

Cooperation for Scalable Supervision of Autonomy in Mixed Traffic

Cameron Hickert , Member, IEEE, Sirui Li , Member, IEEE, and Cathy Wu , Member, IEEE

Abstract—Advances in autonomy offer the potential for dramatic positive outcomes in a number of domains, yet enabling their safe deployment remains an open problem. This work’s motivating question is: In safety-critical settings, can we avoid the need to have one human supervise one machine at all times? The work formalizes this *scalable supervision* problem by considering remotely located human supervisors and investigating how autonomous agents can cooperate to achieve safety. This article focuses on the safety-critical context of autonomous vehicles (AVs) merging into traffic consisting of a mixture of AVs and human drivers. The analysis establishes high reliability upper bounds on human supervision requirements. It further shows that AV cooperation can improve supervision reliability by orders of magnitude and counterintuitively requires fewer supervisors (per AV) as more AVs are adopted. These analytical results leverage queuing-theoretic analysis, order statistics, and a conservative, reachability-based approach. A key takeaway is the potential value of cooperation in enabling the deployment of autonomy at scale. While this work focuses on AVs, the scalable supervision framework may be of independent interest to a broader array of autonomous control challenges.

Index Terms—Autonomous agents, human factors and human-in-the-loop, intelligent transportation systems, scalable supervision.

I. INTRODUCTION

GIVEN the complexity, chaos, and unpredictability of real-world environments, safety is a critical challenge for autonomous systems. This is particularly the case for mixed autonomy systems, which refer to systems in which humans and machines both exhibit control in the same environment [1]. Indeed, in an Allianz Global Assistance survey investigating decreasing interest in self-driving cars (one of the most-discussed

mixed autonomy systems today), over 70% of respondents cited safety as a reason for their lack of interest, representing increased concern from the year prior [2]. Thus, improving safety in mixed autonomy systems is not only an end in itself, but also is a prerequisite for unlocking the benefits that such systems may offer.

The question becomes: How can we achieve sufficient levels of safety and performance even when formal guarantees are elusive in a chaotic world? To investigate this topic, our research borrows the idea of remote supervision from air traffic control, in which a centralized controller coordinates agents for safety.

A. Contributions

his work provides one perspective on the topic of improving safety in mixed autonomy settings by investigating the online human supervision of autonomous vehicles (AVs) in an illustrative merging task. To this end, we present a number of contributions that—to the best of our knowledge—have not been addressed in previous work.

- 1) We formalize the *scaling supervision* problem for online autonomous agents and propose a “scalability” metric: expected number of AVs per supervisor. In the process, we propose the first (to our knowledge) large-scale remote supervision and control framework for AVs and present queuing-theoretic and statistical tools for analysis of such systems.
- 2) We propose a reachability analysis-based method for AV supervision. We provide a high reliability upper bound on the expected supervision time required for a supervisor to monitor a given AV’s merge, as well as a closed-form expression for the probability that the number of supervisors is insufficient when merging tasks are pooled together and distributed across multiple human supervisors.
- 3) Leveraging order statistics, we show that—interestingly—cooperation of AVs enables supervision time *inversely* related to AV adoption (i.e., the proportion of vehicles in a traffic network that are AVs), as well as orders-of-magnitude improvement in supervision requirements. An implication of this result is that cooperation substantially facilitates safe autonomous deployments at scale.

These outcomes suggest the adage, “An ounce of prevention is worth a pound of cure,” applies here. That is, preventing the need for supervision is more powerful than having more supervisors.

Manuscript received 19 April 2022; revised 24 November 2022; accepted 20 February 2023. Date of publication 12 April 2023; date of current version 8 August 2023. This work was supported in part by the MIT Amazon Science Hub, in part by MIT-IBM Watson AI Lab, in part by MIT Energy Initiative (MITEI) Mobility Systems Center, in part by MIT Mobility Initiative, and in part by the National Science Foundation (NSF) under Grant 2149548. This paper was recommended for publication by Associate Editor M. Walter. (Corresponding author: Cameron Hickert.)

Cameron Hickert and Sirui Li are with the Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: chickert@mit.edu; siruil@mit.edu).

Cathy Wu is with the Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA 02139 USA, also with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139 USA, and also with the Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: cathywu@mit.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TRO.2023.3262120>.

Digital Object Identifier 10.1109/TRO.2023.3262120

II. BACKGROUND

A. Motivation

a) Safety in AVs: The question of AV safety is an open and well-documented problem [3], [4], particularly in mixed autonomy traffic settings [5], [6]. Thus, the importance of safety in any deployment of AVs is sufficient justification for investigating supervision of autonomous agents in itself. Human supervision of autonomous and semiautonomous agents has been shown to improve system safety in other settings and is the de facto safety backup method for driver-assist systems on the road today [7], [8], [9]. If AVs replaced even a modest portion of US vehicles, the number of lives depending on AV safety systems could number in the thousands [10].

Moreover, AVs are not the only autonomous system with the potential to inflict human harm. The authors anticipate that the approach and lessons learned in this research can inspire future work to improve the safety and performance of autonomous agents in settings beyond society's roadways, such as in manufacturing and other cyber-physical systems.

This work adopts an operational reliability perspective on safety. That is, the proposed method does not guarantee crash-free behavior, but rather focuses on identifying the resources necessary to achieve human-level safety with an arbitrary desired level of reliability. This is achieved via human supervision of the autonomous agents and thus the system's ultimate safety remains subject to human performance. The concept of remote supervision is borrowed from air traffic control, in which a human supervisor centrally coordinates agents for safety [11].

b) AVs' Benefits to Traffic Systems: While ensuring a sufficient level of safety is an important end in itself, doing so is foundational to the wider adoption of AVs, which can in turn bring benefits including decreased traffic congestion and traffic emissions. AVs can achieve outsized positive impact without dominating roadways. Using AV policies learned via reinforcement learning (RL), previous work shows that even when AVs account for only 5%–10% of vehicles in a traffic network, they can still boost the collective average speed of *all* vehicles in the network by up to 57% in idealized settings [12]. Researchers found that improving traffic flows in California can reduce highway carbon dioxide emissions by up to nearly 20%; this is particularly significant given that transportation accounts for approximately a third of all US carbon dioxide emissions [13].

c) Merging: The case of AVs merging onto a highway with traffic is an appropriate task for investigating how we can scale supervision for a number of reasons. First, it is a common occurrence that can pose difficult control and coordination challenges for AVs [14], [15]. Given the speed at which freeway merges occur, this task can result in deadly crashes. Indeed, successful merging may sometimes rely on implicit communication between drivers or even nontraditional explicit communication, such as hand waves. For these reasons a “sufficient level” of safety in this article refers to that which capable and attentive human drivers provide. This is achieved via the use of human supervision.

Merging also presents unique challenges in mixed autonomy settings. Were a system to consist entirely of AVs, a solution may exist via decentralized or centralized coordination, but that is much more difficult—if not impossible—for the foreseeable future, in which human drivers still populate the roadways [16]. Due to its relative frequency, unpredictability, and potential for impact across entire traffic systems, highway merging has been a particularly important and popular maneuver to study [17], [18], [19].

d) Other applications: The scalable supervision framework, concepts, and formalization presented here are adaptable to other autonomous machine settings. Within mixed autonomy traffic environments, one could consider scalable supervision as a means of supporting dedicated AV lanes on freeways. These lanes have been shown in simulation to improve traffic efficiency while also worsening on-road safety [20]. The potential speed difference between AV-only and standard lanes present a merging problem similar to the case studied here. Scalable supervision would also be useful for the case of accident-prone “hotspots” (construction sites, school zones, etc.) where concerns arise about the potential behavior of AVs [21]. Each of these is a practical instance of the scenario studied here.

Beyond traffic scenarios, reachability-based supervision scaling that enforces safety constraints may be applicable to domains as varied as warehouses, ports, airports, satellite orbits, and disaster response. Our work can be situated among these domains along the axes of spatial scale and reliability requirements. In robotics applications with human supervisors, the spatial scale is on the order of meters or less, but reliability requirements and response time conditions are not as stringent [22], [23]. In human supervision of (increasingly autonomous) air and maritime traffic applications, the reliability requirements are high, but the spatial scale is on the order of tens of kilometers or more, corresponding to a larger temporal scale [24], [25]. Our work seeks to improve supervision of autonomous agents in a direction of both high reliability (up to 99.9999%) and meter-scale spatial resolution. Note that application to other domains requires computation of the relevant reachable sets, but—as we will show—conservative estimates of such sets in many of these applications may provide useful approximations that are more computationally efficient.

More broadly, policies learned via RL have demonstrated superhuman performance in domains from board games and video games to high-altitude balloon control and simulated fighter jet combat [26], [27], [28]. However, deep RL methods lack the convergence guarantees of more traditional RL approaches, and thus cannot be trusted in safety-critical settings [29], [30].

This work is similarly relevant to imitation learning, another approach to AV deployment that struggles to reliably maintain safety [31]. While progress has been made on various imitation learning challenges, in practice humans remain a fail safe [32].

B. Related Work

This research draws from interesting and important related work in at least three areas of the literature: human–artificial

intelligence (AI) teaming, supervision and cognitive models, and reachability analysis.

a) Human–AI teaming: Relevant work includes that from the human–robot interaction subfield in multiagent systems, which explores how to reduce a single human operator’s cognitive load from scaling linearly with the number of robots controlled. A significant portion of this work investigates the possibility for well-designed interfaces to reduce task burden or for learned models of human preferences to automate human–machine control handoff (also called, sliding, shared, or adjustable autonomy) [33]. Our work expands this effort by considering the shared autonomy situation in which strict safety guarantees must be maintained, and extends this line of work into the mixed autonomy traffic setting.

Research, such as that described in [34], seeks to optimize the allocation of limited human assistance (via a decision support system) to multiple robots performing tasks in a given environment. The machines in these works benefit from human assistance, but do not exhibit the strict safety-critical requirements of AVs in traffic. As such, these works also place greater emphasis on computing and comparing the various assistance permutations. Such comparisons are not necessary given our binary approach to safety and strict safety requirements (i.e., no suboptimal safety configurations are permitted).

b) Scalable supervision and cognitive models: To the best of our knowledge, previous work concerned with decreasing the time burden of human supervision of autonomous tasks tends to be directed at empirical studies of a single human monitoring multiple UAVs (unmanned aerial vehicles) [35], [36], [37]. Relatedly, human–robot interaction theory developed in the early 2000s considers “fanout,” the number of robots that a human can effectively control [36], [37], [38], [39], [40], [41]. While philosophically aligned with our work on some levels, in the AV merge setting we isolate a single task to analyze and face safety challenges unique to the mixed autonomy setting. In addition, the mixed autonomy AV setting considered here is arguably higher stakes than a typical UAV setting, as it includes human drivers operating in the same vicinity; crashes would more likely involve human fatalities.

Our work can be seen as complementary to the research described in [42] and [43], which also applies air traffic control principles to human, remote supervision of AV fleets. Our reliability metric parallels the “exceedance probability” in [43], but whereas that work’s queuing analysis assumes batch arrivals of simulation-generated tasks upon which humans provide input (but not direct control), supervision tasks in our case are assigned to individual remote operators who control the vehicle for the entire task duration. More broadly, our approach to defining the conditions requiring supervision is compatible with these works, which assume AVs can predict such scenarios. Finally, these works do not consider how AV cooperation can facilitate supervision scaling.

A closely related work is that on scaled autonomy by Swamy et al., in which the authors investigate how to assist an operator in selecting which robots in a fleet most require teleoperation [44]. Our work differs by seeking to minimize the supervisor’s active control time (the teleoperator in that work is never idle) and the

mixed autonomy setting at the heart of our work must maintain strict safety standards.

Broader research in this direction has considered a single machine supervisor observing vehicles with known paths [45] and algorithmic supervision of human drivers—in which a machine overrides human driving controls when the human is behaving in ways known to be unsafe [46].

c) Reachability analysis: At its most basic, reachability analysis is about identifying the set of states that a dynamical system could enter—the “reachable states” or “reachable set”—given all admissible inputs and parameters [47]. A variety of methods exist for doing so; a common tradeoff among these is that between approximation and computational inefficiency [48], [49]. Methods, such as sampling-based approaches may require less computation, but they are not fully conservative, meaning they are not guaranteed to compute the reachable set [49]. For safety-critical applications, such approximations may be unacceptable.

One method for computing the exact reachable set is Hamilton–Jacobi (HJ) reachability. HJ methods can handle nonlinear system dynamics and allow for formal treatment of bounded disturbances, but suffer exponential computational complexity in terms of the number of state variables [50]. Fortunately, the vehicle dynamics and merging problem structure allow us to adopt a fully conservative approach while avoiding the computational complexity. Our proposed reachability-based method calculates the exact reachable set for each vehicle in $O(1)$ time. However, we note that utilizing other mechanisms of computing conservative reachable sets (such as HJ reachability) could be used in future work generalizing to more scenarios [50]. Other work that seeks to integrate reachability analysis and AVs relaxes the conservatism of their approaches by making assumptions about human driving behaviors [51], or by not accounting for the full range of system dynamics [52]. The need remains for an approach that maintains strict safety guarantees.

