On efficient cooperative strategies between UAVs and humans in a dynamic environment

Ketan Savla[*] Carl Nehme[†] Tom Temple[‡] Emilio Frazzoli[§]

In this paper, we consider an automation support system aiming at facilitating efficient cooperation between unmanned vehicles (UAVs) and remotely located human operators, in a dynamically-changing environment. In particular, we consider the following problem. A number of UAVs must visit and service (e.g., classify) targets that are generated according to a spatio-temporal Poisson process, uniformly in a convex, compact region of the plane. The vehicles can plan their motion autonomously. However, once a vehicle is in the proximity of a target location, it requires assistance from a human to finish the service at that location. The system objective is to minimize the expected waiting time between the appearance of a target, and the time of completion of its service. We model the performance of the human operators to vary with its utilization factor and propose several coordination strategies for joint target assignments for the humans and the UAVs.

I. Introduction

One of the prototypical missions for Uninhabited Aerial Vehicles (UAVs), e.g., in environmental monitoring, security, or military setting, is wide-area surveillance. Low-altitude UAVs in such a mission must provide coverage of a region and investigate events of interest, possibly with the assistance of a human operator, as they manifest themselves. In particular, we are interested in cases in which close-range information is required on targets detected by high-altitude aircraft, spacecraft, or ground spotters, and the UAVs must proceed to the locations to service the target under direct operator supervision, where servicing of targets can include tasks like gathering on-site information, target classification, or engagement, e.g, see Figure 1.

Variations of problems falling in this class have been studied by numerous researchers recently.[1–3] However, prior work assumes that UAVs are perfectly autonomous and do not

[*]Postdoctoral Associate, Laboratory for Information and Decision Sciences, Massachusetts Institute of Technology, 32-D716 MIT, 77 Massachusetts Avenue, Cambridge, MA 02139. ksavla@mit.edu

[†]Graduate Research Assistant, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139. nehme@mit.edu

[‡]Graduate Research Assistant, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139. ttemple@mit.edu

[§]Associate Professor, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139. AIAA Member. frazzoli@mit.edu

require any assistance from a human operator at any time during the mission. Even though one can foresee UAVs to have completely automated guidance and navigation modules in the near future, the role of human operators will be indispensable when the servicing of targets involves on-site decision making, especially in a setting when a wrong decision can have lethal consequences.

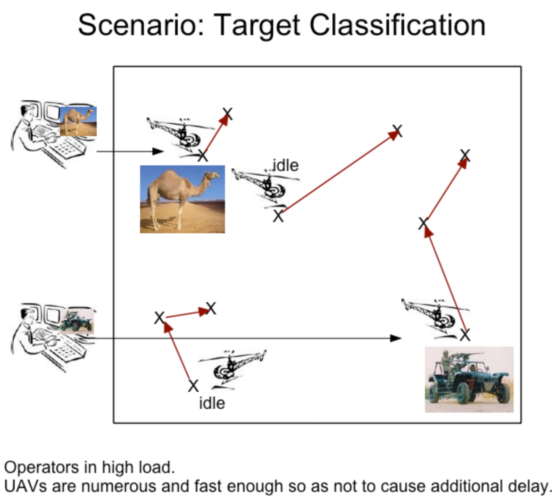

Figure 1. A target classification scenario involving cooperation between human operators and UAVs

The human supervisory control of multiple UAVs has generated a lot of research interest in the recent past (for example, see Cummings *et al.*,[4] and references therein). However, some of the most relevant literature, such as work by Nehme *at al.*,[7] has only examined the capacity of a single operator to supervise multiple unmanned vehicles. In this paper, as illustrated in Figure 1 we propose a novel framework to study vehicle routing problems that involves a heterogeneous team comprised of multiple UAVs and multiple remotely located human operators. We design novel coordination policies among them that maintain a proper balance between the system objective and operator work load. Moreover, we provide an analytical characterization of how the overall system performance depends on the numbers of available human operators and UAVs, respectively.

The contributions of the paper as follows: (i) introduction of a vehicle routing problem which requires active human supervision at target locations along with explicit consideration of load-dependent performance for the human operators; (ii) stability conditions for this coupled system; (iii) design of novel approximation algorithms for the vehicles and the operators for various scenarios.

The rest of the paper is organized as follows. In Section II, we state a few preliminaries and outline the architecture of the support system. In Section III, we model the decision of human operators as a dynamical system and verify it against experimental data. We then study the behavior of this dynamical system model under external input in form of task allocation by the SS in Section IV We design efficient policies for the SS and analyze their performances for asymptotic cases in Sections V and VI. We conclude the paper with Section VII where we state possible avenues for the continuation of this work. The proofs presented in the paper are brief and intended to present only the key ideas. More details will be included in a future extended version of this paper.

II. Preliminaries

In this section, we define preliminary notation and outline the architecture of the automation support system that facilitates cooperation between UAVs and humans.

A. Notation

Let $\mathcal{Q} \subset \mathbb{R}^2$ be a convex, compact domain on the plane, with non-empty interior; we will refer to $\mathcal{Q}$ as the *environment*. Let $\mathcal{A}$ be the area of $\mathcal{Q}$. An external Poisson process generates *targets* over time, which are associated to points in $\mathcal{Q}$, with finite time intensity $\lambda > 0$. Furthermore, the points are sampled from an absolutely continuous spatial distribution described by the density function $\varphi : \mathcal{Q} \to \mathbb{R}_+$. The spatial density function φ is normalized such that $\int_{\mathcal{Q}} \varphi(q)\, dq = 1$. In this paper we concentrate on the case in which the targets are generated uniformly in $\mathcal{Q}$, i.e., $\varphi(q) = 1/\mathcal{A}$.

