approximate temporal difference method and estimating with only local information hurts the quality of the final policy used by each taxi. We derive the upper bound of the negative effect of these approximations through a regret bound analysis, comparing the policy synthesized with the proposed algorithm (13) and the Bellman-optimal solution (6).

Theorem 1: The expected distance between an arbitrary taxi \mathbf{Q} -value estimates, \mathbf{Q}_t^i , and the Bellman-optimal solution \mathbf{Q}_t^b is upper bounded by:

$$\mathbb{E} \|\mathbf{Q}_t^i - \mathbf{Q}_t^b\|_2 \leq \frac{2\sqrt{n_q(\varepsilon + \varsigma)}}{(1 - \gamma)(1 - \sqrt{1 - \lambda_{\min}(\sum_{j \in \mathcal{I}} A_t^{ij})})} \quad (14)$$

where $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue of a matrix.

Proof: First, we write the distributed iteration as an application of the Bellman operator on the previous timestep with a disturbance. Then we solve for the disturbance to derive the final bound.

Step 1: Consider (13) and add and subtract $\alpha \mathbf{R}_t$:

$$\begin{aligned} \mathbf{Q}_{t+1}^i &= \mathbf{Q}_t^i + \alpha (\mathbf{R}_t^i + \gamma F^i \mathbf{Q}_t^i - \mathbf{Q}_t^i) + \alpha \mathbf{R}_t - \alpha \mathbf{R}_t \\ &= \mathbf{Q}_t^i + \alpha (\mathbf{R}_t + \gamma F^i \mathbf{Q}_t^i - \mathbf{Q}_t^i) + \alpha (\mathbf{R}_t^i - \mathbf{R}_t) \\ &= \mathbf{Q}_t^i + \alpha (\mathcal{T} \mathbf{Q}_t^i - \mathbf{Q}_t^i) + \alpha \mathbf{e}_t^i \end{aligned}$$

where $\mathbf{e}_t^i = \mathbf{R}_t^i - \mathbf{R}_t$. The system is contracting at rate $1 - \alpha(1 - \gamma)$, implying the system geometrically converges to an equilibrium about $\mathcal{T} \mathbf{Q}_t^i = \mathbf{Q}_t^i$. In addition, from Banach's fixed point theorem [32], \mathcal{T} contracts to a unique fixed point, \mathbf{Q}_t^b . Applying Discrete Gronwall's lemma [32]:

$$\|\mathbf{Q}_t^i - \mathbf{Q}_t^b\| \leq \frac{\|\mathbf{e}_t^i\|}{1 - \gamma}$$

Note that the decaying initial condition term does not appear because we assume that the policy estimate is initialized with the Bellman solution, i.e. $\mathbf{Q}_0^i = \mathbf{Q}_0^b$.

Step 2: Here we need to bound the value $\|\mathbf{e}_t^i\|$, i.e. the error between the estimated reward and the true reward. We write the dynamics of the error vector \mathbf{e}_t^i by subtracting (7) from (9):

$$\mathbf{e}_{t+1}^i = (I - \sum_{j \in \mathcal{I}} A_t^{ij}) \mathbf{e}_t^i + \mathbf{d}_t$$

where $\mathbf{d}_t = \mathbf{w}_t + \sum_{j \in \mathcal{I}} A_t^{ij} \mathbf{v}_t^j$. By applying Weyl's interlacing eigenvalue theorem [33], we prove that the i^{th} system is contracting at rate lower bounded by $\lambda_t^i = 1 - \lambda_{\min}(\sum_{j \in \mathcal{I}} A_t^{ij})$.

We rewrite the disturbance as the product of an input matrix, M , and the stacked noise, $\mathbf{z}_t \sim \mathcal{N}(0, W)$:

$$\mathbf{d}_t = M_t \mathbf{z}_t$$

where $\mathbf{z}_t = [\mathbf{w}_t; \mathbf{v}_t^1; \dots; \mathbf{v}_t^{n_i}]$, $M_t = [I_{n_q}, A_t^{i1}, \dots, A_t^{in_i}]$ and $W = \text{blkdiag}(\varepsilon I_{n_q}, \varsigma I_{n_q}, \dots, \varsigma I_{n_q})$.

By application of the convergence theorem of discrete stochastic contracting systems [34], [35], the expected error of a single agent is upper bounded by:

$$\begin{aligned} \mathbb{E} \|\mathbf{e}_t^i\| &\leq \frac{2\sqrt{C}}{1 - \sqrt{1 - \lambda_{\min}(\sum_{j \in \mathcal{I}} A_t^{ij})}} \\ C &= \text{trace}(M_t^T M_t W) \end{aligned}$$

It remains to calculate the value C :

$$C = \varepsilon \text{trace}(I_{n_q}) + \varsigma \sum_{j \in \mathcal{I}} \text{trace} \left((A_t^{ij})^T A_t^{ij} \right) \leq n_q (\varepsilon + \varsigma)$$

where we use the linearity of the trace operation to move it outside of the sum, then we use the non-negativity and row-stochasticity of A_t to bound $\sum_{j \in \mathcal{I}} (A_t^{ij})^T A_t^{ij} \leq I_{n_q}$. The final result is found by plugging the result from Step 2 into the result from Step 1. ■

Remark 1: The regret bound is driven by the contraction rate of the system λ_t^i , a combined graph and observability quantity. Intuitively, this corresponds to a non-zero value when taxi i and its neighbors taxis j can measure the entire

state-action vector. We can also consider a batch measurement over a time interval, n_T and an average contraction rate, λ_t^i . This time interval approach exists in the multi-agent adaptive control literature, where $\lambda_t^i > 0$ is analogous to an excitation level in the Collective Persistency of Excitation condition [36].