III. FORMALIZING SCALABLE SUPERVISION

In this section, we provide a formalization of scalable supervision.

A. Formalizing Scalable Supervision

At its core, the two goals of scaling supervision are to reduce the number of necessary human interventions and to reduce the time required for those that must occur. One way to do this (our approach) is by utilizing reachability analysis to identify when a merging AV is in danger of colliding with another vehicle in the system, and activate human control of the AV only at that time. This prevents unnecessary human interventions (e.g., for a merge when no other vehicles are around), and only calls for supervision at the moment transition of control is necessary, achieving the two goals previously. It is also important to note that “safety” in this article is not synonymous with crash-free behavior. This aligns with the operational reliability perspective on safety described previously. Since the safety backup for the AV is a human driver, the ultimate safety of the system remains subject to human fallibility.




Fig. 1. General illustration of the reachability conditions that would activate a human supervisor. The yellow rectangular zones represent the reachable distances over time horizon t for each vehicle. When the merge point (marked by the red cross) is within both vehicles' reachable zones, there is a possibility of an accident in the time horizon and the supervisor activates.

To understand how our approach works, first imagine an AV on an on-ramp merging into a highway on which a human vehicle (HV) is already driving. The objective is to prevent the two vehicles from colliding during the AV's merge, so the reachability question, simply put, is whether it is possible for both the AV and HV to get to the merge point over some finite time horizon t . The intuition here is that a supervisor need not assume control of the AV if a potential collision is far in the future; it is enough to assume control before the potential crash but with enough time to avoid it comfortably. We can vary t to be more or less cautious in any given scenario. In at least one setting humans required 5–8 s to safely assume control of a vehicle; this is the time window we adopt in generating results [53].

The reachability question naturally decomposes into two independent subquestions: 1) Can the HV reach the merge point? and 2) Can the AV reach the merge point? That is, for each vehicle, is the distance it would travel if it were to apply its maximum acceleration over the horizon t greater than or equal to the distance between that vehicle's current position and the merge point? If the answer to *both* subquestions is "yes," then human supervisor control of the AV must be activated to ensure human-level safety. See Fig. 1 for a visual representation.

We can write this more formally: the horizon t_{sup} within which the supervisor must assume control of a merging AV for safety (denoted by the AV subscript) is

$$t_{\text{sup}} = \min\{t \mid t \geq 0, (d_{\text{HV}}(v_{\text{HV},0}, a_{\text{HV,max}}, t) \geq d_{m,\text{HV}}) \wedge (d_{\text{AV}}(v_{\text{AV},0}, a_{\text{AV,max}}, t) \geq d_{m,\text{AV}})\} \quad (1)$$

where, $d_i(v_{i,0}, a_{i,\text{max}}, t)$ is the maximum distance vehicle i can travel in time horizon t and is a function of that vehicle's initial velocity $v_{i,0}$ and its maximum acceleration $a_{i,\text{max}}$, $d_{m,i}$ is the distance between vehicle i (either an HV or an AV) and the merge point m , and the HV is on the highway onto which the AV is merging. The road networks considered here are limited to one-lane roads and thus considering a vehicle's maximum forward distance is sufficient. This amounts to the most conservative reachability for the given time horizon: there is no location that




Fig. 2. Illustrative vehicle merging scenario. A merging AV's path is indicated by the blue arrows; it enters via an on-ramp on the right-hand side of the circular highway and exits via the off-ramp at the bottom. This AV's merge point is indicated with the red marker. Human vehicles are shown in gray. AVs are shown in blue. (Not to scale; modified for visibility.)

the vehicle could reach over t time that is not within its reachability zone. Note that this maintains safety guarantees even when generalizing to cases in which two-dimensional motion is considered (such as by adding lane changes). The maximum distance reachability formulation still provides a conservative reachable zone because side-to-side motion only can reduce the distance along the road that the vehicle travels, rather than increase it.

B. Scenario

Because this work focuses on merging, the analysis begins by considering a large rotary traffic network, which approximates a highway with multiple on-ramps and off-ramps. The ring road at the core of the structure is a well-studied setting in traffic literature and has been shown to mimic traffic congestion patterns that might occur on an infinite roadway [54]. The ring therefore emulates a highway. Merging vehicles enter the system via the on-ramps, merge into the ring road, and then exit via off-ramps (see Fig. 2).

Since this is a mixed autonomy system, we consider both AVs and HVs. Acceleration profiles are parameterized via the widely used intelligent driver model (IDM), which has been shown to emulate the actual behavior of human drivers [55], [56]. We first present analysis for mixed autonomy traffic with unconnected AVs (UCAVs). These are AVs that cannot communicate with other AVs in the system (but still allow for remote supervision). Next, we consider noncooperative connected AVs (NCAVs)—also using IDM—to investigate the supervision scaling implications of connected AVs. These are modeled such that they can communicate with other connected AVs in the system, but do not alter their behavior to accommodate merging connected AVs. We will see that they allow for linear improvements in supervision scaling. Finally, we analyze cooperative connected

TABLE I

VEHICLE TYPES AND KEY PROPERTIES. ALL AVS ARE ASSUMED TO ALLOW FOR REMOTE SUPERVISION AND CONTROL; THE CONNECTIVITY REFERENCED IN THE TABLE (“UNCONNECTED,” “CONNECTED,”) REFERS TO AN AV’S ABILITY TO COMMUNICATE WITH OTHER AVS IN THE SYSTEM, IN PARTICULAR THE MERGING AV

Vehicle type	Acronym	Properties		
		Requires remote supervision for dangerous merges	Shares near-term trajectory w/ merging AV (resulting in truncated reachability zone)	Adjusts behavior to accommodate AV merges
Human vehicle	HV			
Unconnected AV	UCAV	✓		
Noncooperative connected AV	NCAV	✓	✓	
Cooperative connected AV	CCAV	✓	✓	✓

AVs (CCAVs): these represent supervision-aware AVs that not only can communicate with other connected AVs in the system, but also seek to assist merging connected AVs to effectively join highway traffic while reducing the need for human supervision. Via analytical results we show how the teaming of CCAVs enables orders-of-magnitude improvement in supervision requirements. See Table I for a summary of the vehicle types and their properties. Note that all AVs (including UCAVs) are assumed to allow for remote supervision and control.

IV. THEORETICAL ANALYSIS

By applying reachability analysis to the problem of AV merging, a number of bounds can be derived for the settings in which we are interested. To analytically quantify the risk and scalability of supervision, we take a queuing theoretic approach in this section to analyze long-term steady-state conditions.

This analysis begins with a description of an upper bound on the supervision requirements for the scenario in which a group of remote supervisors manage merges for an arbitrarily large number of on-ramps, as well as a closed-form expression for calculating the probability that AVs need supervision but cannot receive it. This is of particular interest because it allows for the characterization of the risk associated with a system with a fixed number of supervisors. Conversely, it allows one to calculate the number of supervisors necessary to achieve a desired level of supervision safety. After the provision of the upper bound and analysis of the UCAV and NCAV settings, an analysis of the “typical case” is provided, wherein the expected supervision scaling gains via (supervision-aware) CCAVs are described.

Let C be a *potential* conflict event. Here this refers to the joint event that a merge point falls within the reachable zone of at least one vehicle in the ring (event C_{in}) and that the same merge point falls within the reachable zone of a vehicle outside the ring that is attempting to merge in (event C_{out}). C_{in} and C_{out} thus correspond to the subquestions outlined in Section III-A. In different settings one may consider alternative dangerous situations and methods other than reachability analysis to indicate them. One necessary component of the remote supervisor analysis is a characterization of $P(C_{in})$, that is, the probability that the merge point falls within the reachable zone of at least one vehicle on the ring. The analysis further below investigates this by considering how an upper bound on $P(C_{in})$ varies in the settings in which the in-ring vehicles are entirely HVs, a mix of HVs and NCAVs, and finally a mix of HVs and (supervision-aware) CCAVs. Note the NCAVs are *not* antagonistic; they do not seek to inhibit merges, but simply do not alter their path to accommodate them.

It is important to first note that we can rewrite (1) using kinematics. We allow each vehicle to take any nonnegative velocity $v \in [0, \infty)$ and only assume finite vehicle acceleration $a_i \in [a_{i,min}, a_{i,max}]$. This provides conservative, physics-informed safety: it is impossible for a vehicle to cover more distance than it does when maximally accelerating over the time horizon’s duration. Thus, preserving this view of reachability requires accounting for the worst-possible case, and the reachable set for vehicle i at position 0 is $[0, d_i(v_{i,0}, a_{i,max}, t)]$.

This conservative approach to reachability has the ancillary benefit that vehicle i ’s reachable distance in the given time horizon can be written simply

$$d_i(v_{i,0}, a_{i,max}, t) = v_{i,0}t + \frac{1}{2}a_{i,max}t^2. \quad (2)$$

While this work uses IDM values for the vehicles’ maximum accelerations, in practice acceleration profiles can be derived from data for determining a pragmatic upper bound for a given setting [57]. Similarly, one could consider decoupling t in the above from vehicle dynamics. For example, it (or even the reachable distance itself) could be provided by a prediction system that estimates vehicles’ forward progress [58]. This work uses the kinematics-based approach to maintain strict safety requirements, but we expect much of the analysis to apply to less conservative approaches. In light of this, in the statements below these components are left as variables (although sometimes omitted from explicit inclusion for ease of exposition). Specific values are only applied during the generation of the numerical results and are sourced from relevant literature. With this in mind, we turn to the challenge of characterizing the multiple-supervisor, multiple-merge case.

A. Multiple Supervisors Monitoring an AV Fleet

We consider a future setting in which a team of human supervisors located remotely could monitor a fleet of AVs and assume remote control if necessary. A natural question would be: how many supervisors are necessary?

Theorem 1: Suppose we have k remote supervisors and n on-ramps on which AVs appear and trigger the on-ramp supervision condition with independent arrival processes $\text{Poisson}(\lambda_1), \dots, \text{Poisson}(\lambda_n)$. Suppose the service rate of each remote supervisor follows $\text{Exp}(\mu)$.

The fraction of AVs that require supervision but cannot immediately receive it (and thus go unsupervised) is given by

$$P_k = \frac{(\lambda/\mu)^k/k!}{\sum_{i=0}^k (\lambda/\mu)^i/i!} \quad (3)$$

where, $\lambda = \sum_{j=1}^n \lambda_j P(C_{in})$, $P(C_{in})$ is the probability that an in-ring vehicle will trigger its supervision condition, and $\lambda < \mu$ in order to have a valid steady state probability.

Proof: Allow different on-ramps to have different arrival rates $\lambda_1, \dots, \lambda_n$. Because the supervision tasks are allocated to a centralized group of remote supervisors, the problem reduces from a tasking situation involving n separate queues to a single-queue tasking case. Given the additivity of independent Poisson random variables [see (30) in the appendix], the arrival rate of AVs requiring supervision follows $\text{Poisson}(\sum_{j=1}^n \lambda_j P(C_{in}))$. Poisson arrivals are known to successfully model true vehicle arrivals and are commonly used highway on-ramp arrival modeling [59], [60], [61], [62]. Exponential service rates are used in modeling air and roadway traffic supervision tasks in [42] and [63]. Set

$$\lambda = \left(\sum_{j=1}^n \lambda_j \right) P(C_{in}).$$

Let X be the random variable denoting the number of AVs requiring control at an arbitrary point in time. Without loss of generality, assume $\lambda_1 = \dots = \lambda_n = \hat{\lambda}/n$ for some constant $\hat{\lambda}$, then the AVs arriving on the on-ramp trigger the supervision condition following the distribution $X \sim \text{Poisson}(\sum_{i=1}^n \hat{\lambda}/n) = \text{Poisson}(\hat{\lambda})$ [again, see (30) in the appendix]. Let Y be a random variable denoting the number of AVs that need to be supervised. We know $Y|X \sim \text{Binomial}(X, P(C_{in}))$ where $P(C_{in})$ can be obtained from the in-ring vehicle's reachability condition and will be derived further in the following.

By (29) (see Appendix), we have that $Y \sim \text{Poisson}(\hat{\lambda}P(C_{in}))$ is the arrival process for AVs that need to be supervised. Then, the problem reduces to an $M/M/k$ queue with arrival process $\text{Poisson}(\hat{\lambda}P(C_{in}))$, service rate $\text{Exp}(\mu)$ for each supervisor, and finite capacity k . The finite capacity is due to the fact that when all supervisors are busy, any additional on-ramp AVs requiring supervision will be rejected immediately. These cases cannot wait on the queue for future service as their needs are immediate; that is, to maintain safety, once the on-ramp and in-ring conditions are triggered for a given merge, a supervisor must immediately supervise the merge. Thus, the rejected cases that the pool of k supervisors cannot immediately service represent dangerous situations.