This process can be thought of as a collection of functions $\{\mathcal{P} : \overline{\mathbb{R}}_+ \to 2^{\mathcal{Q}}$ such that, for any $t > 0$, $\mathcal{P}(t)$ is a random collection of points in $\mathcal{Q}$, representing the targets generated in the time interval $[0, t)$, and such that

- the total number of targets generated in two disjoint time-space regions are *independent* random variables;
- the *expected* number of targets generated in a measurable region $S \subseteq \mathcal{Q}$ during a time interval of length Δt is given by:

$$\mathrm{E}[\mathrm{card}\left((\mathcal{P}(t + \Delta t) - \mathcal{P}(t)) \cap \mathcal{S}\right)] = \lambda \Delta t \cdot \frac{\mathrm{Area}(\mathcal{S})}{\mathcal{A}}.$$

Each target has a service request associated with it. These service requests are identical and are to be fulfilled by a team of m UAVs and remotely located n human operators. The team is clearly heterogeneous, and each class of agent provides a unique capability: the vehicles provide access to on-site information, and the human operators provide the cognitive capabilities necessary to process such information and complete the service. In order to service a request, one of m UAVs needs to travel to the target location and, once at the target location, the UAV needs assistance from a human operator to finish the service. Once an operator starts assisting a UAV, the time until completion is assumed to be purely determined by the cognitive capability of the human operator. The vehicles are assumed to be point masses that move with speed v. Moreover, all the vehicles and are assumed to be identical and to have unlimited capacity. The operators are also assumed to be identical.

B. Automation Support System (SS)

We propose an automation Support System (SS) to facilitate the cooperation between vehicles and the human operators. The overall architecture of the SS is depicted in Figure 2, where we show the interaction between the SS, vehicles and the operators for a particular task of target classification.

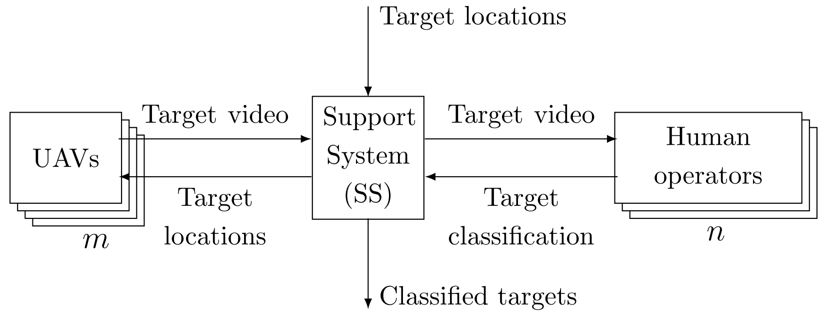

Figure 2. Support System (SS) Architecture

The detailed interaction between the vehicles, human operators and the SS is as follows. The target locations are transmitted to the SS instantaneously with their generation. The SS assigns loitering stations to the vehicles, i.e., locations where the vehicles wait when they do not have any targets alloted to them. During the course of the mission, the SS allocates target locations to each vehicle in batches. Every vehicle deals with the batches in the order that they are received. For each batch of target locations assigned to it, the vehicle determines the shortest tour through them starting with its current location. We assume that each vehicle navigates autonomously along straight line paths between two locations. Hence, the shortest tour for the vehicle over a collection of target locations would be the Euclidean Traveling Salesperson tour over those target locations. Once a vehicle has reached a target location, it issues an *attention request* to the SS, indicating that they need assistance from a human operator. The SS directs these attention requests to human the operators, not necessarily as soon as they are received. Moreover, the SS could also assign *secondary* targets, e.g., health monitoring of the UAVs, etc., to the operators. [a]

A human operator finishes the tasks assigned to it (attention request or secondary task) in the order they are received, picking a task as soon as it is *idle*. When an operator picks up an attention request, it is connected directly to the UAV which issues that request. The operator and the UAV then cooperate to finish the service of that target. Once the service is over, the operator notifies the SS about service completion.

As mentioned before, the UAVs navigate autonomously and do no require assistance from the SS or human operators for the same. In summary, a typical task assignment for the SS consists of managing assignment of loitering locations and target locations to the vehicles, and attention requests and secondary tasks to the operators. In the course of this paper, we would be interested in designing such task assignment policies for the SS.

C. Problem Formulation

Let π denote a generic task assignment policy for the SS. The objective is the design of task assignment policies that allow the team of vehicles and operators to fulfill service requests efficiently. A policy π is said to be *stabilizing* if, under its effect, the expected number of

[a]Not directing the attention requests from the vehicles to the operators instantaneously or assigning secondary tasks to an operator might seem to delay the *primary* task. However, the reason for doing that is related to maintaining proper operating conditions for the operators

outstanding targets does not diverge over time, i.e., if

$$\overline{n}_\pi = \lim_{t\to+\infty} \mathrm{E}\left[n(t): \text{ SS executes policy } \pi\right] < +\infty,$$

where $n(t)$ is the number of outstanding targets at time t. Intuitively, a policy is stabilizing if the SS is able to ensure that targets are serviced at a rate that is, on average, at least as fast as the rate at which new service requests are generated. Let T_j be the time elapsed between the generation of the j-th service request, and the time it is fulfilled. If the system is stable, then the following balance equation (also known as Little's formula[5]) holds:

$$\overline{n}_\pi = \lambda \overline{T}_\pi, \tag{1}$$

where $\overline{T}_\pi := \lim_{j\to+\infty} \mathrm{E}[T_j]$ is the system time under policy π, i.e., the expected time a service request must wait before being fulfilled, given that the SS follows the strategy defined by π. Note that the system time $\overline{T}_\pi$ can be thought of as a measure of the quality of service collectively provided by the SS.