Remark 2: The proposed online method is used to estimate the optimal policy in a dynamic environment, i.e. the reward model, \mathbf{R}_t is time-dependent. In this case, the optimal policy is non-stationary, i.e. $\mathbf{Q}_{t+1}^b \neq \mathbf{Q}_t^b$. In order to guarantee the convergence of the TD-algorithm, we require that there exists a timescale separation between the convergence of the TD-algorithm and the dynamics of \mathbf{Q}^b :

$$\|\mathbf{Q}_{t+1}^b - \mathbf{Q}_t^b\| \ll (1 - \gamma)(1 - \sqrt{1 - \lambda_{\min}(\sum A_t^{ij})})$$

Remark 3: The adjacency matrix A_t dictates that each agent takes a convex combination of the neighboring measurements. Further, the estimation gain matrices, K_t^i are chosen with a Distributed Kalman Information Filter, whose proof of optimality with respect to mean-squared-error can be found in [37], [38]. Thus, the agents weigh the neighboring measurements appropriately.

C. Hybrid Temporal Difference Algorithm

We define the switching condition in line 4 of Algorithm 2 with two parameters, δ_e , the estimated error in the system, and δ_d , the user specified desired error in the system.

Proposition 1: If we define:

$$\delta_e = \frac{2\sqrt{n_q(\varepsilon + \varsigma)}}{(1 - \gamma)(1 - \sqrt{1 - \lambda_{\min}(\sum_{j \in \mathcal{I}} A_t^{ij})})} \quad (15)$$

the expected policy sub-optimality will be bounded by δ_d .

Nominally, the system evolves with the distributed temporal difference method (13), computing δ_e at each timestep. Each agent is able to compute this value because ε , ς , and γ are known system parameters and each agent keeps track of its own $\lambda_{\min}(\sum_{j \in \mathcal{I}} A_t^{ij})$ values. Applying the result from Theorem 1, the expected policy suboptimality is identically δ_e , so, if δ_e exceeds the desired error, δ_d , the desired sub-optimality is violated. However, if this condition occurs at time t , the system resets all taxis with a central policy update (6), i.e. $\mathbf{Q}_t^i = \mathbf{Q}_t^b$, $\forall i \in \mathcal{I}$. Therefore, the H-TD² algorithm maintains the distance between the estimated policy and a true regret policy to user specification. In effect, δ_d controls the trade-off of computational expense to policy sub-optimality, where $\delta_d \rightarrow 0$ produces a solution with no regret but maximum computational effort and $\delta_d \rightarrow \infty$ produces a solution with potentially infinite regret with little computational effort.

D. Game Theoretic Task Assignment

In this section, we propose a game-theoretic task assignment to coordinate the taxis according to the \mathbf{Q}_t^i -values estimated in Sec. V-B. The \mathbf{Q}_t^i policy does not account for the actions of the other taxis and, without additional coordination, the taxis would behave greedily by all going to the highest value cell, increasing the overall customer waiting time. To avoid this

behavior, we design a potential game and a local action profile iteration to maximize each agent's marginal utility. This is implemented in Algorithm 2, Lines 10-18.

First, we introduce a global action profile, \mathbb{A}_t and a local action profile for the i^{th} taxi, \mathbb{A}_t^i :

$$\mathbb{A}_t = \{a_t^j \mid \forall j \in \mathcal{I}\}, \quad \text{and} \quad \mathbb{A}_t^i = \{a_t^j \mid \forall j \in \mathcal{I}_t^i\} \quad (16)$$

where the j^{th} -neighbor taxi communicates its action, a_t^j , to the i^{th} -taxi, where a_t^j is defined in (1).

Next, we introduce the current global fleet distribution Ω_t , i.e. the number of taxis in each cell, as a function of the action profile:

$$\Omega_t(s, \mathbb{A}_t) = \frac{1}{n_i} \sum_{j \in \mathcal{I}} \mathbb{I}(s = f(s_j^j, a_t^j)) \quad (17)$$

where \mathbb{I} denotes the indicator function and f is the dynamics model specified in (1). The local fleet distribution, Ω_t^i , is found with the same calculation but summing only over the neighboring agents, $j \in \mathcal{I}_t^i$. By defining the radius of communication as $R_{\text{comm}} = 3ds$ where ds is the length of a cell in the environment, we guarantee that the i^{th} taxi can always calculate the fleet distribution in neighboring cells, \mathcal{S}^i , within one action of the current cell of the i^{th} taxi.

Next, we describe the desired fleet distribution using the \mathbf{Q} -values as computed in (9). Consider the following Boltzmann exploration strategy [31] strategy to synthesize a desired distribution, Ω^* , from the \mathbf{Q} -value estimates:

$$\Omega^*(s, Q_t^i) = \frac{\exp(\beta \max_{a \in \mathcal{A}} Q_t^i(s, a))}{\sum_{a' \in \mathcal{A}} \exp(\beta Q_t^i(s, a'))} \quad (18)$$

Recall that \mathcal{A} are the local actions available to each taxi defined in (1) and $\beta \in \mathbb{R}$ is an exploration/exploitation design constant. For simplicity of notation, we have written the action-values in its functional form, Q_t^i .