For an $M/M/k$ queue with no waiting space, both the steady state probability and loss formula are known [64]. The expression for the steady state probability is

$$P_q = \frac{(\lambda/\mu)^q/q!}{\sum_{i=0}^k (\lambda/\mu)^i/i!} \quad (4)$$

where, $q \in \{0, \dots, k\}$. Then, q represents the number of AVs requiring supervision, k is the total number of supervisors (as previously defined), the arrival rate is $\lambda = \hat{\lambda}P(C_{in})$, and μ is kept as the service rate. The loss formula here represents the fraction of AVs that require supervision but cannot immediately receive it (and thus go unsupervised), and the result follows. ■

Based on (3), the expected number of supervisors needed is the smallest k such that $P_k \leq \delta$ for some risk tolerance of interest $0 < \delta \ll 1$. Various examples of the relationships encoded

in (3) are visualized in Figs. 3 and 4 for a mixed HV/UCAV traffic network setting (described in the following) with fixed inflow and service rates. These trends were found to hold across a variety of configurations. In short, these answer the question, "How many supervisors do I expect to need?" in the UCAV case, assuming the questioner can define their risk threshold δ .

In the previous, observe that—due to the functional reduction of the n inflows to a single task queue serviced by the k supervisors—the exact number of on-ramps does not directly affect the steady state probability or loss formula, except insofar as the individual arrival rates $\lambda_1, \dots, \lambda_n$ contribute to the global arrival rate.

We recognize that the exponential distribution of service times inherent to a Poisson process can present unrealistic scenarios (service rarely takes zero time). Indeed, this is a known limitation when using Poisson processes for modeling physical systems. A truncated exponential distribution can mitigate this issue and can be handled explicitly using discrete event simulation. The numerical analysis below adopts rates μ substantially larger than the reachability horizon to implicitly account for this shift.

Finally, note the importance of the *in-ring reachability* $P(C_{in})$ in determining whether an AV requires supervision. This follows naturally from the problem formalization, in which both the in-ring and on-ramp supervision conditions (reachability conditions) must be met in order for a merging AV to require supervision. The following sections provide insight into this term. The abovementioned equations are provided as functions of $P(C_{in})$, among other variables. In the following lemmas, corollaries, and theorem, we consider various cases where we can compute $P(C_{in})$ analytically to investigate how the system's supervisability under various conditions.

B. In-Ring Reachability for Mixed HVs and UCAVs

Lemma 1: Given a single-lane ring road of circumference c with R (HVs or UCAVs) distributed uniformly at random (but not necessarily independently) along the length of the ring, the probability that an arbitrary fixed point on that ring is reachable over time horizon t by a vehicle is bounded above as follows:

$$P(C_{in}) \leq \frac{\sum_{i=1}^R d_i(v_{i,0}, a_{i,\max}, t)}{c} \quad (5)$$

where, $d_i(v_{i,0}, a_{i,\max}, t)$ is the reachable range for vehicle i over the time horizon.

Proof: Suppose X is the fixed merge point and the R unconnected vehicles are composed of $R - S$ HVs and S UCAVs. Let H_i be the asymmetric distance (following traffic flow) between the i th in-ring unconnected vehicle and X . Then the event

$$\begin{aligned} C_{in} &= \cup_{i=1}^{R-S} \{\text{HV } i \text{ triggers supervision cond.}\} \\ &\cup \left(\cup_{j=1}^S \{\text{UCAV } j \text{ triggers sup. cond.}\} \right) \\ &= \cup_{i=1}^R \{H_i \leq d_i(v_{i,0}, a_{i,\max}, t)\} \end{aligned} \quad (6)$$

where, the first equality comes from the fact that the event C_{in} happens if any of the unconnected vehicles triggers the supervision condition, and the second equality comes from the reachability analysis for each HV and UCAV. Because the

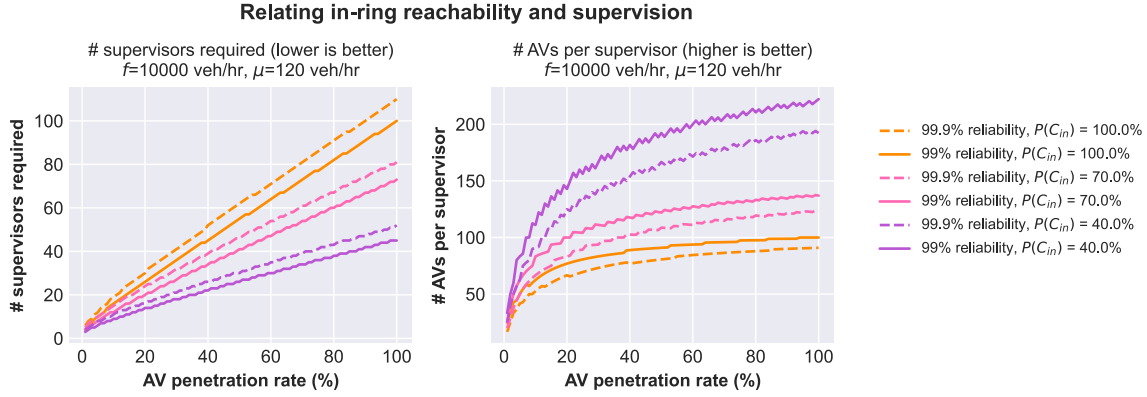


Fig. 3. Illustration of the relationships in (3) for the “baseline” case with UCAVs. f is the total vehicle inflow rate. “AV penetration rate” refers to the percentage of vehicles in the system (including inflows) that are autonomous. Note that as the AV penetration rate grows, so does the number of supervisors needed to maintain a given safety threshold. This is because the number of merging AVs grows with greater penetration rates. Also note how altering $P(C_{in})$ affects the plots through its effect on λ . The μ value represents monitors that can supervise a merge in an average of 30 s. The selection rationale for the values of μ and f is described in Section IV-I.

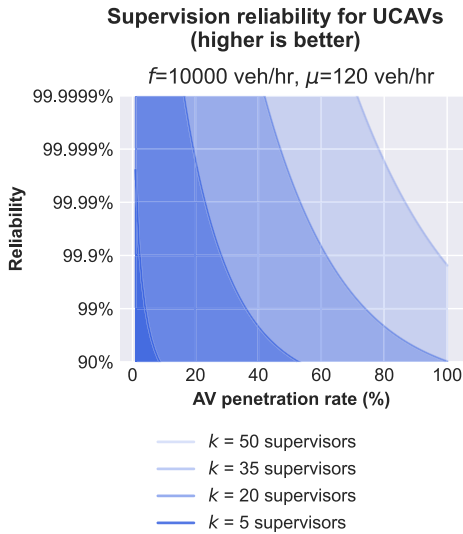


Fig. 4. Visualization of the reliability thresholds in the UCAV case, where the “reliability” metric is calculated as $1 - P_k$. For example, 99.9% reliability means that of 1000 AVs requesting a supervisor due to reachability conditions being violated, 1—on average—will not receive it. This does not necessarily mean a crash occurs, but that one is possible and no human supervisor is available to monitor the situation. As AV penetration increases, the system’s ability to maintain adequate supervision drops monotonically. Note the logarithmic y-axis.

UCAVs cannot communicate with other vehicles in the system, merging AVs must adopt the same conservative approach applied to HVs. Indeed, from the merging AV’s perspective, HVs and UCAVs are functionally indistinguishable.

Applying a union bound to the joint probability event in (6), we have

$$\begin{aligned}
 P(C_{in}) &\leq \sum_{i=1}^R P(\{H_i \leq d_i(v_{i,0}, a_{i,\max}, t)\}) \\
 &= \sum_{i=1}^R \frac{d_i(v_{i,0}, a_{i,\max}, t)}{c} \quad (7)
 \end{aligned}$$

where, the last equality comes from our assumption that each vehicle’s location is distributed uniformly at random along the length of the ring, i.e., $H_i \sim U([0, c])$. This uniform distribution is a product of the infinite uniform roadway that the ring road simulates [54]; even if traffic occurs, it is equally likely to do so at each point on the road. Furthermore, because this work focuses on merging into free-flowing highways, nonuniform congestion is beyond our scope. However, note this does not require each vehicle’s location be independent of other vehicles’ positions; this is investigated further in the following. ■

The uniform distribution assumption previously may be satisfied both by a) a situation in which traffic is flowing freely around the ring, and b) a setting with stop-and-go traffic where the congestion is equally likely to occur at any point within the ring. A case not covered by the lemma is one in which congestion routinely occurs around the merge point, but this case is beyond the current scope of work given our focus on safety in high-speed merges.

C. Joint In-Ring and On-Ramp Reachability for Mixed HVs and UCAVs

Lemma 2: Given a single-lane ring road of circumference c with R (HV or UCAVs) distributed uniformly at random (but not necessarily independently) along the length of the ring, and a merging vehicle q distributed uniformly at random along an on-ramp of length $l_{\text{on-ramp}}$, the joint probability that the on-ramp’s merge point into the ring road is reachable over time horizon t by both an in-ring vehicle and the on-ramp vehicle is bounded above as follows:

$$\begin{aligned}
 P(C) &= P(C_{in}) \times P(C_{out}) \\
 &\leq \frac{\sum_{i=1}^R d_i(v_{i,0}, a_{i,\max}, t)}{c} \times \frac{d_q(v_{q,0}, a_{q,\max}, t)}{l_{\text{on-ramp}}} \quad (8)
 \end{aligned}$$

where, $d_i(v_{i,0}, a_{i,\max}, t)$ is the reachable range for vehicle i over the time horizon t .

Proof: Without loss of generality, we assume a sufficiently long on-ramp, i.e., $l_{\text{on-ramp}} \geq d_i(v_{i,0}, a_{i,\max}, t)$.

The abovementioned follows from analyzing the joint probability of independent events. The rightmost term $P(\text{reach}_{\text{on-ramp}})$ is the probability that an arbitrary point on the on-ramp is within the reachable zone of merging vehicle q over horizon t , and follows similar logic to that used for writing $P(C_{\text{in}})$. ■

Consequently, this also relies upon the assumption that the likelihood of finding the merging vehicle in a given position is evenly distributed across the on-ramp. This assumption is more tenuous in this case than it was in Lemma 1 because in situations of interest (such as when traffic exists on the highway) the AV may have to yield. This would cause the AV to spend a disproportionately large amount of time just before the merge point. A more thorough on-ramp analysis would also have to account for the interaction between AV behavior and HV responses; this is an interesting direction for future work.

D. In-Ring Reachability for Mixed HVs and NCAVs

Now consider how a mixed autonomy system with NCAVs may improve the situation. NCAVs may communicate their near-term trajectories but are not assumed to alter their trajectories to avoid triggering supervision in the system. That is, they do not actively accommodate merging vehicles. The benefit of this connectivity is that in-ring NCAVs' reachable zones can be reduced from where they *might* be to where they *will* be, along with a desired safety buffer. This can have a substantial impact, as the kinematics-based reachability accounts for a vehicle's maximum possible acceleration and is thus likely to overestimate the true distance. This decreases the probability that they trigger the in-ring reachability condition. Further discussion of this point—with particular emphasis on the NCAVs' ability to predict their trajectories—is provided after the corollary's presentation.

Corollary 1: Given a single-lane ring road of circumference c with $R - S$ HVs and S NCAVs distributed uniformly at random (but not necessarily independent of each other) along the length of the ring, the probability that an arbitrary point on that ring is reachable over time horizon t by a vehicle is bounded above as follows:

$$P(C_{\text{in}}) \leq \frac{\sum_{i=1}^{R-S} d_i(v_{i,0}, a_{i,\max}, t) + \sum_{j=1}^S l_j}{c} \quad (9)$$

where, $d_i(v_{i,0}, a_{i,\max}, t)$ is the reachable zone for vehicle i over the time horizon and l_j is the length of NCAV j . Note that one may define l_j to include a safety buffer.

Proof: Let H_i be the asymmetric distance (following traffic flow) between the i th in-ring HV ($i = 1, \dots, R - S$) and the fixed merge point at the current time $t = 0$. Let \tilde{A}_j be the asymmetric distance between the j th in-ring NCAV ($j = 1, \dots, S$) and the fixed merge point at future time $t = t_f$. Due to the connectivity, we know the near-term trajectory of the NCAVs and hence the location of the in-ring NCAVs up to future time $t = t_f$, so NCAV j triggers supervision if its future location at time $t = t_f$ has distance less than its length l_j (i.e., if it is directly

at the merge point). l_j could be expanded to include a safety buffer. In addition, the corollary's assumption says $H_i \sim U[0, c]$ and $\tilde{A}_j \sim U[0, c]$.