At this point we can finally state our problem: we wish to devise policies that are (i) stabilizing, and (ii) yield a quality of service (i.e., system time) achieving, or approximating, the theoretical optimal performance given by

$$\overline{T}_{\text{opt}} = \inf_{\pi \text{ stabilizing}} \overline{T}_\pi.$$

III. Modeling Human Decision Time

In this section, we model the cognitive capability of the human operators in terms of time required to finish the service once it starts assisting the corresponding UAV; we shall refer to this time as the *decision time* for the human operator. Specifically, we focus our attention on the variation of the decision time under different load conditions. Unlike an inanimate processor or sensor whose performance level can be assumed to be more or less indifferent to its past history, it is intuitive to expect the performance of a human operator to depend on its past utilization. We study the following question: what is the decision time $s(t)$ at time t for a task, given the history of utilization of that human operator up to time t? A solution to this question is of great importance in determining efficient target allocation strategies for the human operators by the SS.

Inspired by the Yerkes-Dodson law[6] and recent work by Nehme *et al.*,[7] we seek to model the relationship between $s(t)$ and $u(t)$; unlike earlier work, we propose an approach based on a dynamical system approach. The Yerkes-Dodson law is a hypothetical relationship implying that, as depicted in Figure 3, the human performance on a particular cognitive task improves with increases in the mental arousal state up to a certain point, beyond which it starts decreasing. There is no universal definition of mental arousal rate and its quantification depends on the problem at hand.

For our problem domain, we take inspiration from the work by Donmez *et al.*,[8] where the authors analyze data from experiments with human participants in a setting similar to the one in this paper. The experiment was based on a software tool developed by the Humans and Automation Laboratory at the Massachusetts Institute of Technology, allowing human operators to supervise a team of unmanned vehicles in a simulated operational environment.

Human participants were engaged in a number of supervisory tasks, ranging from waypoint selection to target assignment, and visual classification. A more detailed description can be found in Donmez *et al.*[8]

On the basis of the data from these experiments, the authors report the existence of a U-shaped relationship between the level of *situational awareness* of an operator and its *utilization factor*, reminiscent, although in a different context, of the Yerkes-Dodson relationship. The situational awareness was measured in terms of the reaction time of an operator to respond to a threat area once it intersected with the path of a vehicle. The utilization factor was calculated as the proportion of time the operator actively interacted with the display (e.g., adding a way point, engaging in a visual target, etc.) during the course of the experiment. The utilization factor was computed as the percentage of busy time over fixed intervals of 150 seconds.

Our first objective in this paper is to identify a model that is suitable for on-line operator scheduling; the model in Donmez *et al.* is non-causal and is better suited for post-mission analysis of the data.In the following analysis, we concentrate on the data corresponding to classification tasks in the experiment;[8] we would like to remark that our analysis is *not* meant to be a full investigation of human performance in such tasks, but rather as a justification of our proposed model on the available experimental data. In fact, the available data refer to an experiment designed to investigate the situational awareness of a single operator, required to perform a variety of task; extensions to multiple-operators and/or single-task scenarios (as the ones considered in this paper) are not supported by the experimental data but must be understood as an extrapolation of the available results. The validity of such extrapolation needs to be investigated specifically. For example, our model does not capture the effects of operator fatigue or other factors. Extensions to multiple operators, visual tasks, and other scenarios are currently being pursued by the Humans and Automation Laboratory at MIT.

Let $u : \mathbb{R} \to \{0, 1\}$ be a indicator function defined so that $u(t)$ is 1 if an operator was busy at time t and 0 otherwise. We seek to identify the relationship between the history of the operator activity, i.e., $\{u(\theta) : \theta \in [-\infty, t)\}$ and the decision time $s(t)$ for the operator if he/she were assigned a task at time t.

We consider the utilization factor as the output of y a dynamical system, driven by the input $u : \mathbb{R} \to \{0, 1\}$. In particular, we consider a first-order linear lag, of the form

$$\begin{aligned} \frac{d\,x(t)}{d\,t} &= \frac{u(t) - x(t)}{\tau}, \qquad x(0) = x_0, \\ y(t) &= x(t), \end{aligned} \tag{2}$$

where τ is a time constant that determines the extent to which past experiences affect the utilization factor for the human operator, x is the state of the system, and x_0 is an initial condition summarizing the effect of the input for $t < 0$.

Notice that, given the definition of u as an indicator function for operator activity, as long as $x_0 \in [0, 1]$, then $y(t) \in [0, 1]$, for all $t \geq 0$. Moreover, assuming that the time constant τ is much larger than the average decision time for a single task, the output of the system y approximates to the long-term average of u.