The goal of the game theoretic task assignment is to find an action profile, \mathbb{A}_t through local iteration methods that minimizes the distribution distance between the current distribution, Ω_t and the desired distribution Ω^* . We describe a potential and noncooperative game meaning that the taxis will try to converge to a Nash equilibrium with a high potential function value. The global and marginal potential functions, Φ and J are defined as follows:

$$\Phi(\mathbb{A}_t) = - \sum_{s \in \mathcal{S}} (\Omega^*(s, Q_t^i) - \Omega(s, \mathbb{A}_t))^2 \quad (19)$$

$$J(\mathbb{A}_t^i) = - \sum_{s \in \mathcal{S}^i} (\Omega^*(s, Q_t^i) - \Omega(s, \mathbb{A}_t^i))^2 \quad (20)$$

For the calculation of J , we only require the indices that correspond to neighboring cells of the current cell of the i^{th} -taxi, thereby permitting a local calculation.

Remark 4: The game's utility function, $\Phi(\mathbb{A}_t)$ has an analogy to sample-based planners if each taxi in the fleet is considered as a sampled action of a stochastic policy, $\Omega^*(s, \mathbf{Q}_t^i)$. This choice of utility function is interesting because it could be the utility function chosen by a centralized algorithm but we can maximize it with local calculations through J .

Remark 5: Note that J is indeed the marginal contribution on the global potential function:

$$\Phi(\mathbb{A}'_t) - \Phi(\mathbb{A}_t) = J(\mathbb{A}'_t) - J(\mathbb{A}_t)$$

where $\mathbb{A}'_t = \{a_t^j \mid \forall j \in \mathcal{I}/\{i\}\} \cup \{a_t^{i'}\}$ is the global alternative action set.

We use a game-theoretic reinforcement learning technique, binary log-linear learning [28] to iterate to an action set \mathbb{A}_t , shown in line 10-19 of Algorithm 2. At each timestep, t , the action set, \mathbb{A}_t is randomly initialized. While all other taxi's actions are held, a randomly selected i^{th} -taxi chooses between the previously held action, a_t^i , and an alternate action $a_t^{i'}$ with probability $p_t^i(\mathbb{A}_t^i, \mathbb{A}_t^{i'})$:

$$p_t^i(\mathbb{A}_t^i, \mathbb{A}_t^{i'}) = \frac{\exp(J(\mathbb{A}_t^i)/\tau)}{\exp(J(\mathbb{A}_t^i)/\tau) + \exp(J(\mathbb{A}_t^{i'})/\tau)} \quad (21)$$

where $\mathbb{A}_t^{i'} = \{a_t^j \mid \forall j \in \mathcal{I}_t/\{i\}\} \cup \{a_t^{i'}\}$ is the local alternative action set. The coefficient $\tau \in \mathbb{R}_{>0}$ is a design parameter specifying how likely taxi i chooses a sub-optimal action, to specify the trade-off between exploration and exploitation. The action set is chosen once the iteration has converged, completing the game-theoretic task assignment. Then, the cell-based action a_t^i is converted to a dispatch vector \mathbf{u}_t^i , where \mathbf{u}_t^i is a randomly sampled position vector in the cell after the dispatch action is taken.

VI. NUMERICAL EXPERIMENTS

A. Baseline and Variants

We compare our H-TD² solution with a receding horizon control (RHC) baseline dispatch algorithm, adapted from the baseline in [11]. For this section, we use independent notation from the rest of the paper, matching the notation in [11]. The RHC dispatch algorithm is formulated as the following linear program:

$$\begin{aligned} \mathbf{u}^* &= \operatorname{argmax} \sum_{t=t_0}^{t_0+t_{\text{rhc}}} \gamma^{t-t_0} \sum_{i=1}^M \min(\bar{w}_{t,i} - \mathbf{x}_{t,i}, 0) \\ \text{s.t.} \quad &\sum_{j=1}^M u_{ij,t} = \mathbf{x}_{t+1,i}, \quad \sum_{i=1}^M u_{ij,t} = \mathbf{x}_{t,i} \\ &u_{ij,t} = 0 \quad \forall j \notin S^i, \quad \mathbf{x}_{t_0,i} = X_{0,i} \end{aligned} \quad (22)$$

where the state variable, $\mathbf{x}_t \in \mathbb{Z}^{|\mathcal{S}|}$, is the number of free taxis in each cell, the control variable $u_{ij,t}$ is the number of taxis moving from cell i into cell j at time t and t_{rhc} is the RHC planning horizon. The first two constraints are conservation constraints: (i) the number of taxis in cell i is the number of taxis moving into cell i , and (ii) the number of taxis moving from cell i is equal to the number of taxis previously in cell i . The third constraint is that each taxi can only move to a neighboring cell. The fourth constraint is the initial condition. For large fleets, a proper assumption from [11] is to relax $u_{ij,t}$ from integers to real numbers, resulting in a linear program in $u_{ij,t}$. The reward is the difference of customer demand and taxi supply, where $\bar{w}_{t,i}$ is the expected customer demand and is computed from offline training data with a calculation

proposed in [9], [11]. The baseline is chosen to demonstrate the advantage of fast adaption over learned prediction: static prediction methods (either explicit in model-based or implicit in model-free) are vulnerable to events occurring outside of the training domain.