Then we similarly have the event C_{in} as the event that any of the HVs or NCAVs trigger supervision

$$\begin{aligned} C_{\text{in}} &= \left(\bigcup_{i=1}^{R-S} \{\text{HV } i \text{ triggers sup. cond.}\} \right) \\ &\quad \cup \left(\bigcup_{j=1}^S \{\text{NCAV } j \text{ triggers sup. cond.}\} \right) \\ &= \left(\bigcup_{i=1}^{R-S} \{H_i \leq d_i(v_{i,0}, a_{i,\max}, t)\} \right) \\ &\quad \cup \left(\bigcup_{j=1}^S \{\tilde{A}_j \leq l_j\} \right). \end{aligned} \quad (10)$$

Applying a union bound to the joint probability event in (10) results in

$$\begin{aligned} P(C_{\text{in}}) &\leq \sum_{i=1}^{R-S} P(\{H_i \leq d_i(v_{i,0}, a_{i,\max}, t)\}) \\ &\quad + \sum_{j=1}^S P(\{\tilde{A}_j \leq l_j\}) \\ &= \sum_{i=1}^{R-S} \frac{d_i(v_{i,0}, a_{i,\max}, t)}{c} + \sum_{j=1}^S \frac{l_j}{c}. \end{aligned} \quad (11)$$

The t_f term need not explicitly appear in these expressions because an NCAV's location is independent of the merge point, and thus t_f is unnecessary for assessing probabilities (as opposed to individual instances).

It is worth further examining connected AVs' ability to predict and communicate their trajectory. Trajectory planning is a well-studied problem in AVs [65]. However, mixed autonomy settings present unique challenges to trajectory planning—while AVs can plan their own trajectories over a given time horizon in isolation, the human drivers on the road are unpredictable. How can AVs predict their own trajectories when sharing the road with human drivers?

First, note that faster-than-expected HVs in this setting do not pose a problem for the connected AVs' trajectory planning; indeed, if they speed further ahead, the AVs have more space. The instances which might be problematic are those in which an HV quickly slows.

However, even here the problems subside upon further analysis. If an HV sharply brakes far ahead of the merge point, the connected AVs behind it are also far from the merge point, and thus do not pose a collision risk for the merging vehicle. If an HV sharply brakes far after the merge point, no connected AV near the merge point will be substantially affected, especially not immediately.

Thus the instances of concern are further limited to those cases in which an HV rapidly brakes near the merge point when an on-ramp vehicle is about to merge. However, if the HV of concern is just prior to the merge point, note that it will have already triggered the supervisor, so any adjustment of a connected AV's propensity to trigger the supervision condition

is redundant. (Recall that once a supervisor is triggered, it supervises the entire merge.)

The remaining case is that in which an HV has just passed the merge point when it rapidly decelerates. Yet the only way an HV manages to pass the merge point without triggering supervision is if the merge point is beyond the on-ramp vehicle's reachable zone during the entire time that same point is within the HV's reachable zone. Therefore, even in this final case of concern, if the HV behaves erratically, there remains the entirety of the reachability time horizon for the connected AV to respond—and likely significantly more, given that the merging vehicle's reachable zone assumes its maximum acceleration.

Note that in these cases ultimate safety still depends on the ability of the human supervising the merging AV to respond to traffic conditions appropriately (as one would expect for any human driver). The abovementioned analysis simply explains the in-ring NCAVs' ability to reduce the uncertainty in their reachability zones from where they might be to where they will be. This reduces the odds supervision requirements will be triggered while still preserving the time horizon for a supervisor to respond, should one be needed. Of course, as discussed previously, unforeseen events can occur, but that is beyond the scope of this research focusing on the merge event.

E. In-Ring Reachability for Mixed HVs and cooperative (supervision-Aware), Connected AVs (CCAVs): Worst Case

Now consider how a mixed autonomy system with connected AVs that cooperate to avoid triggering supervision (that is, supervision-aware CCAVs) may improve upon the previous cases. Recall that both CCAVs and NCAVs may communicate their near-term trajectories to merging AVs. However, whereas NCAVs do not alter their trajectories to accommodate merging AVs, CCAVs can alter their trajectory to assist AVs to merge without triggering human supervision.

We first consider a worst case improvement. Then, in Sections IV-F and IV-G, we consider the improvement given more typical cases with the CCAVs.

Corollary 2: Given a single-lane ring road of circumference c with $R - S$ HVs and S (supervision-aware) CCAVs distributed uniformly at random along the length of the ring, the probability that an arbitrary point on that ring is reachable over time horizon t by a vehicle is bounded above as follows:

$$P(C_{in}) \leq \frac{\sum_{i=1}^{R-S} d_i(v_{i,0}, a_{i,max}, t)}{c} \quad (12)$$

where, $d_i(v_{i,0}, a_{i,max}, t)$ is the reachable zone for vehicle i over the time horizon.

That is, the previous upper bound in Corollary 1 is improved by dropping the $\sum_{j=1}^S \frac{l_j}{c}$ term.

Proof: Again, let H_i be the distance between the i th in-ring HV ($i = 1, \dots, R - S$) and the fixed merge point at the current time $t = 0$. We assume $H_i \sim U[0, c]$.

Given the fixed merge point and a perfect control of its trajectory during the length \tilde{t} planning interval, so long as \tilde{t} is sufficiently long, there exists a control input (sequence of accelerations over the planning horizon) such that the CCAV is

not at the merge point at $t = t_f$ with probability 1. For example, this could consist of slowing down or speeding up to create a gap for the merging AV. See more discussion of \tilde{t} below. So the event $\{\text{CCAV } j \text{ triggers supervision condition}\} = \emptyset$, i.e., does not occur. Hence, the above upper bound improves from the previous one in Corollary 1 with the $\sum_{j=1}^S \frac{l_j}{c}$ term.

In the worst case, the order of HVs and CCAVs (with respect to the merge point) is adversarially distributed, such that 1) all the HVs are asymmetrically closer than all CCAVs to the merge point (when moving in the forward direction), and 2) all HVs have nonoverlapping reachability zones. In this case the CCAVs cannot influence any HV's behavior at the merge point, so it is possible for any HV to trigger the supervision condition. Furthermore, the HVs' combined reachability zone is achieving its maximum coverage over the ring.

The event that an in-ring vehicle triggers supervision is thus due to the HVs

$$C_{in} = \cup_{i=1}^{R-S} \{\text{HV } i \text{ triggers sup. cond.}\}. \quad (13)$$

The union bound gives

$$P(C_{in}) \leq \frac{\sum_{i=1}^{R-S} d_i(v_{i,0}, a_{i,max}, t)}{c}. \quad (14)$$

■

Note that \tilde{t} is distinct from the reachability time horizon t , and in practice would likely be much smaller. \tilde{t} simply corresponds to the case in which the planning interval is too short for the CCAV to avoid the merge point. For example, if a merging AV appears on the on-ramp when an in-ring, supervision-aware CCAV is just before the merge point and moving quickly, there may not exist sufficient time for the in-ring CCAV to brake—and thus block any HVs behind it from interfering with the merge—before its momentum carries it past the merge point. Note also that \tilde{t} should also account for reasonable response times in vehicles following the CCAV. Indeed, the analysis in the following examples assumes the blocking behavior occurs at a rate sufficient to allow trailing vehicles (both AVs and HVs) to respond safely.

F. In-Ring Reachability for Mixed HVs and CCAVs: Uniform Case

Recognizing that the situation described in Lemma 2 is adversarial, additional improvement in supervision scalability can be obtained with a typical (less adversarial) distribution on A_j and H_j .

In a typical mixed autonomy case, the order of HVs and CCAVs relative to the merge point is interspersed. As the CCAVs are supervision-aware and fully cooperative, the CCAV closest to the merge point may stop to accommodate a merge, and thus any vehicle after that CCAV cannot trigger supervision.

To model this, let us denote the order statistics $A_{(1)} \leq \dots \leq A_{(S)}$. Without loss of generality and for ease of exposition, consider the restriction to $c = 1$ and the shorthand notation $d_i := d_i(v_{i,0}, a_{i,max}, t)$.

We consider three different distribution schemes for the in-ring vehicles, beginning with the uniform case described in Section IV-F and continuing to two nonuniform induced traffic

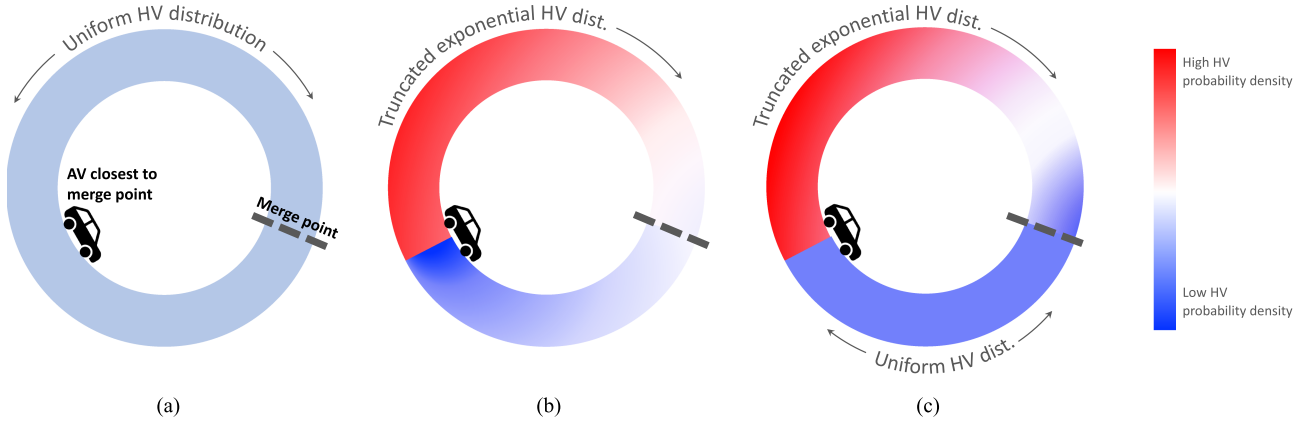


Fig. 5. Examples of three different HV distributions in the mixed HV/CCAV traffic setting: (a) is the case in Example 1, in which HVs are equally likely to be at any point in the ring, regardless of CCAV location(s), (b) considers a case from the induced traffic distribution scenario (see Example 2), in which a CCAV forms a “platoon” of vehicles behind it, such that the distribution of HVs is most dense immediately after the CCAV and thereafter decays smoothly around the ring, (c) corresponds to Example 3. It recognizes that (b) is slightly adversarial: HVs in front of the lead CCAV (between the lead CCAV and the merge point), are more likely to be close to the merge point than far from it. Of course, given asymmetric traffic flows, it is unrealistic for a CCAV to determine the location of HVs in front of it. Therefore, (c) utilizes the exponential distribution for HVs behind the lead CCAV, while allowing a uniform distribution for HVs in front of it. Each subfigure here can be conceptualized as taking a vertical “slice” from its corresponding subfigure in Fig. 6 and wrapping it around the ring road. As such, each diagram here is a visualization of *only one possible configuration* for each distribution; more general representations (i.e., the results if all vertical slices for each example were stacked next to each other) are shown in Fig. 6.

distributions described in Sections IV-G and IV-H. Instances of these three example cases are shown in Fig. 5. In the first case, HVs are equally likely to be at any point in the ring. The second and third cases consider instances in which HVs are more likely to follow closely behind a CCAV.

Theorem 2: Assume CCAV locations are independent from each other (i.e., $A_j \perp A_k$ for all $j, k \in \{1, \dots, S\}$) and all HVs’ and CCAVs’ locations are independent $H_i \perp A_j$. Also assume each vehicle’s location is distributed uniformly at random along the length of the ring, i.e., $A_j \sim U[0, 1]$ for all $j \in \{1, \dots, S\}$ and $H_i \sim U[0, 1]$ for all $i \in \{1, \dots, R - S\}$. Then the probability that an arbitrary point on the ring is reachable over time horizon t by a vehicle is upper bounded as follows:

$$P(C_{\text{in}}) \leq \frac{1 - (1 - d_i)^{S+1}}{S + 1}. \quad (15)$$

Proof: Since $A_j \stackrel{i.i.d.}{\sim} U[0, 1]$ for all $j \in \{1, \dots, S\}$, we have from order statistics $A_{(1)} = \min_j A_j \sim \text{Beta}(1, S)$, and the probability density function for $A_{(1)}$ is $f_{A_{(1)}}(a) = S(1 - a)^{S-1}$, $a \in [0, 1]$, where a denotes the distance of the closest CCAV from the merge point [66]. As $H_i \sim U[0, 1]$, the probability density function for H_i is $f_{H_i}(h) = 1$, $h \in [0, 1]$. By independence, we have the joint probability density function $f_{A_{(1)}, H_i}(a, h) = f_{A_{(1)}}(a)f_{H_i}(h) = S(1 - a)^{S-1}$, $a \in [0, 1]$, $h \in [0, 1]$ (and 0 otherwise).