For a given utilization factor y, the decision time is modeled as a random variable $s(y)$. It is of interest to compute the statistical properties of such random variable; in particular, we are interested in modeling its expected value $\bar{s}(y)$, which we will assume to be a continuous

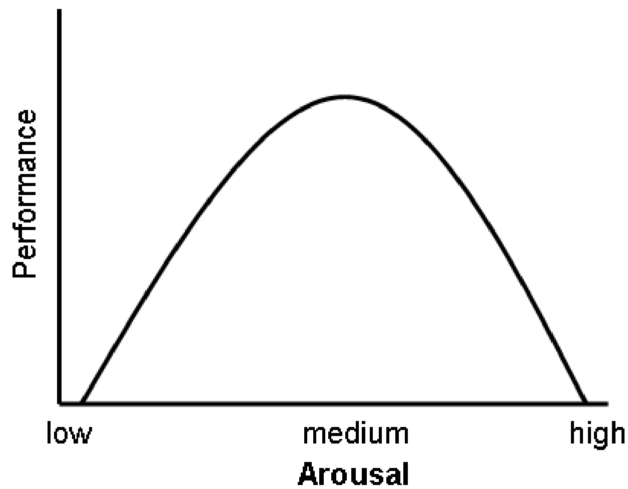

Figure 3. The Yerkes-Dodson Law Graph

function $\bar{s} : [0,1] \to \mathbb{R}_+$ of the utilization factor y. In order to study $s(y)$, we group the data samples—consisting of pairs (u_i, s_i), where u_i is the utilization factor at the inception of the i-th task, and s_i is the corresponding decision time—into bins corresponding to 5% intervals for the utilization factor. Interestingly, the distribution of decision times in each bin is very well modeled by a log-normal distribution, i.e., the logarithms of the decision times are normally distributed, see Figure 4. Moreover, the standard deviation of the logarithms of the decision time is almost constant in the interval of interest for the utilization factor.

The average of the decision times within each bin is computed, and associated to the midpoint of the bin. Finally, a weighted least-square procedure is used to find a quadratic empirical approximation to the function $\bar{s}(y)$. Weights are proportional to the number of samples in each bin. Our model in Equation (2), with the choice $\tau = 300$ s and $x_0 = 0.5$ yielded a U-shaped relationship between the decision time and the utilization factor; moreover, this relationship is well described by a quadratic law. The bin averages along with the quadratic approximation are depicted in Figure 4. The resulting quadratic approximation for the average decision time is

$$\bar{s}(y) = c_2 y^2 + c_1 y + c_0 = (229y^2 - 267y + 99.0) \text{ s}. \tag{3}$$

The minimum average decision time is evaluated as $s^* = 21$ s, achieved for $y = 0.583$.

We remark again that this analysis is not meant to be an conclusive quantitative investigation for human performance in classification tasks, since the experiments were designed with other purposes in mind. On the other hand, the good agreement between the experimental data and our model provides some evidence justifying our modeling choices for the dynamic utilization factor (2) and the quadratic approximation (3) of the average decision time.

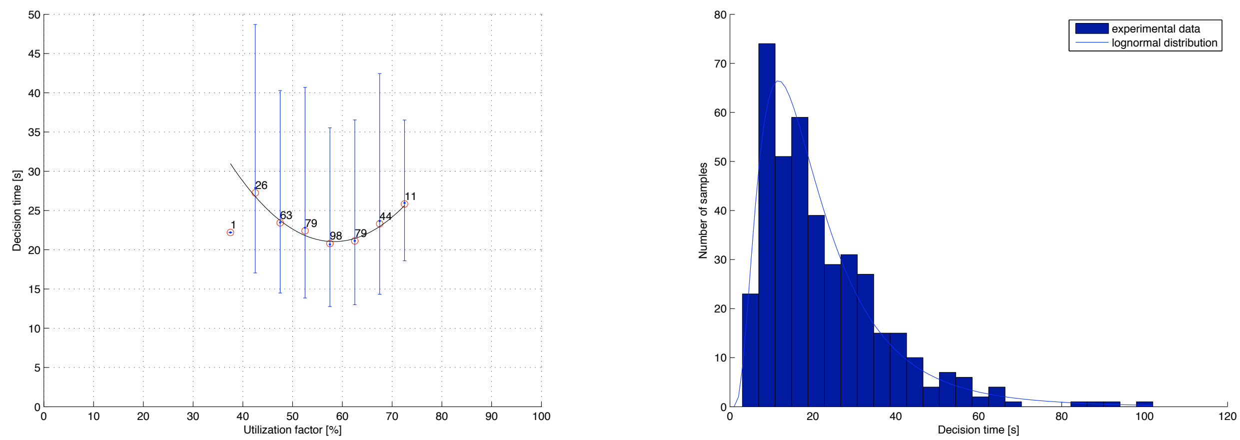

Figure 4. Left: Plot of decision times in classification tasks: circles represent the average of the data in each 5% utilization factor bin, and the solid line is the weighted least-squares quadratic approximant; the numbers indicate the number of samples in each bin; the error bars show 1-σ confidence intervals (notice that the outlier corresponds to a bin with a single sample). Right: the distribution of decision time is well modeled by a log-normal distribution.

IV. Behavior of the human operator as a dynamic server—the deterministic case

In the previous section, we developed a model to predict the decision time for a human operator based on his/her utilization history. The utilization factor at any instant is a function of the task allocation strategy of the SS for the human operator as well as the decision time on the tasks undertaken in the past. Hence, in order to understand efficient task allocation strategies by the SS for the human operator, one needs to understand the relation of the decision time as a function of task allocation strategy by the SS.