We study the effect of each component of the policy estimation algorithm by comparing variants of the policy: *Centralized Temporal Difference* (C-TD), *Distributed Temporal Difference* (D-TD), and *Bellman-optimal*. All of these variants control the fleet with binary log-linear learning (21), but they differ in how \mathbf{Q} is synthesized: C-TD uses (8) and (12), D-TD uses (9) and (13), and the Bellman-optimal policy is synthesized with (6). Equivalently, the D-TD and Bellman-optimal algorithms can be interpreted as the limiting behavior of H-TD² corresponding to the respective cases where $\delta_d = \infty$ and $\delta_d = 0$.

B. Customer Demand Datasets

We consider two datasets of customer requests: a synthetic dataset for a Gridworld environment and the real customer taxi dataset from the city of Chicago [39]. Recall each customer request is defined as follows: $\mathbf{c}^k = [t_r^k, t_d^k, \mathbf{p}^{k,p}, \mathbf{p}^{k,d}]$.

The synthetic dataset is generated as follows: At each timestep, $t \in [t_0, t_0 + \Delta_t]$, the customer request model is sampled n_c times, where n_c is the number of customers per timestep of the simulation. The time of the request, t_r^k is uniformly randomly sampled in the timestep. The pickup location, $\mathbf{p}^{k,p}$, is found by sampling a 2-dimensional Gaussian Mixture Model (GMM), where translating Gaussian distributions capture the underlying dynamic customer demand. The dropoff location, $\mathbf{p}^{k,d}$ is a randomly sampled position in the map, and the duration of the trip, t_d^k is given by $\eta(\mathbf{p}^{k,p}, \mathbf{p}^{k,d})$. The GMM model is parameterized by the number of Gaussian distributions, n_G , the speed of the distributions, v_G , the variance, σ_G , and the initial position, and unit velocity vector of the centroid for each distribution.

The real customer dataset is taken from the city taxi dataset of Chicago [39] filtered by start and end timestamp. The Socrata API permits importing raw data in the \mathbf{c}^k format, where the location data is specified in longitude, latitude coordinates. For this experiment, we load a city map of Chicago as a shapefile and perform minor geometric processing with the Shapely Python toolbox.

C. Results

1) *Gridworld Simulations:* We present variants of our H-TD² algorithm against a RHC baseline in a Gridworld environment as shown in Fig. 3. This flexible, synthetic environment permits us to test on a range of system parameters.

First, we introduce the algorithms in a small-scale simulation. We synthesize a dataset with parameters: $n_c = 5$, $n_G = 2$, $v_G = 0.02625$, $\sigma_G = 0.014$, and randomly initialize the mean and direction of the distributions. Our simulation parameters are: $n_i = 100$, $|\mathcal{S}| = 85$, $\gamma = 0.9$, $\alpha = 0.75$, $n_T = 10$, $\zeta = 0.014$, $\varepsilon = 0.0187$, $\delta_d = 0.025\|\mathbf{Q}_0^b\|$, $\beta = 150$, $\bar{v}_{\text{taxi}} = 0.125$, $\tau = 0.0001$, and $t_{\text{RHC}} = 10$. We run this experiment for each of the algorithms for 5 trials. This experiment is shown in Fig. 3, modified with $n_i = 1000$.

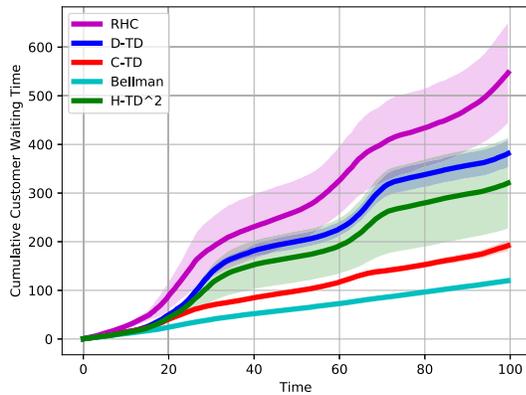


Fig. 4. Cumulative customer waiting time for different algorithms in the small-scale Gridworld environment. The algorithms behave as expected: in descending order of performance, Bellman, centralized temporal difference, H-TD², distributed temporal difference, followed by the receding horizon control baseline. The simulation is run 5 times, and the mean with standard deviations is visualized in the plot.

Next, we collect statistics on the cumulative reward of each algorithm, and plot the results in Fig. 4. The algorithms behave as expected: in descending order of performance, Bellman, centralized temporal difference, H-TD², distributed temporal difference, followed by the receding horizon control baseline. At the cost of computational effort, the user can tune the performance of H-TD² between the distributed temporal difference and Bellman-optimal solution by changing δ_d . The simulation is run 5 times, and the mean with standard deviations is visualized in the plot.