We can therefore write the following:

$$\begin{aligned} & \{\text{HV } i \text{ triggers sup. cond.}\} \\ &= \{\text{HV } i \text{ triggers sup. cond.} \cap H_i \leq A_{(1)}\} \\ & \cup \{\text{HV } i \text{ triggers sup. cond.} \cap H_i > A_{(1)}\} \quad (16) \end{aligned}$$

where, $P(\{\text{HV } i \text{ triggers supervision condition} \cap H_i > A_{(1)}\}) = 0$ in the current setting because any HV further away

from the merge point than the nearest CCAV will be blocked by that CCAV from triggering the supervisor. Therefore

$$\begin{aligned} & \{\text{HV } i \text{ triggers sup. cond.} \cap H_i \leq A_{(1)}\} \\ &= \{H_i \leq d_i \cap H_i \leq A_{(1)}\}. \quad (17) \end{aligned}$$

One can compute $P(\{H_i \leq d_i \cap H_i \leq A_{(1)}\})$ as

$$\begin{aligned} & P(\{H_i \leq d_i \cap H_i \leq A_{(1)}\}) \\ &= \int_{a=0}^{d_i} \int_{h=0}^a f_{A_{(1)}, H_i}(a, h) dh da \\ &+ \int_{d_i}^1 \int_0^{d_i} f_{A_{(1)}, H_i}(a, h) dh da \\ &= \int_0^{d_i} \int_0^a S(1 - a)^{S-1} dh da \\ &+ \int_{d_i}^1 \int_0^{d_i} S(1 - a)^{S-1} dh da \\ &= \frac{1 - (1 - d_i)^{S+1}}{S + 1}. \quad (18) \end{aligned}$$

A visual representation of this double integration for three examples is provided in Fig. 6. In both the uniform distribution case and the nonuniform “platoon” distribution case (in which a number of HVs follow closely behind a CCAV) *one benefit* of CCAVs is represented by the shaded triangle. It corresponds to the portion of the ring at which in-ring HVs’ reachability zones include the merge point, but which are blocked by an in-ring CCAV, and thus do not pose a danger to the merging vehicle. *An additional benefit* may come in the form of the distribution shift of in-ring vehicles that the CCAV can cause, illustrated in the image with the color gradients, in Fig. 6(b). Here, not only

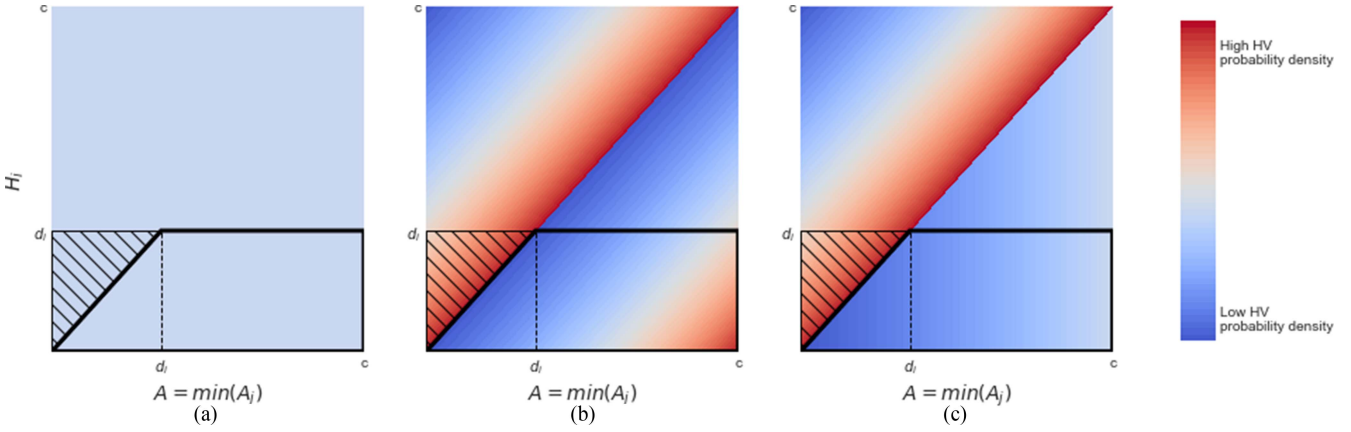


Fig. 6. Visual representation of the HVs' conditional distributions $f_{H_i|A_{(1)}}(h|a)$ for the three examples described in Sections IV-F and IV-G, along with a visualization of the double integration from (18). The two axes represent the random variables H_i and $A_{(1)} := \min A_j$. The bold black border encapsulates the area captured by the double integral, and the shaded triangle represents the portion of the ring for which HVs that would otherwise pose a threat to a merging vehicle will be blocked. The colors represent the varying probability density. (a) Probability in the uniform distribution case (see Example 1). The even tone indicates an equal likelihood of HVs at all ring positions relative to the blocking CCAV. (b) Conditional probability in the nonuniform “platoon” distribution case when HVs follow closely behind the CCAV (see Example 2). The red indicates a higher likelihood of HVs immediately behind the blocking CCAV and the blue indicates a lower likelihood of HVs immediately preceding it. Thus in this second case the CCAV produces a greater supervision scaling effect in the red zone, but is also disadvantageous in that the HVs between the CCAV and the merge point are more likely to be closer to the merge point. (c) HVs' nonuniform conditional distribution, as described in Section 3.

does the CCAV block certain HVs from threatening the merge, but it also shifts the in-ring vehicle distribution such that an HV is less likely to be in the vicinity in the first place. This case will be analyzed further in Example 2.

Fig. 5 can aid in interpreting Fig. 6. Both images can be interpreted geometrically by imagining the colors as representing a third dimension on the image (as if rising out of the page toward the reader). Taking the red to represent a greater volume of probability mass, we see that the shaded triangle in Fig. 6 would also have more probability mass, and thus translate into greater supervision scaling than in the uniform distribution case. At the same time, the distribution in Fig. 6(b) includes an adversarial element: the HVs between the CCAV nearest the merge point and the merge point itself are more likely to be distributed closely to the merge point, which increases the odds that they trigger the in-ring reachability condition. Previously (in the NCAV case), $P(H_i \leq d_i) = d_i$, so the absolute improvement in each term inside the union bound is

$$\begin{aligned} & P(H_i \leq d_i) - P(\{H_i \leq d_i \cap H_i \leq A_{(1)}\}) \\ &= d_i - \frac{1 - (1 - d_i)^{S+1}}{S + 1} \end{aligned} \quad (19)$$

and the relative improvement of each term inside the union bound is

$$\begin{aligned} & \frac{P(H_i \leq d_i) - P(\{H_i \leq d_i \cap H_i \leq A_{(1)}\})}{P(H_i \leq d_i)} \\ &= \frac{d_i - \frac{1 - (1 - d_i)^{S+1}}{S + 1}}{d_i} \\ &= 1 - \frac{1 - (1 - d_i)^{S+1}}{d_i(S + 1)} \end{aligned} \quad (20)$$

where, $d_i \in [0, 1]$ as we normalize $c = 1$. This monotonically increases in d_i and S —that is, the relative improvement increases with larger reachability zones or more CCAVs.

Example 1. (Uniform vehicle distribution): As a numerical example, when $d_i = 0.1$ and $S = 5$, the absolute probability [given via (18)] is 0.078. The previous absolute probability was 0.1, and so the absolute improvement [from (19)] is 0.022, and the relative improvement [from (20)] is 0.219, or 21.9%.

Interpreting this in terms of the reachability bound

$$\begin{aligned} P(C_{in}) &\leq \sum_{i=1}^{S-A} \{\text{HV } i \text{ triggers sup. cond.}\} \\ &= \sum_{i=1}^{S-A} P(\{H_i \leq d_i \cap H_i \leq A_{(1)}\}) \end{aligned} \quad (21)$$

each term inside the sum is roughly 21.9% less than it was previously. Thus, this effectively drops 21.9% of terms from the set $\{1, \dots, S - A\}$ that the sum is taken over, and hence produces an additional 21.9% improvement gain relative to the upper bound.

G. In-Ring Reachability for Mixed HVs and CCAVs: Induced Traffic Cases

Example 2. (An induced traffic vehicle distribution): As the CCAV nearest the merge point X may slow down or stop the HVs behind it, consider the scenario in which HVs may be more likely to be behind—and close to—the first CCAV, recalling Fig. 6(b). For example, imagine a CCAV leading a platoon of HVs.

Accordingly, we model the “wrapped-around” asymmetric distance between the HV following a CCAV using a truncated




Fig. 7. Visualization of the cooperation-based benefits of CCAVs for the situation described in Section 2. The left two plots show the probability that the merge point is within the reachable zone of HV_i in the UCAV and NCAV cases (dotted line) and in the case in which CCAVs are present (solid line, $P(H_i \leq d_i, H_i \leq A_{(1)})$) for two different d values. The dashed line removes the blocking effect for comparison. The CCAVs' relative improvement over the UCAV and NCAV cases is shaded in the two left plots and explicitly represented with corresponding colors in the bar chart. Note these plots consider a single hypothetical HV—not the system's overall probability of triggering supervision. See Appendix A-C for the source data.

Exp(1) distribution with support $[0, 1]$.¹ As such, the probability density function of H_i given $A_{(1)}$ can be written as

$$f_{H_i|A_{(1)}}(h|a) = \begin{cases} \frac{\exp(-(h-a))}{1-\exp(-1)} & \text{if } a \leq h \leq 1 \\ \frac{\exp(-(h+1-a))}{1-\exp(-1)} & \text{if } 0 \leq h < a. \end{cases} \quad (22)$$

Note that the $(h+1-a)$ term occurs when $h < a$, and therefore corresponds to the unlikely case when an HV is closer to the merge point than a CCAV. In this setting, this adjustment is necessary to ensure that when $h = 0$ the probability produced is equal to the $a \leq h$ case when $h = 1$. For a visual explanation see Fig. 6(b): this adjustment is what ensures the probability distribution smoothly “wraps around” from the top of the diagram to the bottom.

For tractable computation assume $f_{A_{(1)}}(a) = \frac{S \exp(-Sa)}{1-\exp(-S)}$ to be the truncated Exp(S) random variable on $[0, 1]$. The resulting joint distribution is

$$\begin{aligned} f_{A_{(1)}, H_i}(a, h) &= f_{H_i|A_{(1)}}(h|a) f_{A_{(1)}}(a) \\ &= \begin{cases} \frac{\exp(-(h-a)) S \exp(-Sa)}{1-\exp(-1) 1-\exp(-S)} & \text{if } a \leq h \leq 1 \\ \frac{\exp(-(h+1-a)) S \exp(-Sa)}{1-\exp(-1) 1-\exp(-S)} & \text{if } 0 \leq h < a. \end{cases} \end{aligned} \quad (23)$$

We can compute the probabilities of interest using the same integration as in Case 1, but with a different joint probability density function. Closed-form probabilities can be obtained as

$$\begin{aligned} P(\{H_i \leq d_i \cap H_i \leq A_{(1)}\}) &= \frac{e^{(S-d_i S)}(e^{d_i S} - e^{d_i S} + S - 1)}{(e-1)(e^S - 1)(S-1)} \\ &\quad + \frac{(e^{d_i} - 1)S(e^{d_i+S} - e^{d_i S+1})}{(e-1)(e^S - 1)(S-1)} \end{aligned} \quad (24)$$

$$\begin{aligned} P(\{H_i \leq d_i \cap H_i > A_{(1)}\}) &= \frac{e^{S+1}(e^{-d_i S} - e^{-d_i S} + S - 1)}{(e-1)(e^S - 1)(S-1)} \end{aligned} \quad (25)$$

$$P(\{H_i \leq d_i\}) = P(\{H_i \leq d_i \cap H_i \leq A_{(1)}\})$$

$$+ P(\{H_i \leq d_i \cap H_i > A_{(1)}\}). \quad (26)$$

As a numerical example, with $d_i = 0.1$, we obtain a $\sim 44\%$ improvement from the previous upper bound when $S = 5$ and a $\sim 68\%$ improvement when $S = 15$.

See Fig. 7 for the numerical results with $d_i = 0.1$ and $d_i = 0.01$, respectively, in a setting with 16 in-ring vehicles total. Although no in-ring HV would actually be present in the 100% AV penetration setting, values are reported for illustrative purposes. The data used for the plots are included in Tables II and III in the appendix. In reality, the true distribution is likely to lie somewhere between the two cases outlined in this section. Note the relative improvement column in the tables is computed by comparing the $P(H_i \leq d_i, H_i \leq A_{(1)})$ column as AV penetration rate increases with $P(H_i \leq d_i) = d_i$, which corresponds to the case in which the probability of a given HV's reachability zone intersecting the merge point is unaffected by the CCAVs. To illustrate the relative impact of the CCAVs' blocking benefit relative to their distribution shift benefit, we ablate the effect of blocking.