In this section, we start by studying this relationship for a simple task allocation strategy by the SS: for operator i, $i \in \{1, \ldots, n\}$, the SS assigns tasks to operator i at a rate λ_i. Furthermore, each operator is responsible only for a *single class of tasks*. One can deduce, from stability requirements, that $\sum_{i=1}^{n} \lambda_i = \lambda$. The utilization profile of an operator, however, depends not only on the rate of arrival of tasks, but also on the inter-arrival times of the tasks. In general, the inter-arrival times of tasks from the SS depends, among other things, on the inter-arrival time of attention requests from the vehicles, which is a complicated process. Hence, we start by analyzing the simple case when the inter-arrival times of task from SS is constant.

We focus our analysis on a single operator i. For brevity in notation, we drop the subscript from λ_i, i.e., we study the dynamics of decision time for a human operator when the inter-arrival time of tasks is $1/\lambda$. We further make the simplifying assumption that, for every y, the service time $s(y)$ is a deterministic quantity $s(y) = \bar{s}(y)$, and $\bar{s}(y)$ is a quadratic function:

$$s(y) = \bar{s}(y) = c_2 y^2 + c_1 y + c_0,$$

where c_0, c_1 and c_2 are constants such that s is strictly convex and positive on the interval $[0, 1]$ and has a minimum in the interval $(0, 1)$. These last conditions ensure that $\bar{s}(y)$ is of the form discussed in Section III.

Let us begin our analysis with the case in which the operator has no backlog; in this case, upon the arrival of a new task, the operator is able to service that task with no queueing delay. Let us indicate with $y_i = x_i$ the utilization factor of the operator upon arrival of the i-th task. The utilization factor upon arrival of the $(i+1)$-th task is then evaluated by integration of (2) over the time period $[0, 1/\lambda]$, with initial condition $x_0 = x_i$ (recall that the system is time invariant, so we can choose arbitrarily the origin of the time axis). The utilization factor when the operator has completed service of the i-th task, at time $t' = \bar{s}(x_i)$ is

$$x(t') = 1 - (1 - x_i)e^{-t'/\tau}.$$

Consequently, we get

$$x_{i+1} = x(t')e^{-(1/\lambda - t')/\tau},$$

and finally

$$x_{i+1} = (1 - (1 - x_i)e^{-\bar{s}(x_i)/\tau})e^{(\bar{s}(x_i) - 1/\lambda)/\tau} = \left(x_i - 1 + e^{\bar{s}(x_i)/\tau}\right) e^{-\frac{1}{\lambda\tau}}.$$

Equilibrium conditions are attained when $x_{i+1} = x_i$, assuming $\bar{s}(x_i) < 1/\lambda$ (so that if the queue was empty at the arrival of the i-th task, it will also be empty at the arrival of the $(i+1)$-th task). In other words, equilibrium conditions are found when

$$x = \left(x - 1 + e^{\bar{s}(x)/\tau}\right) e^{-\frac{1}{\lambda\tau}},$$

i.e., when

$$\bar{s}(x) = \tau \log\left(1 - (1 - e^{\frac{1}{\lambda\tau}})x\right).$$

Let us define

$$r(x) := \tau \log\left(1 - (1 - e^{\frac{1}{\lambda\tau}})x\right);$$

notice that $0 = r(0) \leq r(x) \leq r(1) = 1/\lambda$, for all $x \in [0,1]$; furthermore, when $\tau \to \infty$, $r(x) \approx x/\lambda$.

Equilibrium points can be found as the solutions of the equation $\bar{s}(x) = r(x)$, i.e., as the intersections between the curves $\bar{s}$ and r. Moreover, x_{i+1} will be no greater than x_i when $\bar{s}(x_i) > r(x_i)$, and viceversa.

This lets us draw a first conclusion: a sufficient condition for queue instability is the following. If there exists x_{d} such that

$$\bar{s}(x) \geq \frac{1}{\lambda}, \qquad \forall x \geq x_{\mathrm{d}},$$

then the queue is unstable, for all initial conditions $x_0 \geq x_{\mathrm{d}}$. In other words, should the operator utilization factor exceed x_{d}, the utilization factor will increase until it reaches 100%, the queue size will diverge, since the operator will not be able to service tasks at a rate that is higher than the rate at which they are generated.

Moreover, if $\bar{s}(x) \geq r(x)$ for all $x \in [0,1]$, then the human operator will not be able to stabilize the queue, for any initial condition. In this case, $x_{i+1} \geq x_i$, with $x_{i+1} = x_i$ if and only if $x_i = 1$: the utilization factor will increase until it becomes greater than x_{d}, (and eventually reach 100%), at which point the queue size will diverge. This condition can be stated as a bound on the maximum task generation rate that the human operator can stabilize: let us indicate with $\lambda_{\max}$ the smallest λ such that $\bar{s}(x) \geq r(x)$ for all $x \in [0,1]$.

On the other hand, if the equation $\bar{s}(x) = r(x)$ admits only one solution, then it must be the case that $\bar{s}(1) \leq r(1) = 1/\lambda$. In this case, the human operator will be able to stabilize the queue, for all initial conditions. In other words, if the queue is empty, it will stay so; if it is not, it will be eventually emptied.