To explain the performance difference between the variants of our method, we show an error trace of the policy in Fig. 5. For a given policy \mathbf{Q} , we calculate the error with respect to the Bellman solution: $e_t^{\mathbf{Q}} = \|\mathbf{Q}_t^{\mathbf{Q}} - \mathbf{Q}_t^{\mathbf{b}}\| / \|\mathbf{Q}_t^{\mathbf{b}}\|$. For this experiment, we set the δ_d parameter is set to 2.5% of the norm of the Bellman solution, indicated by the dashed horizontal line. Initially, the H-TD² and distributed temporal difference (D-TD) algorithms behave identically, until the trigger condition is satisfied and the H-TD² requests a global Bellman-optimal update, thereby bringing its error to zero. As expected, the centralized-temporal difference method, C-TD, generally tends to estimate the \mathbf{Q} -values better than its distributed counterpart, D-TD.

Next, we test the proposed algorithm and baseline's scalability and performance across a wide range of taxi-densities and plot the results in Fig. 6. We fix the parameters from the small-scale simulation and only change the number of taxis and number of customers, where we maintain the ratio $n_i/n_c = 10$. In the top figure, the average reward is shown across a variety of taxi density regimes, where the H-TD² algorithm outperforms a RHC baseline by almost a factor of 2 in all taxi-density regimes. In the bottom figure, we show that the computational time is approximately linear with number of taxis (and taxi-density) across 3 orders of magnitude. The computational complexity of both RHC and H-TD² scales with the spatial resolution of the simulation, and in practice, we limit the maximum number of cells to 200.

2) *Chicago City Simulations*: We present the H-TD² against an RHC baseline using real customer data from the city of

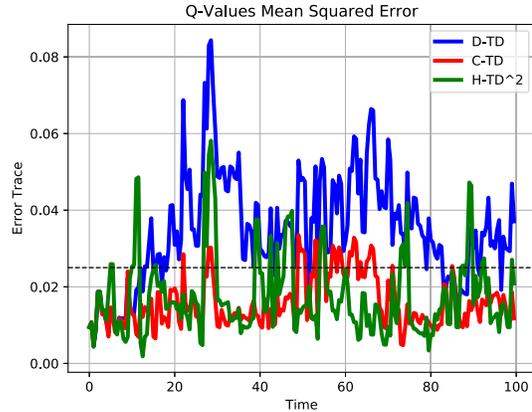


Fig. 5. \mathbf{Q} -value error trace for different algorithms with respect to the Bellman-optimal. Initially, the H-TD² and distributed temporal difference algorithms behave identically, until the trigger condition is satisfied and the H-TD² requests a global Bellman-optimal update, bringing the error to zero. The δ_d parameter is set to 2.5% of the norm of the Bellman solution and is shown with a dashed black horizontal line.

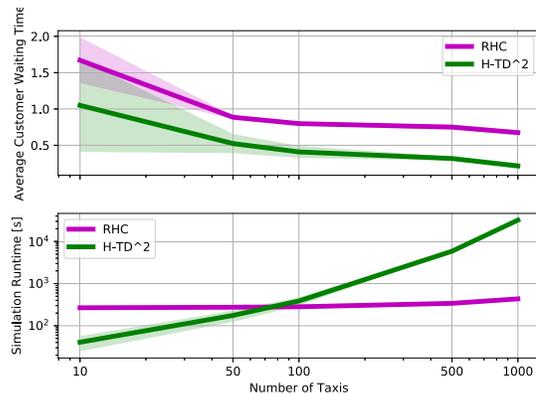


Fig. 6. Performance and scalability analysis of H-TD² and RHC against number of taxis. In the top subplot, the average reward is shown across a variety of taxi density regimes, and the proposed algorithm outperforms a receding horizon control baseline by at least 50% in all taxi-density regimes. In the bottom subplot, the computational time is approximately linear with number of taxis (and taxi-density) across 3 orders of magnitude.

Chicago public dataset [39], in a Chicago map environment. We show that our algorithm outperforms the baseline in practical datasets and demonstrate that online algorithms are robust in irregular urban mobility events.

In Fig. 7, we present the Chicago city taxi customer demand across an irregular event: Game 5 of the 2016 Baseball World Series. The map cells show the number of customer pickup requests, and the green star is Wrigley Field's (baseball stadium) location. Below, we plot the customer demand over time for the cell containing Wrigley Field. We show that a reward model trained using data from the day before would not accurately predict the behavior of the next day.

We evaluate the algorithms and plot the results in Fig. 8. Our simulation parameters are: $n_i = 2000$, $|\mathcal{S}| = 156$, $\gamma = 0.8$, $\alpha = 0.1$, $n_T = 10$, $\zeta = 0.0001$, $\varepsilon = 0.0001$, $\delta_d = 0.025\|\mathbf{Q}_0^{\mathbf{b}}\|$, $\beta = 1$, $\bar{v}_{\text{taxi}} = 22$ miles per hour, $\tau = 0.0001$, and $t_{\text{RHC}} = 10$. We train a reward model using the data from October 29th, 2016 with a total of 91, 165 customer requests. Then, we collect 54, 115 customer requests from October 30th, 2016, which we reveal real-time to the H-TD² and RHC

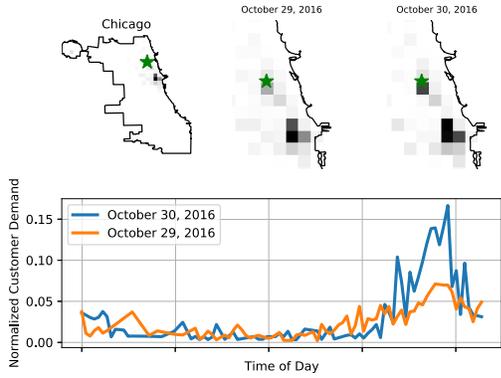


Fig. 7. Chicago city taxi customer demand across an irregular event: Game 5 of the 2016 Baseball World Series. The map cells show the number of customer pickup requests, and the green star is Wrigley Field's location. Below, we plot the customer demand over time for the cell containing Wrigley Field to show that a reward model trained using data from the day does not accurately predict the future behavior.