Example 3 (A more realistic CCAV case):

H. Supervision Requirements for a Single Rotary

In the abovementioned example we assume both the HVs behind and in front of the first CCAV follow a truncated exponential distribution. Such a distribution is slightly unrealistic (and slightly adversarial) for HVs in front of the first CCAV because it implicitly says that those HVs have a greater probability of being closer to the merge point than far away (that is, closer to the first CCAV). Recall that we define the asymmetric distance following the traffic flow direction. A less adversarial distribution would say the HVs before the first CCAV follow a uniform distribution between that CCAV and the merge point. In fact, here we consider such a joint distribution for HVs and CCAVs, where

¹Specifically, we define $X = \begin{cases} h-a & \text{if } h > a \\ 1-(a-h) & \text{if } 0 \leq h \leq a \end{cases}$ and we assume

$X \sim \text{Exp}(1)$.




Fig. 8. Visualization of the cooperation-based benefits of CCAVs for the situation described in Section 3, similar in format to the abovementioned figure. The dashed line represents the gains (relative to the dotted line) that the in-ring distribution shift incurred by CCAVs produces. The solid line compounds that with the benefit from their ability to sometimes block HVs that would otherwise trigger supervision. See Appendix A-C’s tables for the source data.

the conditional probability density function of H_i given $A_{(1)}$ is

$$f_{H_i|A_{(1)}}(h|a) = \begin{cases} \frac{\exp(-(h-a))}{1-\exp(-1)} & \text{if } a \leq h \leq 1 \\ \frac{e^a-1}{a(e-1)} & \text{if } 0 \leq h < a. \end{cases} \quad (27)$$

where, $\frac{e^a-1}{a(e-1)}$ is set so that the conditional probability $f_{H_i|A_{(1)}}(h|a)$ integrates to 1 for each a . That is, this is a truncated Exp(1) random variable on $[0,1]$ where the random variable is the asymmetric distance from an HV to a CCAV (following the traffic flow) if the HV is behind the first CCAV, but conditionally uniform if the HV is before the first CCAV. A visual representation of this conditional probability density function can be seen in Fig. 6(c).

Again, for tractable computation assume $f_{A_{(1)}}(a) = \frac{\text{Sexp}(-Sa)}{1-\exp(-S)}$ to be the truncated Exp(S) random variable on $[0,1]$. The resulting joint distribution is

$$f_{A_{(1)}, H_i}(a, h) = f_{H_i|A_{(1)}}(h|a)f_{A_{(1)}}(a) = \begin{cases} \frac{\exp(-(h-a)) \text{Sexp}(-Sa)}{1-\exp(-1) 1-\exp(-S)} & \text{if } a \leq h \leq 1 \\ \frac{e^a-1 \text{Sexp}(-Sa)}{a(e-1) 1-\exp(-S)} & \text{if } 0 \leq h < a. \end{cases} \quad (28)$$

The results of this computation are shown in Fig. 8 and the corresponding tables in the appendix. Table IV contains the result when fixing $d_i = 0.1$. Table V contains the corresponding result when fixing $d_i = 0.01$.

We can see that the relative improvements are larger than those in Example 2, as the distribution of HVs before the first CCAV is less adversarial.

We can leverage this work detailing the distribution of in-ring vehicles—and more specifically, the $P(H_i \leq d_i, H_i \leq A_{(1)})$ values, from which we construct $P(C_{in})$ and therefore λ —with (3) to get more direct insight into quantities and trends of interest.

I. Supervision Requirements for a Traffic Network

The analysis to this point has been centered on a single ring road. However, returning to our objective of asking, “How many supervisors are necessary?,” particularly in the setting in which

a team of human supervisors remotely supervise a fleet of AVs, we are interested in a broader traffic network.

Using our kinematics-based reachability distance and the on-ramp density for interstate highways from [67], we find that reasonable values for d range from 0.08 to 0.11 in Washington and California interstate systems, and thus will proceed with $d = 0.1$ for this section of the analysis. However, we observe that the precise value of d to use will vary according to context and desired analysis; still, the trends outlined below persisted for a range of values. The CCAV $P(C_{in})$ values are thus drawn from Table IV. On-ramp flow rates f can range from 0 to $\sim 1,900$ veh/h before single-lane ramps become saturated, and the mean time for a supervisor to monitor a merge is estimated to be between 15 and 60 s, including the time to safely gain situational awareness and assume effective control of the vehicle remotely from the autonomous driver [43], [68].

Thus, we set $f = [7000, 10\ 000, \text{ and } 17\ 000]$ for our investigations and highlight $f = 10\ 000$ for our initial plots. This value can equivalently refer to a network with 10 on-ramps each producing 1000 veh/h, or 100 on-ramps each producing 100 veh/h flowing onto a free-flowing interstate.

Results are presented in Figs. 9 and 10. These trends were found to hold across a variety of configurations. We see that CCAVs enable supervision time that is on the whole inversely related with AV penetration, and that the number of AVs each supervisor can monitor at a given reliability threshold monotonically increases along with the AV penetration rate. The plots also illustrate the extent to which cooperation among the connected AVs boosts reliability in supervision. In particular, note the logarithmic scale on the y-axis in these plots. The “reliability” metric is calculated as $1 - P_k$. This means that at a reliability threshold of 99.9%, 1 in 1000 vehicles—on average—for which the reachability conditions require supervision would not be able to receive supervision (because there is not a supervisor available to immediately provide it). Observe that this does not necessarily mean a crash would occur, given the conservative nature of the reachability constraints. Future work via microsimulation could gather data for analysis of this.

Of particular note is the improvement of the CCAV case over the NCAV case. Fig. 11 illustrates how CCAVs can provide

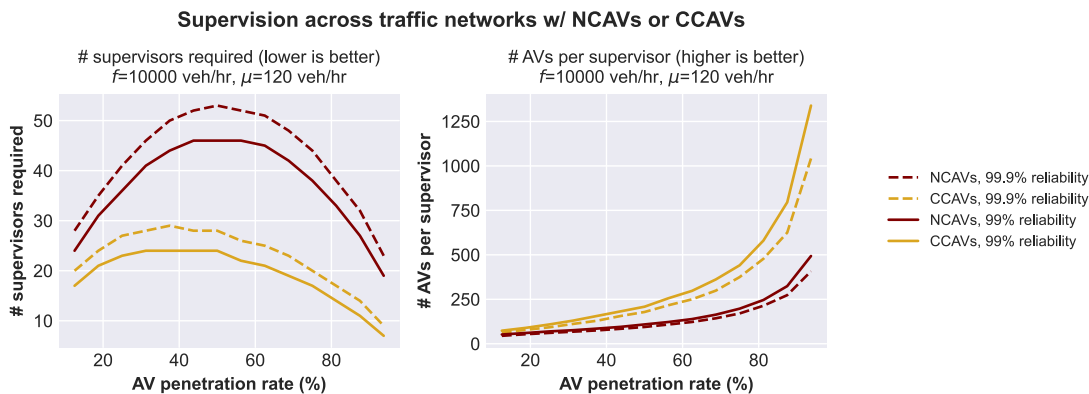


Fig. 9. With connected AVs the number of supervisors necessary to achieve desired high reliability thresholds no longer increases monotonically with rising AV penetration rates, as occurred in the UCAV (baseline) case shown in Fig. 3. The CCAVs substantially outperform the NCAVs. In fact, the CCAV case exhibits an overall trend inversely related with increasing AV penetration. Increasing AV presence boosts the number of AVs requiring supervision at merge time, but also enables AVs to have greater supervision-reducing effects within the ring.

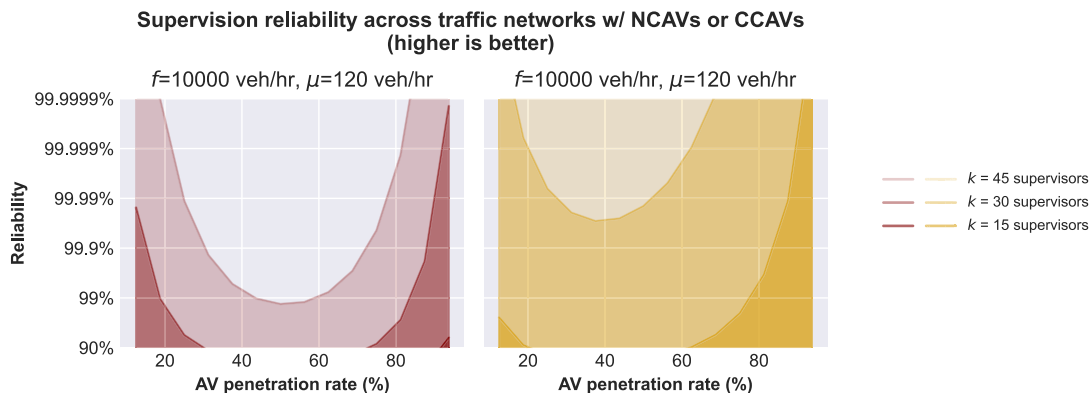


Fig. 10. A visualization of the reliability thresholds NCAVs (left) and CCAVs (right) can achieve at a given task arrival and service rate. Other rates were found to exhibit similar trends. Comparing these plots—while making special note of the logarithmic y-axis—illustrates how CCAVs enable a fixed number of supervisors to achieve a substantially higher reliability threshold.




Fig. 11. Under connected AV cooperation, the operator team successfully covers 99.9999% of cases when the AV penetration rate is 30%. In contrast, without connected AV cooperation, the same team covers just over 90% of cases.

orders of magnitude more reliability without additional human supervisors (note the log scale).

The shape of the various curves can be explained by three interacting effects as the AV penetration rate increases. See Fig. 12. Exerting upward pressure on the number of supervisors required to maintain a given reliability level as AV penetration increases is the growing number of merging AVs that present potential supervision needs. Exerting downward pressure for both the noncooperative and CCAV cases is the fact that as AV penetration increases, a greater proportion of vehicles have the truncated reachability zones that connectivity grants, reducing the average area of the highway that is covered by reachability zones. Finally, for the CCAVs, there exist the benefits discussed in the presentation of the other plots: a) the ability to shift the in-ring vehicles' collective distribution such that each vehicle's reachable zone is more likely to overlap with that of another vehicle, thus leaving more of the ring open to accommodate merges, and b) the ability of CCAVs to block would-be dangerous HVs from encroaching on merging AVs as they join the highway flow. This final effect contributes both to the goal of reducing the

Illustration of interacting effects: reliability achieved with $k=45$ supervisors (higher is better)




Fig. 12. As the AV penetration rate increases, three notable factors affect the system’s supervisability: (1) more AVs to supervise, (2) a greater share of vehicles with truncated reachability zones, and (3) more AVs to assist merges. (1) makes supervision more difficult. (2) and (3) ease supervision. The UCAV curve (in blue) is affected only by 1, as the AVs neither assist merges nor have truncated reachability zones. The NCAV curve (in red) shows what happens with only 1 and 2, since NCAVs do not assist merges. The CCAV curve (in yellow) is affected by 1, 2, and 3. Recall that Effect 3 is the composite of the two benefits described in Figs. 7 and 8.

number of necessary human interventions and reducing the time required for those that must occur. The truncated reachability effect contributes to the goal of reducing the number of human interventions. The curves presented in all four figures can be understood in these terms.

We observe that the “dip” in the reliability curves consistently occurs at a similar point, indicating the mixed autonomy cases with relatively equal proportions of human and AVs provide the most challenging supervision scenarios. Also notable is that the dip shifts leftward in the CCAV setting relative to the NCAV setting. This may have to do with the nonlinear influence of vehicles on others in the system and is an open avenue for future study.

While the values here are representative of a network larger than a single rotary, we found that similar qualitative behavior occurs even on the much smaller case of a single rotary with four on- and off-ramps.

V. IMPLICATIONS FOR AUTONOMOUS DRIVING

It is worth pausing for a moment to appreciate the broader context of these improvements and the leveraging effect that can be achieved in the traffic setting by adopting autonomy with supervision assistance for the driver. The results suggest preventing the need for supervision is more powerful than having more supervisors. Not only does this provide further motivation for investigating scalable supervision of mixed autonomy systems in the specific context of AVs, but it also serves as a case study for determining the potential for supervision scaling to achieve practical gains in a given setting. In so doing, this analysis provides a framework that could be used to evaluate other tasks for scalable supervision.

A. Improved Safety and User Experience

AVs provide a compelling case for investigating scalable supervision because short, particularly dangerous events—such as merging or navigating a construction zone—are interspersed

within longer, easier events—such as cruising a highway. Imagine a future where AVs can safely handle mundane driving like occurs on an interstate highway. An intervention that prevents a human driver from lifting a finger during a five-seconds-long merge may thus not only spare the person interruption for those five seconds, but indeed may have in effect spared her interruption over 20 min if the merge interlinked 10-min stretches on two separate highways.