A more complex case is that in which the equation $\bar{s}(x) = r(x)$ admits more than one solution. We will concentrate on the case in which the operator's time constant is large with respect to the task inter-arrival time, i.e., $\tau >> 1/\lambda$. In this case,

$$x_{i+1} = \left(x_i - 1 + e^{\bar{s}(x_i)/\tau}\right) e^{-\frac{1}{\lambda\tau}} \approx \left(1 - \frac{1}{\lambda\tau}\right) x_i + \frac{\bar{s}(x)}{\tau}$$

and there are at most two equilibrium points, corresponding to the roots of the quadratic equation $\bar{s}(x) = x/\lambda$. Let us indicate the two equilibria as x_1 and x_2, with $x_1 \leq x_2$. Under the stated convexity assumptions for $\bar{s}$, it can be recognized that $x_{i+1} > x_i$ for all $x_i < x_1$, and all $x_i > x_2$, whereas $x_{i+1} < x_i$ for $x_1 < x_i < x_2$. Furthermore, x_2 is an unstable equilibrium point, and x_1 is a stable equilibrium. In other words, for τ large enough, the queue is unstable for any initial conditions such that $x_0 \geq x_2$, and stable otherwise, with the utilization factor converging asymptotically to x_1.

Summarizing, we have the following cases for a *deterministic* queueing system, where the service time depends on the utilization factor of the server (the human operator):

(i) If $\lambda < 1/\bar{s}(1)$, then the queue is always stable.

(ii) If $\lambda \geq \lambda_{\max}$ then the queue is always unstable.

(iii) If $1/\bar{s}(1) \leq \lambda < \lambda_{\max}$, and $\tau\lambda >> 1$, then the queue is conditionally stable. In particular, all initial conditions that are greater than both equilibrium points lead to instability. For all initial condition smaller than at least one equilibrium point the utilization factor converges to the smallest equilibrium.

See Figure 5 for an illustration of the above cases. For the model studied in the previous section (3), we would get

- **[Light load]** $1/\bar{s}(1) = 1/60\ \text{s}^{-1}$: task arrival rates smaller than about one task per minute per operator are always stabilizable.
- **[Instability]** $\lambda_{\max} = 1/33.9\ \text{s}^{-1}$: task arrival rates greater than about 1.8 tasks per minute per operator cannot be stabilized.
- **[Minimum operator delays]** $\lambda^* = 1/36\ \text{s}^{-1}$: minimal operator delays are in principle obtained when the task generation rate is about 1.7 tasks per minute per operator, corresponding to a stable equilibrium condition with utilization factor of about 58%. Notice that the domain of attraction of such equilibrium extends only to an utilization factor of about 74%.

While these results are only valid in the deterministic case, they provide some useful guidelines for sizing a supervisory team, given the rate of arrival of external tasks. Stochastic models of the proposed human-operator queueing system will be the subject of future work.

Our discussion so far has focused on studying the performance of a human operator with respect to its utilization. Let $\bar{s}^*$ and $\sigma^2_{s^*}$ be the mean and the variance, respectively

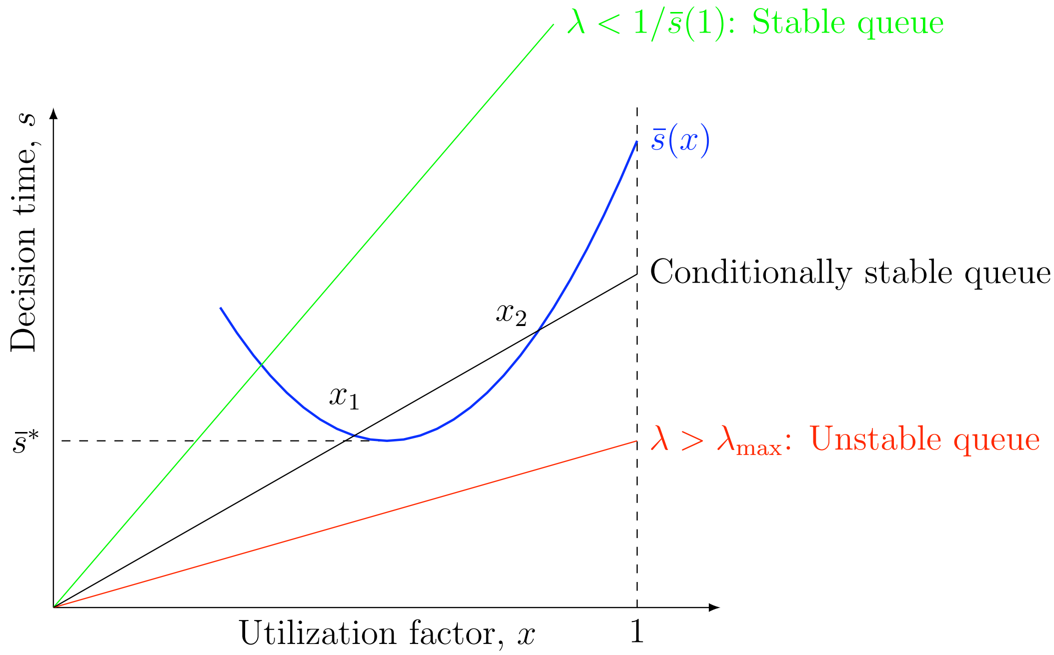

Figure 5. Stability of the deterministic dynamic queueing model

of the decision time for the visual task for an operator under *ideal* conditions. However, in the general case, it is not always possible for the SS to make sure that the operators are working under ideal conditions while making sure that the waiting time for a task does not grow unbounded. Without dealing with this issue in detail, we assume that the mean and the variance of the actual decision time for an operator, as guaranteed by the SS for all $\lambda \leq \lambda_{\max}$, are upper bounded by $\bar{s}$ and σ_s^2 respectively, and that assignment policies for the human operators by the SS are in place. In the next few sections, we now focus on the task assignment policies for the vehicles and analyze the performance of those policies. For analysis purposes, we shall concentrate our attention on a few asymptotic cases; the reason being that these scenarios lend themselves to succinct and meaningful analysis. Note, however, that the policies themselves can be implemented in all scenarios.