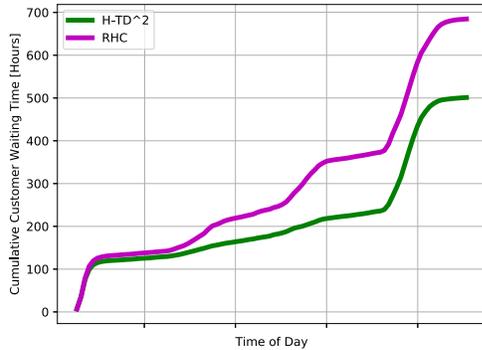


Fig. 8. Cumulative customer waiting time for the H-TD² and the RHC baseline in a Chicago city environment with a fleet of 2,000 taxis servicing 54,115 real customer requests during the 2016 Major League Baseball World Series. The H-TD² algorithm has a total customer waiting time of 504hours, an improvement of 26% over the RHC baseline.

dispatch algorithms. In total, the H-TD² algorithm has a total customer waiting time of 501hours an improvement of 26% over the RHC baseline of a cumulative customer waiting time of 684hours. This result demonstrates the robustness of adaptive algorithms to irregular events.

VII. CONCLUSION

In this work, we present H-TD², a novel approach to taxi dispatch that exploits the natural topology of transportation network systems with hybrid behavior: a fleet of taxis compute lightweight policy updates locally and, episodically, the policy of all taxis are reset with a centralized batch update from the dispatch center. We derive a regret bound on the local policy estimation that introduces a trigger condition to the batch update, permitting the user to explicitly specify the computational expense/policy bounded sub-optimality trade-off. Unlike state-of-the-art taxi dispatch methods, this algorithm is adaptive, model-free and coordinated, permitting the fleet to adapt in a flexible manner to regular and non-regular events. Compared to a receding horizon control baseline, the proposed algorithm decreases average customer waiting time by 50% on a synthetic dataset across a wide range of system parameters, and by 26% on real customer data

from the city of Chicago taxi dataset over irregular customer distributions during Game 5 of the 2016 Baseball World Series. In future work, we propose coupling H-TD² with a lower-level motion planner such as [40].

APPENDIX FLEET SIMULATOR

In this section we describe the implementation of the fleet simulator described in Algorithm 1. In particular, the fleet is initialized with uniform random initial positions in the state space and in the dispatch operating mode (Line 1). The customer requests are either synthetically generated in Gridworld or drawn from the Chicago city taxi dataset (Line 3). The taxis are assigned to customers with a distributed matching algorithm (Line 4). After the taxi's choose a dispatch action, the environment updates as follows:

The dynamical model of the i^{th} -taxi servicing customer \mathbf{c}^k is defined as:

$$\mathbf{p}_t^i = \begin{cases} \mathbf{p}_{t_r}^i + \frac{t - t_r^k}{t_p^{k,i} - t_r^k} (\mathbf{p}^{k,p} - \mathbf{p}_{t_r}^i) & t_r^k \leq t < t_p^{k,i} \\ \mathbf{p}^{k,p} + \frac{t - t_p^k}{t_d^k} (\mathbf{p}^{k,d} - \mathbf{p}^{k,p}) & t_p^{k,i} \leq t < t_p^{k,i} + t_d^k \end{cases} \quad (23)$$

where this model is evaluated for all servicing taxis, i.e. $\forall i \in \mathcal{I}_s$, $t_p^{k,i} = t_r^k + \eta(\mathbf{p}^{k,p}, \mathbf{p}_{t_r}^i)$ is the pickup time of customer k by taxi i , and η is the estimated time of arrival function for the environment.

The dynamic model of the i^{th} -taxi dispatched by dispatch action a_t^i is given by:

$$\mathbf{p}_{t+1}^i = \mathbf{p}_t^i + \frac{\Delta_t}{\eta(\mathbf{u}_t^i, \mathbf{p}_t^i)} (\mathbf{u}_t^i - \mathbf{p}_t^i) \quad (24)$$

where this model is evaluated for all free agents, i.e. $\forall i \in \mathcal{I}_f$, Δ_t is the simulation timestep, and \mathbf{u}_t^i is the desired position vector received from the dispatch algorithm.

The *Estimated Time of Arrival (ETA)*, η accepts arbitrary position vectors, returns a scalar time value and is defined as:

$$\eta(\mathbf{p}^1, \mathbf{p}^2) = \frac{\|\mathbf{p}^1 - \mathbf{p}^2\|}{\bar{v}_{\text{taxi}}} \quad (25)$$

where \bar{v}_{taxi} is the average taxi speed, computed offline as an input parameter.