Certainly, unforeseen events can occur at any point in the driving process, but that is beyond the scope of this research, which focuses on the illustrative and prespecified event of merging. Other instances that present increased risks or unpredictability but can be anticipated include construction zones, school zones, corridors particularly congested with foot traffic, and known traffic accident hotspots [21].

B. Informed Fleet Operations

Beyond the improved driving experience and safety for an individual, in the further future such scaling might be a crucial element enabling the widespread deployment of AVs, and thus crucial for achieving the environmental and social benefits outlined previously. Consider a company that seeks to provide a fully hands-off driving experience. Even if the company’s AVs are generally quite effective at navigating roadways, given the difficulty of some tasks like merging and consumers’ desire for additional safety, the firm may look for additional safety measures. If it were sufficiently economical (and assuming low-latency connections, etc.), one way to do so might be to hire remote supervisors to monitor the AV fleet and provide real-time control when necessary. Thus, the question of interest for this company would not only be reductions in the number of interventions and the time required over some baseline, but how many remote supervisors need to be hired as a whole to supervise the entire fleet of AVs.

One aim of this work is to provide rigorous methodology for informing such estimates by refining the estimated time required for supervision of the AVs. When tailored to the locale

and services of an AV operator, the outcome of the proposed methodology could be incorporated into an estimate for fleet operations.

Based on some moderate assumptions, we might expect that—even if we were to require that a supervisor monitor *each and every* AV merge, regardless of reachability conditions and without any of the scaling techniques used in this work—one of these remote supervisors might be able to supervise between 18 and 66 AVs for the task of highway merging in the United States. The logic behind this estimate is as follows. In the US approximately 33% of total vehicle miles traveled (VMTs) occur on interstates, other freeways, or expressways, where “access and egress points are limited primarily to on- and off-ramps” [69], [70], [71]. Making the conservative assumptions that on these roads the average speed is 50 mph and the average trip length is 15 miles, the total time for each trip is 0.3 h. If the time it takes for a supervisor to assume control and execute a merge task safely is 60 s, and assuming one merge per trip, then a supervisor is needed for 5.55% of every trip. Taking the inverse of this, we see that—if tasks are allocated without gaps or delays—one supervisor could theoretically monitor merges for 18 highway vehicles, even assuming that every merge must be supervised. That is, this value is *without* any of the scalable supervision benefits discussed in this work. Since VMTs on interstates, other freeways, and expressways account for only a third of all VMTs, one supervisor operating in such a manner could supervise up to 52 vehicles (since the majority never make a highway merge). Of course, the scaling potential is substantially greater once we allow for the supervision methods outlined previously.

VI. CONCLUSION

This work explored an operational reliability perspective on improving safety in mixed autonomy settings by investigating human supervision of AVs in a simulated merging task. We investigated the question: Can we do better than the present-day industry standard of persistent supervision? Our findings indicate that clever allocation of human resources for supervision tasks may ease near-term adoption of imperfect autonomous agents in safety-critical environments. Key conclusions include the following.

- 1) AV teaming can enable supervision that is both safe and scalable, with orders-of-magnitude improvement in supervision requirements. This cooperation may enable autonomous systems’ deployment at scale in certain settings.
- 2) Queuing-theoretic analysis and order statistics allow for the establishment of high reliability upper bounds on human supervision requirements in the given setting.
- 3) Scalable supervision provides a method for achieving human-level safety even when autonomous system performance guarantees are elusive.

Future work could use traffic microsimulators to empirically investigate the distribution produced, and could also evaluate the impact of the supervision-aware AVs’ cooperative behavior on traffic flows. These simulations could extend to investigate how learning can be incorporated to exhibit behaviors similar to—

even superior to—those discussed in this work. More generally, continued research in this direction could adopt alternative reachability analysis tools to investigate supervision scaling’s potential in other mixed autonomy settings.

APPENDIX USEFUL STATEMENTS

We provide useful probability statements that we adopt in our proofs. The statements can be found in a standard probability textbook [66].

A. Binomial Distribution Parameterized by a Poisson Random Variable

$$\begin{aligned} \text{Let } N \sim \text{Poisson}(\lambda), X|N \sim \text{Binomial}(N, p) \\ \text{then, } X \sim \text{Poisson}(\lambda p). \end{aligned} \quad (29)$$

B. Poisson Distribution of Summed Independent Random Variables

$$\begin{aligned} \text{Let } X \sim \text{Poisson}(\lambda), Y \sim \text{Poisson}(\mu) \text{ and } X, Y \\ \text{are independent.} \\ \text{Then, } X + Y \sim \text{Poisson}(\lambda + \mu). \end{aligned} \quad (30)$$

C. Tables With Numerical Data for Examples in Section IV-G

See Table II through Table V.

TABLE II
DATA FOR THE CASE DESCRIBED IN SECTION II. RELATIVE REDUCTION IN SUPERVISION TIME FOR S CCAVs RELATIVE TO CONNECTED AV BASELINE $P(H_i \leq D_i) = 0.1$ FOR NON-UNIFORM DISTRIBUTION WHEN $d_i = 0.1$ AND THE TOTAL NUMBER OF IN-RING VEHICLES IS 16. THE COLUMN SECOND FROM THE LEFT IS AN INDICATOR OF THE DISTRIBUTION SHIFT CAUSED BY THE CCAVs. THE COLUMN SECOND FROM THE RIGHT IS THE OVERALL ABSOLUTE SUPERVISION PROBABILITY THAT THE CCAVs TYPICALLY ACHIEVE, AND IS THUS A FUNCTION OF BOTH GAINS VIA THE IMPLICIT DISTRIBUTION SHIFT INDICATED IN THE PRIOR COLUMN AND GAINS VIA THE COOPERATIVE BLOCKING BEHAVIOR ITSELF. THE RIGHTMOST COLUMN INDICATES THE RELATIVE % IMPROVEMENT OVER THE UPPER BOUND $d_i = 0.1$

S	$P(H_i \leq d_i)$	$P(H_i \leq d_i, H_i \leq A_{(1)})$	Relative improvement
2	0.0914	0.0749	25.12%
3	0.0894	0.0675	32.46%
4	0.0888	0.0614	38.60%
5	0.0891	0.0564	43.64%
6	0.0902	0.0522	47.80%
7	0.0917	0.0487	51.29%
8	0.0934	0.0457	54.26%
9	0.0952	0.0432	56.85%
10	0.0970	0.0409	58.13%
11	0.0989	0.0388	61.18%
12	0.1007	0.0370	63.03%
13	0.1025	0.0353	64.72%
14	0.1042	0.0337	66.27%
15	0.1058	0.0323	67.71%
16	0.1074	0.0310	69.03%

TABLE III
DATA FOR THE CASE DESCRIBED IN SECTION II WHEN $d_i = 0.01$

S	$P(H_i \leq d_i)$	$P(H_i \leq d_i, H_i \leq A_{(1)})$	Relative improvement
2	0.0086	0.0084	16.00%
3	0.0081	0.0078	21.86%
4	0.0077	0.0074	26.43%
5	0.0074	0.0070	29.90%
6	0.0072	0.0067	32.52%
7	0.0071	0.0065	34.53%
8	0.0070	0.0064	36.12%
9	0.0069	0.0063	37.40%
10	0.0069	0.0062	38.47%
11	0.0069	0.0061	39.38%
12	0.0069	0.0060	40.17%
13	0.0069	0.0059	40.88%
14	0.0069	0.0058	41.51%
15	0.0069	0.0058	42.10%
16	0.0069	0.0057	42.63%

TABLE IV
DATA FOR THE CASE DESCRIBED IN SECTION III WHEN $d_i = 0.1$

S	$P(H_i \leq d_i)$	$P(H_i \leq d_i, H_i \leq A_{(1)})$	Relative improvement
2	0.0805192	0.0639507	36.05%
3	0.0813434	0.0594455	40.55%
4	0.0828135	0.0554341	44.57%
5	0.084711	0.0519237	48.08%
6	0.0868491	0.0488549	51.14%
7	0.0890942	0.0461509	53.85%
8	0.0913594	0.0437418	56.26%
9	0.0935922	0.0415718	58.43%
10	0.0957621	0.0395984	60.40%
11	0.0978524	0.0377901	62.21%
12	0.0998548	0.036123	63.88%
13	0.101766	0.0345786	65.42%
14	0.103586	0.0331427	66.86%
15	0.105317	0.0318034	68.20%
16	0.106961	0.0305513	69.45%

TABLE V
DATA FOR THE CASE DESCRIBED IN SECTION III WHEN $d_i = 0.01$

S	$P(H_i \leq d_i)$	$P(H_i \leq d_i, H_i \leq A_{(1)})$	Relative improvement
2	0.00714992	0.00696878	30.31%
3	0.00694662	0.00670019	33.00%
4	0.00680191	0.00648492	35.15%
5	0.00670619	0.00631588	36.84%
6	0.00664762	0.00618278	38.17%
7	0.00661575	0.00607607	39.23%
8	0.00660267	0.00598828	40.12%
9	0.00660279	0.00591401	40.86%
10	0.00661227	0.00584953	41.50%
11	0.0066285	0.00579226	42.08%
12	0.0066497	0.00574043	42.60%
13	0.0066746	0.00569277	43.07%
14	0.00670233	0.00564842	43.52%
15	0.00673222	0.00560671	43.93%
16	0.00676381	0.00556716	44.33%

ACKNOWLEDGMENT

The authors are grateful to Dr. Eric Horvitz for inspiring our exploration of AV teaming. They would also like to thank Zhongxia Yan for feedback and assistance with network design and Jiaqi Zhang for discussion about and assistance with elements of analysis.

REFERENCES


- [1] C. Wu, *Learning and Optimization for Mixed Autonomy Systems-A Mobility Context*. Berkeley, CA, USA: UC Press, 2018.
- [2] Allianz Partners, "Self-driving cars hit a speedbump; Interest in autonomous vehicle technology slows down," 2018. [Online]. Available: <https://www.prnewswire.com/news-releases/self-driving-cars-hit-a-speedbump-interest-in-autonomous-vehicle-technology-slows-down-300714345.html>
- [3] C. Bila, F. Sivrikaya, M. Khan, and S. Albayrak, "Vehicles of the future: A survey of research on safety issues," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 5, pp. 1046–1065, May 2017.
- [4] A. Nascimento et al., "A systematic literature review about the impact of artificial intelligence on autonomous vehicle safety," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 12, pp. 4928–4946, Dec. 2020.
- [5] N. Arbabzadeh and M. Jafari, "A data-driven approach for driving safety risk prediction using driver behavior and roadway information data," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 2, pp. 446–460, Feb. 2018.
- [6] B. Vanholme, D. Gruyer, B. Lusetti, S. Glaser, and S. Mammar, "Highly automated driving on highways based on legal safety," *IEEE Trans. Intell. Transp. Syst.*, vol. 14, no. 1, pp. 333–347, Mar. 2013.
- [7] W. Saunders, G. Sastry, A. Stuhlmüller, and O. Evans, "Trial without error: Towards safe reinforcement learning via human intervention," in *Proc. 17th Int. Conf. Auton. Agents and Multiagent Syst.*, 2017, pp. 2067–2069.
- [8] J. Y. C. Chen and M. J. Barnes, "Human-agent teaming for multirobot control: A review of human factors issues," *IEEE Trans. Human-Mach. Syst.*, vol. 44, no. 1, pp. 13–29, Feb. 2014.
- [9] A. Ziebinski, R. Cupek, D. Grzechca, and L. Chruszczyk, "Review of advanced driver assistance systems (ADAS)," in *Proc. AIP Conf. Proc.*, 2017, Art. no. 1906.
- [10] U.S. DoT Federal Highway, "Administration highway statistics 2019: Licensed drivers, U.S.," Bureau of Transportation Statistics, 2021. [Online]. Available: <https://www.fhwa.dot.gov/policyinformation/statistics/2021/dv1c.cfm>
- [11] V. D. Hopkin, "Air traffic control," in *Human factors in Aviation*, E. L. Winer and D. C. Nagel, Eds. Cambridge, MA, USA: Academic Press, 1988, ch. 19, pp. 639–663.
- [12] C. Wu, A. R. Kreidieh, K. Parvate, E. Vinitzky, and A. M. Bayen, "Flow: A modular learning framework for mixed autonomy traffic," *IEEE Trans. Robot.*, vol. 38, no. 2, pp. 1270–1286, Apr. 2022.
- [13] M. Barth and K. Boriboonsomsin, "Real-world carbon dioxide impacts of traffic congestion," *Transp. Res. Rec.*, vol. 2058, pp. 163–171, 2008.
- [14] H. Julia, E. Kosmatopoulos, and P. Ioannou, "Collision avoidance analysis for lane changing and merging," *IEEE Trans. Veh. Technol.*, vol. 49, no. 6, pp. 2295–2308, Nov. 2000.
- [15] M. Zhou, X. Qu, and S. Jin, "On the impact of cooperative autonomous vehicles in improving freeway merging: A modified intelligent driver model-based approach," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 6, pp. 1422–1428, Jun. 2017.
- [16] J. Rios-Torres and A. Malikopoulos, "Automated and cooperative vehicle merging at highway on-ramps," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 4, pp. 780–789, Apr. 2017.
- [17] N. Chen, B. van Arem, T. Alkim, and M. Wang, "A hierarchical model-based optimization control approach for cooperative merging by connected automated vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 12, pp. 7712–7725, Dec. 2021.
- [18] J. Rios-Torres and A. Malikopoulos, "A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 5, pp. 1066–1077, May 2017.
- [19] Y. Zheng, B. Ran, X. Qu, J. Zhang, and Y. Lin, "Cooperative lane changing strategies to improve traffic operation and safety nearby freeway off-ramps in a connected and automated vehicles environment," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 11, pp. 4605–4614, Nov. 2020.
- [20] H. Yu, S. Tak, M. Park, and H. Yeo, "Impact of autonomous-vehicle-only lanes in mixed traffic conditions," *Transp. Res. Rec.*, vol. 2673, pp. 430–439, 2019.
- [21] S. Erdogan, I. Yilmaz, T. Baybura, and M. Gullu, "Geographical information systems aided traffic accident analysis system case study: City of Afyonkarahisar," *Accident Anal. Prevention*, vol. 40, no. 1, pp. 174–181, 2008.
- [22] R. Hoque et al., "Fleet-DAGger: Interactive robot fleet learning with scalable human supervision," in *Proc. 6th Conf. Robot Learn.*, 2022, pp. 368–380.