V. The light load case

We start by studying the problem for the case when the rate of generation of targets is very small, i.e., when $\lambda \to 0^+$.

The Weber function

In order to study the problem in this case, we need to introduce a problem from geometric optimization. Given the set $\mathcal{Q} \subset \mathbb{R}^2$ and a set of points $p = \{p_1, p_2, \ldots, p_m\} \in \mathcal{Q}^m$, the *Voronoi* cell of an agent i, denoted as $\mathcal{V}_i$, is defined to be the set of points that are closest to an agent i than any other agent, i.e.,

$$\mathcal{V}_i(p) = \{q \in \mathcal{Q} \mid \|p_i - q\| \le \|p_j - q\| \quad \forall j \in \{1, \dots, n\}\}.$$

The expected distance between a random point q, sampled from a uniform distribution over $\mathcal{Q}$, and the closest point in p is given by

$$\mathcal{H}_m(p, \mathcal{Q}) := \int_{\mathcal{Q}} \frac{1}{\mathrm{A}} \min_{i \in \{1, \dots, m\}} \|p_i - q\| \, dq = \sum_{i=1}^{m} \int_{\mathcal{V}_i(p)} \frac{1}{\mathrm{A}} \|p_i - q\| \, dq. \tag{4}$$

The function $\mathcal{H}_m$ is known in the locational optimization literature as the *continuous multi-median function*; see[9,10] and references therein.

The m-median of the set $\mathcal{Q}$ is the global minimizer

$$p_m^*(\mathcal{Q}) = \underset{p \in \mathcal{Q}^m}{\operatorname{argmin}} \, \mathcal{H}_m(p, \mathcal{Q}).$$

We let $\mathcal{H}_m^*(\mathcal{Q}) = \mathcal{H}_m(p_m^*(\mathcal{Q}), \mathcal{Q})$ be the global minimum of $\mathcal{H}_m$. It is straightforward to show that the map $p \mapsto \mathcal{H}_1(p, \mathcal{Q})$ is differentiable and strictly convex on $\mathcal{Q}$. Therefore, it is a simple computational task to compute $p_1^*(\mathcal{Q})$. It is convenient to refer to $p_1^*(\mathcal{Q})$ as the median of $\mathcal{Q}$. On the other hand, the map $P \mapsto \mathcal{H}_m(P, \mathcal{Q})$ with $m > 1$ is differentiable (whenever $(p_1, \dots, p_m)$ are distinct) but not convex, thus making the solution of the continuous m-median problem hard in the general case. In fact, it is known[9,11] that the discrete version of the m-median problem is NP-hard for $d \ge 2$. However, gradient algorithms for the continuous m-median problems can be designed[12] by means of the equality

$$\frac{\partial \mathcal{H}_m(p, \mathcal{Q})}{\partial p_i} = \int_{\mathcal{V}_i(p)} \frac{1}{\mathrm{A}} \frac{p_i - q}{\|p_i - q\|} \, dq. \tag{5}$$

We will not pursue the issue of computation of the m-median and of the corresponding $\mathcal{H}_m^*(\mathcal{Q})$ further, but will assume that these values are available.

Lower bound

The following lemma gives the lower bound on the optimum value of the system time, $\overline{T}_{\text{opt}}$.

Lemma V.1. *For given $\mathcal{Q}$, λ, m and n, the optimum system time satisfies the following lower bound.*

$$\overline{T}_{\text{opt}} \ge \frac{\mathcal{H}_m^*(\mathcal{Q})}{v} + \bar{s}^*.$$

Proof. Let us consider the j-th target, generated at time t_j, at a random position $e_j \in \mathcal{Q}$. The time necessary to service the j-th target is at least the sum of the minimum time taken by the closest UAV to move from its location at time t_j to the target's position e_j and the on-site service time spent by the vehicle under the direct supervision of a human operator. The location of the vehicle and the load factor for the human operators at time t_j is in general unknown, since it depends on the chosen control policy, and on the history of generated targets. If we assumed that the vehicle is always in such a location that it minimizes the *a priori* expected Euclidean distance to a randomly-generated target, we would get a lower bound on the travel time for the vehicle to reach the target location: such a point is the median of the set $\mathcal{Q}$. Similarly, if we assumed that the operators are always under *ideal* working conditions it minimizes the on-site time spent at the target location. In other words,

$$\overline{T}_{\mathrm{opt}} \geq \min_{p \in \mathcal{Q}^m} \int_{\mathcal{Q}} \frac{1}{\mathcal{A}} \min_{i \in \{1,\dots,m\}} \frac{\|p_i - q\|}{v}\, dq + \bar{s}^* = \frac{\mathcal{H}_m^*(\mathcal{Q})}{v} + \bar{s}^*.$$

□

Lemma V.1 holds for any policy, and any value of λ. However, it is most useful in the light-load case and as such it is reported in this section. We now propose a policy and then analyze its performance.