Remark 6: Our proposed dispatch algorithm is agnostic to choice of models f_f, f_s, η so these can be augmented with a road network models or a data-driven approach.

ACKNOWLEDGMENT

The authors would like to thank Salar Rahili, who proposed the use of binary log-linear learning in an early version of this work.

REFERENCES

- [1] S. C. Ho, W. Szeto, Y.-H. Kuo, J. M. Y. Leung, M. Petering, and T. W. Tou, "A survey of dial-a-ride problems: Literature review and recent developments," *Transp. Res. B, Methodol.*, vol. 111, pp. 395–421, May 2018.
- [2] V. Pillac, M. Gendreau, C. Gu  ret, and A. L. Medaglia, "A review of dynamic vehicle routing problems," *Eur. J. Oper. Res.*, vol. 225, no. 1, pp. 1–11, 2013.

- [3] R. Claes, T. Holvoet, and D. Weyns, "A decentralized approach for anticipatory vehicle routing using delegate multiagent systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 12, no. 2, pp. 364–373, Feb. 2011.
- [4] K. Menda *et al.*, "Deep reinforcement learning for event-driven multi-agent decision processes," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 4, pp. 1259–1268, Apr. 2019.
- [5] A. O. Al-Abbasi, A. Ghosh, and V. Aggarwal, "DeepPool: Distributed model-free algorithm for ride-sharing using deep reinforcement learning," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 12, pp. 4714–4727, Dec. 2019.
- [6] J. Ke, F. Xiao, H. Yang, and J. Ye, "Optimizing online matching for ride-sourcing services with multi-agent deep reinforcement learning," *CoRR*, vol. abs/1902.06228, Feb. 2019.
- [7] R. Luo, T. J. van den Boom, and B. De Schutter, "Multi-agent dynamic routing of a fleet of cybercars," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 5, pp. 1340–1352, May 2018.
- [8] Z. Zhou, B. De Schutter, S. Lin, and Y. Xi, "Two-level hierarchical model-based predictive control for large-scale urban traffic networks," *IEEE Trans. Control Syst. Technol.*, vol. 25, no. 2, pp. 496–508, Mar. 2017.
- [9] F. Miao *et al.*, "Taxi dispatch with real-time sensing data in metropolitan areas: A receding horizon control approach," *IEEE Trans. Autom. Sci. Eng.*, vol. 13, no. 2, pp. 463–478, Apr. 2016.
- [10] R. Zhang and M. Pavone, "Control of robotic mobility-on-demand systems: A queueing-theoretical perspective," *Int. J. Robot. Res.*, vol. 35, nos. 1–3, pp. 186–203, Jan. 2016.
- [11] T. Oda and C. Joe-Wong, "MOVI: A model-free approach to dynamic fleet management," in *Proc. IEEE INFOCOM*, Apr. 2018, pp. 2708–2716.
- [12] X. Tang *et al.*, "A deep value-network based approach for multi-driver order dispatching," in *Proc. KDD*, 2019, pp. 1780–1790.
- [13] Z. Xu *et al.*, "Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach," in *Proc. KDD*, 2018, pp. 905–913.
- [14] M. L. Littman and C. Szepesvári, "A generalized reinforcement-learning model: Convergence and applications," in *Proc. Int. Conf. Mach. Learn.*, 1996, pp. 1–18.
- [15] B. C. Csáji and L. Monostori, "Value function based reinforcement learning in changing Markovian environments," *J. Mach. Learn. Res.*, vol. 9, pp. 1679–1709, Aug. 2008.
- [16] C. Szepesvári and M. L. Littman, "A unified analysis of value-function-based reinforcement-learning algorithms," *Neural Comput.*, vol. 11, no. 8, pp. 2017–2060, 1999.
- [17] K. G. Vamvoudakis and F. L. Lewis, "Online actor-critic algorithm to solve the continuous-time infinite horizon optimal control problem," *Automatica*, vol. 46, no. 5, pp. 878–888, May 2010.
- [18] T. T. Doan, S. T. Magalur, and J. Romberg, "Finite-time analysis of distributed TD(0) with linear function approximation on multi-agent reinforcement learning," in *Proc. ICML*, in Proceedings of Machine Learning Research, vol. 97, 2019, pp. 1626–1635.
- [19] L. Busoniu, R. Babuska, and B. De Schutter, "Multi-agent reinforcement learning: A survey," in *Proc. 9th Int. Conf. Control, Automat., Robot. Vis.*, 2006, pp. 1–6.
- [20] S. Bandyopadhyay, S.-J. Chung, and F. Y. Hadaegh, "Probabilistic and distributed control of a large-scale swarm of autonomous agents," *IEEE Trans. Robot.*, vol. 33, no. 5, pp. 1103–1123, Oct. 2017.
- [21] H. W. Kuhn, "The Hungarian method for the assignment problem," in *50 Years of Integer Programming*. Berlin, Germany: Springer, 2010, pp. 29–47.
- [22] D. P. Bertsekas, "The auction algorithm: A distributed relaxation method for the assignment problem," *Ann. Oper. Res.*, vol. 14, no. 1, pp. 105–123, Dec. 1988.
- [23] D. P. Bertsekas and D. A. Castañón, "Parallel synchronous and asynchronous implementations of the auction algorithm," *Parallel Comput.*, vol. 17, nos. 6–7, pp. 707–732, Sep. 1991.
- [24] J. Bellingham, M. Tillerson, A. Richards, and J. P. How, "Multi-task allocation and path planning for cooperating UAVs," in *Cooperative Control: Models, Applications and Algorithms*, S. Butenko, R. Murphy, and P. M. Pardalos, Eds. Boston, MA, USA: Springer, 2003, pp. 23–41.
- [25] D. Morgan, G. P. Subramanian, S.-J. Chung, and F. Y. Hadaegh, "Swarm assignment and trajectory optimization using variable-swarm, distributed auction assignment and sequential convex programming," *Int. J. Robot. Res.*, vol. 35, no. 10, pp. 1261–1285, 2016.
- [26] D. Dionne and C. A. Rabbath, "Multi-UAV decentralized task allocation with intermittent communications: The DTC algorithm," in *Proc. Amer. Control Conf.*, Jul. 2007, pp. 5406–5411.
- [27] P. B. Sujit and R. Beard, "Distributed sequential auctions for multiple UAV task allocation," in *Proc. Amer. Control Conf.*, Jul. 2007, pp. 3955–3960.
- [28] J. R. Marden and J. S. Shamma, "Revisiting log-linear learning: Asynchrony, completeness and payoff-based implementation," *Games Econ. Behav.*, vol. 75, no. 2, pp. 788–808, Jul. 2012.
- [29] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (Wiley Series in Probability and Statistics). Hoboken, NJ, USA: Wiley, 1994.
- [30] M. L. Puterman and M. C. Shin, "Modified policy iteration algorithms for discounted Markov decision problems," *Manage. Sci.*, vol. 24, no. 11, pp. 1127–1137, Jul. 1978.
- [31] C. Szepesvári, *Algorithms for Reinforcement Learning* (Synthesis Lectures on Artificial Intelligence and Machine Learning). San Rafael, CA, USA: Morgan & Claypool, 2010.
- [32] A. M. Stuart and T. Humphries, *Dynamical Systems and Numerical Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 1996.
- [33] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2012.
- [34] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for non-linear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [35] H. Tsukamoto and S.-J. Chung, "Robust controller design for stochastic nonlinear systems via convex optimization," *IEEE Trans. Autom. Control*, early access, Nov. 16, 2020, doi: [10.1109/TAC.2020.3038402](https://doi.org/10.1109/TAC.2020.3038402).
- [36] P. M. Wensing and J.-J. Slotine, "Cooperative adaptive control for cloud-based robotics," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, May 2018, pp. 6401–6408.
- [37] R. Olfati-Saber, "Kalman-consensus filter: Optimality, stability, and performance," in *Proc. IEEE Int. Conf. Decis. Control*, Dec. 2009, pp. 7036–7042.
- [38] S. Bandyopadhyay and S.-J. Chung, "Distributed Bayesian filtering using logarithmic opinion pool for dynamic sensor networks," *Automatica*, vol. 97, no. 1, pp. 7–17, Nov. 2018.
- [39] *Chicago Data Portal*. Accessed: Mar. 29, 2021. [Online]. Available: <https://data.cityofchicago.org/Transportation/Taxi-Trips/wrvz-psev>
- [40] B. Rivière, W. Hönig, Y. Yue, and S. Chung, "GLAS: Global-to-local safe autonomy synthesis for multi-robot motion planning with end-to-end learning," *IEEE Robot. Autom. Lett.*, vol. 5, no. 3, pp. 4249–4256, Jul. 2020.