- [23] T. Sheridan, "Human supervisory control of robot systems," in *Proc. IEEE Int. Conf. Robot. Automat.*, 1986, pp. 808–812.
- [24] B. Kirwan, "The role of the controller in the accelerating industry of air traffic management," *Saf. Sci.*, vol. 37 pp. 151–185, 2001.
- [25] J. van den Broek, J. R. Griffioen, and M. van der Drift, "Meaningful human control in autonomous shipping: An overview," *IOP Conf. Ser.: Mater. Sci. Eng.*, vol. 929, no. 2020, Art. no. 012008.
- [26] M. Bellemare et al., "Autonomous navigation of stratospheric balloons using reinforcement learning," *Nature*, vol. 588, pp. 77–82, 2020.
- [27] D. Hambling, *AI Outguns a Human Fighter Pilot*. Amsterdam, The Netherlands: Elsevier, 2020.
- [28] D. Silver et al., "Mastering the game of go without human knowledge," *Nature*, vol. 550, pp. 354–359, 2017.
- [29] E. Bacci and D. Parker, "Probabilistic guarantees for safe deep reinforcement learning," *Int. Conf. Formal Model. Anal. Timed Syst.*, vol. 18, pp. 231–248, 2020.
- [30] M. Bouton, J. Karlsson, A. Nakhaei, K. Fujimura, M. J. Kochenderfer, and J. Tumova, "Reinforcement learning with probabilistic guarantees for autonomous driving," 2019, *arXiv:1904.07189*.
- [31] L. L. Mero, D. Yi, M. Dianati, and A. Mouzakitis, "A survey on imitation learning techniques for end-to-end autonomous vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 9, pp. 14128–14147, Sep. 2022.
- [32] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer, "HG-Dagger: Interactive imitation learning with human experts," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 8077–8083.
- [33] D. S. Drew, "Multi-agent systems for search and rescue applications," *Curr. Robot. Rep.*, vol. 2, pp. 189–200, 2021.
- [34] A. Dahiya, N. Akbarzadeh, A. Mahajan, and S. L. Smith, "Scalable operator allocation for multi-robot assistance: A restless bandit approach," *IEEE Trans. Control Netw. Syst.*, vol. 9, no. 3, pp. 1397–1408, Sep. 2022.
- [35] M. Cummings, S. Bruni, S. Mercier, and P. Mitchell, *Automation Architecture for Single Operator, Multiple UAV Command and Control*. Cambridge, MA, USA: MIT Press, 2007.
- [36] M. Cummings, C. Nehme, J. Crandall, and P. Mitchell, "Predicting operator capacity for supervisory control of multiple UAVs," *Innov. Intell. Mach.-1*, vol. 70, pp. 11–37, 2007.
- [37] B. Kidwell, G. Calhoun, H. Ruff, and R. Parasuraman, "Adaptable and adaptive automation for supervisory control of multiple autonomous vehicles," in *Proc. Hum. Factors Ergonom. Soc. Annu. Meeting*, 2012, vol. 56, pp. 428–432.
- [38] J. Humann and K. Pollard, "Human factors in the scalability of multirobot operation: A review and simulation," in *Proc. IEEE Int. Conf. Syst., Man Cybern.*, 2019, pp. 700–707.
- [39] J. Conesa-Muñoz, M. Soto, P. Santos, and A. Ribeiro, "Distributed multi-level supervision to effectively monitor the operations of a fleet of autonomous vehicles in agricultural tasks," *Sensors*, vol. 15, pp. 5402–5428, 2015.
- [40] D. Olsen and M. Goodrich, "Metrics for evaluating human-robot interactions," in *Proc. PERMIS*, 2003, pp. 4–5.
- [41] D. Olsen Jr and S. Wood, "Fan-out: Measuring human control of multiple robots," in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst.*, 2004, pp. 231–238.
- [42] R. C. Hampshire, S. Bao, W. S. Lasecki, A. Daw, and J. Pender, "Beyond safety drivers: Applying air traffic control principles to support the deployment of driverless vehicles," *PLoS One*, vol. 15, no. 5, 2020, Art. no. e0232837.
- [43] A. Daw, R. C. Hampshire, and J. Pender, "Beyond safety drivers: Staffing a teleoperations system for autonomous vehicles," 2019, *arXiv:1907.12650*.
- [44] G. Swamy, S. Reddy, S. Levine, and A. Dragan, "Scaled autonomy: Enabling human operators to control robot fleets," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 5942–5948.
- [45] H. Ahn, A. Rizzi, and A. Colombo, "Robust supervisors for intersection collision avoidance in the presence of legacy vehicles," *Int. J. Control, Automat. Syst.*, vol. 18, pp. 384–393, 2020.
- [46] J. Rios-Torres and A. A. Malikopoulos, "A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 5, pp. 1066–1077, May 2017.
- [47] M. Althoff, G. Frehse, and A. Girard, "Set propagation techniques for reachability analysis," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 4, pp. 369–395, 2020.
- [48] A. Kurzhanski and P. Varaiya, "Ellipsoidal techniques for reachability analysis," in *Proc. Int. Workshop Hybrid Syst.: Comput. Control*, 2000, pp. 202–214.
- [49] L. Liebenwein, C. Baykal, I. Gilitschenski, S. Karaman, and D. Rus, "Sampling-based approximation algorithms for reachability analysis with provable guarantees," *RSS*, 2018.
- [50] S. Bansal, M. Chen, S. Herbert, and C. Tomlin, "Hamilton-Jacobi reachability: A brief overview and recent advances," in *Proc. IEEE 56th Annu. Conf. Decis. Control*, 2017, pp. 2242–2253.
- [51] G. Bahati, M. Gibson, and A. Bayen, "Multi-adversarial safety analysis for autonomous vehicles," *RSS Workshop Robust Autonomy*, 2020.
- [52] K. Leung et al., "On infusing reachability-based safety assurance within planning frameworks for human–robot vehicle interactions," *Int. J. Robot. Res.*, vol. 39, pp. 1326–1345, 2020.
- [53] M. Johns et al., "Exploring shared control in automated driving," in *Proc. 11th IEEE/ACM Int. Conf. Hum.-Robot Interact.*, 2016, pp. 91–98.
- [54] Y. Sugiyama et al., "Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam," *New J. Phys.*, vol. 10, 2008, Art. no. 033001.
- [55] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Phys. Rev. E*, vol. 62, 2000, Art. no. 1805.
- [56] M. Treiber and A. Kesting, "Traffic flow dynamics," in *Traffic Flow Dynamics: Data, Models And Simulation*, Berlin Heidelberg, Germany: Springer-Verlag, 2013.
- [57] H. Rakha, M. Snare, and F. Dion, "Vehicle dynamics model for estimating maximum light-duty vehicle acceleration levels," *Transp. Res. Rec.*, vol. 1883, pp. 40–49, 2004.
- [58] K. Brown, K. Driggs-Campbell, and M. J. Kochenderfer, "A taxonomy and review of algorithms for modeling and predicting human driver behavior," 2020, *arXiv:2006.08832*.
- [59] V. R. Rengaraju and V. T. Rao, "Vehicle-arrival characteristics at urban uncontrolled intersections," *J. Transp. Eng.*, vol. 121, pp. 317–323, 1995.
- [60] P. Sarla, M. R. Doodipala, and P. Endla, "Performance of vehicle arrival traffic data at fuel station using queuing system," in *Proc. IOP Conf. Ser.: Mater. Sci. Eng.*, 2020, vol. 981, Art. no. 022080.
- [61] P. Mirchandani and N. Zou, "Analytical modeling of ramp meter systems," in *Proc. IEEE Intell. Transp. Syst. Conf.*, 2006, pp. 432–436.
- [62] J. Ding, H. Peng, Y. Zhang, and L. Li, "Penetration effect of connected and automated vehicles on cooperative on-ramp merging," *IET Intell. Transport Syst.*, vol. 14, pp. 56–64, 2020.
- [63] R. Zhang, F. Rossi, and M. Pavone, "Analysis, control, and evaluation of mobility-on-demand systems: A queueing-theoretical approach," *IEEE Trans. Control Netw. Syst.*, vol. 6, no. 1, pp. 115–126, Mar. 2019.
- [64] D. Gross, *Fundamentals of Queueing Theory*. Hoboken, NJ, USA: Wiley, 2008.
- [65] Y. Rasekhipour, A. Khajepour, S. K. Chen, and B. Litkouhi, "A potential field-based model predictive path-planning controller for autonomous road vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 5, pp. 1255–1267, May 2017.
- [66] G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes*. Oxford, England: Oxford Univ. Press, 2020.
- [67] P. Pilko, J. G. Bared, P. K. Edara, and T. Kim, "Safety assessment of interchange spacing on urban freeways," FHWA-HRT-07-031, Office Res., Develop., Technol., US Federal Highway Administration, 2007. [Online]. Available: <https://www.fhwa.dot.gov/publications/research/safety/07031/07031.pdf>
- [68] J. W. Hess, "Capacities and characteristics of ramp-freeway connections," *Highw. Res. Rec.*, vol. 3, no. 27, pp. 69–115, 1963. [Online]. Available: <https://onlinepubs.trb.org/Onlinepubs/hrr/1963/27/27-004.pdf>
- [69] N. McGuckin and A. Fucci, "Summary of travel trends: 2017 national household travel survey," U.S. Department of Transportation, Federal Highway Administration, 2018. [Online]. Available: https://nhts.ornl.gov/assets/2017_nhts_summary_travel_trends.pdf
- [70] Federal Highway Administration, "Average annual PMT, VMT person trips and trip length by trip purpose," U.S. Bureau of Transportation Statistics, 2018. [Online]. Available: <https://www.bts.gov/content/average-annual-pmt-vmt-person-trips-and-trip-length-trip-purpose>
- [71] Federal Highway Administration (FHWA), "2015 status of the nation's highways, bridges, and transit conditions & performance," in *Report to Congress*. Washington, DC, USA: FHWA, 2017.



Cameron Hickert (Member, IEEE) received the B.S. degree in physics from the University of Denver, Denver, CO, USA, in 2016, and the M.E. degree in computational science and engineering from Harvard University, Cambridge, MA, USA, in 2021. He is currently working toward the Ph.D. degree in social and engineering systems with the Massachusetts Institute of Technology, Cambridge, MA, USA.

His research interests include safe autonomy, reinforcement learning, and distributed cyber-physical systems.



Sirui Li (Member, IEEE) received the B.S. degree with majors in computer science and mathematics from Washington University in St. Louis, St. Louis, MO, USA, in 2019. She is currently working toward the Ph.D. degree in social and engineering systems with the Massachusetts Institute of Technology, Cambridge, MA, USA.

Her research interests include the areas of machine learning for combinatorial optimization and control analysis for transportation systems.



Cathy Wu (Member, IEEE) received the B.S. and M.Eng. degrees from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2012 and 2013, respectively, and the Ph.D. degree from the University of California Berkeley, Berkeley, CA, USA, in 2018, all in electrical engineering and computer sciences.

She is an Assistant Professor with MIT in LIDS, CEE, and IDSS. She completed the Postdoc with Microsoft Research, Redmond, WA, USA. Her research interests include the intersection of machine learning, control, and mobility. Her recent research focuses on how learning-enabled methods can better cope with the complexity, diversity, and scale of control and operations in mobility systems. She is interested in developing principled computational tools to enable reliable decision-making in sociotechnical systems.