Benjamin Rivière (Graduate Student Member, IEEE) received the B.S. degree in mechanical engineering from Stanford University in 2017 and the M.S. degree in aeronautics from Caltech in 2018. He is currently pursuing the Ph.D. degree with the California Institute of Technology. His research focus is planning for high dimensional systems with informed search, data-driven, and control-theoretic methods. He was a recipient of the 2020 Honorable Mention for the IEEE ROBOTICS AND AUTOMATION LETTERS Best Paper Award.



Soon-Jo Chung (Senior Member, IEEE) received the B.S. degree (*summa cum laude*) from the Korea Advanced Institute of Science and Technology, Daejeon, South Korea, in 1998, and the S.M. degree in aeronautics and astronautics and the Sc.D. degree in estimation and control from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2002 and 2007, respectively. He is currently the Bren Professor of aerospace and control and dynamical systems and a Jet Propulsion Laboratory Research Scientist with the California Institute of Technology. He was on the Faculty of the University of Illinois at Urbana-Champaign (UIUC) from 2009 to 2016. His research focuses on spacecraft and aerial swarms, autonomous aerospace systems, theory and application of complex nonlinear dynamics, control, estimation, guidance, and navigation of autonomous space, and air vehicles. He received the UIUC Engineering Deans Award for Excellence in research, the Beckman Faculty Fellowship of the UIUC Center for Advanced Study, the U.S. Air Force Office of Scientific Research Young Investigator Award, the National Science Foundation Faculty Early Career Development Award, and three Best Conference Paper Awards from the IEEE and the American Institute of Aeronautics and Astronautics. He is also an Associate Editor of IEEE TRANSACTIONS ON ROBOTICS, IEEE TRANSACTIONS ON AUTOMATIC CONTROL, and *Journal of Guidance, Control, and Dynamics*.