The Median Based (MB) Assignment Policy

The loitering locations for the vehicles are the m median locations for the region $\mathcal{Q}$. The SS forms batches of one target each and assigns it to the vehicle whose loitering location is closest to the target location.

Let $\overline{T}_{\mathrm{MB}}$ be the system time obtained by implementing the median based assignment policy. We now analyze the performance of the policy in the following theorem.

Theorem V.2. *For all m and n,*

$$\overline{T}_{MB} \leq \frac{\mathcal{H}_m^*(\mathcal{Q})}{v} + \bar{s}, \;\; \text{as } \lambda \to 0^+.$$

Proof. Suppose at time $t = 0$, there are n_0 outstanding targets. For a sufficiently large time, the probability that all of these targets have been serviced is close to one. Also, the probability that no new target appears during that time is also high. Hence after a sufficiently long time, all the vehicles will be positioned at their respective median locations. Therefore, the expected distance that a vehicle will have to travel to the location of a new target would be $\mathcal{H}_m^*(\mathcal{Q})/v$. This, when added to the mean decision time by the human operator, completes the proof.

□

As a consequence, we can conclude that, when the target arrival rate is sufficiently small, the system time does not depend on the number of human operators, but decreases as the number of available UAVs increases. More precisely, the system time decreases as $m^{-1/2}$; the efficiency gains are due to the better geometric coverage of the environment of interest achievable with more vehicles. Clearly, in such situations, it is beneficial to use more UAVs than human operators.

VI. The heavy load case

We now consider the case when the targets are generated frequently so that the vehicles and/or human operators are always busy, i.e., λ is large. We start with a new lower bound.

Lemma VI.1. *Let $\rho_v^* = \lambda\bar{s}^*/n$. Then there exists a constant $\gamma \geq 2/(3\sqrt{2}\pi) \approx 0.266$ such that*

$$\overline{T}_{\mathrm{opt}} \geq \gamma^2 \frac{\lambda\mathcal{A}}{m^2 v^2 (1-\rho_v^*)^2} - \frac{\bar{s}^*(1-2\rho_v^*)}{2\rho_v^*}.$$

Proof. A lower bound on the system time is obtained by assuming that a human operator is always available to process information when required and that the human operators are always working under ideal work load conditions. We then use the proof for the lower bound for the system time for the DTRP from Bertsimas *et. al.*[13] □

We now propose a policy which is inspired largely by the DTRP policies from Bertsimas *et. al.*[13] and that performs particularly well when one has the a large team of human operators, specifically when n is large and $\lambda\bar{s^*}/m \to 1^-$.

The Equipartition Based (EB) Assignment Policy

Select a location $\bar{q}$ in $\mathcal{Q}$ arbitrarily and let this be the loitering location for all the vehicles. For some fixed integer $l \geq 1$, divide $\mathcal{Q}$ into l sub-regions of equal area using radial cuts centered at $\bar{q}$ (i.e., form l wedges of area $\mathcal{A}/l$). As the targets are generated, form batches of k/l targets each, with targets in each batch coming from the same sub-region. Once a batch is formed, deposit them in a queue. The batches are assigned to the first available vehicle in a First Come First Serve (FCFS) basis. Optimize over k.

Let $\overline{T}_{\mathrm{EB}}$ be the system time obtained by implementing the equipartition based assignment policy. We now analyze the performance of the policy in the following theorem.

Theorem VI.2. *Let $\rho_v^* = \lambda\bar{s^*}/m$. Then,*

$$\overline{T}_{EB} \leq \frac{\beta^2}{2}\frac{\lambda\mathcal{A}}{m^2v^2(1-\rho_v^*)^2} - \frac{\beta^2}{2\gamma^2}\frac{\bar{s^*}(1-2\rho_v^*)}{2\rho_v^*} \quad \text{as } \rho_v^* \to 1^- \text{ and } n \to +\infty,$$

where $\beta \approx 0.72$ is the Euclidean TSP constant.

The interpretation of this result is the following. Assume that there are no constraints on the number of operators that can be used. Then, it is possible to select a number m^* of operators in such a way that the rate at which each operator is assigned targets is $\lambda_i = \lambda/m^* \approx \lambda^*$, as close as possible to the rate yielding minimum operator delays. In this case, there are no advantages in using more human operators than m^* at the same time. On the other hand, adding more autonomous vehicles will reduce the system time in an inverse quadratic fashion. This, when compared to the light load case, implies that the system time is more sensitive to the number of UAVs when the UAVs are under heavy load.

We now consider the case when the operators are under heavy load, i.e., when λ is the largest *sustainable* target arrival rate by the SS for given number of human operators.

We now state the result on the performance of the stochastic median queue policy under heavy load conditions for operators and light load conditions for vehicles.

Theorem VI.3. *If m and n are such that $m >> n$ then,*

$$\overline{T}_{SM} \leq \frac{\mathcal{H}_m^*(\mathcal{Q})}{v} + \frac{\lambda\sigma_s^2}{n^2(1-\rho_h)} + \bar{s}, \quad \text{as } \lambda \to \lambda_{max}.$$

In this case, human operators are the main bottleneck in the system. While performance can be improved both by adding human operators and by adding vehicles, the effect of increasing the former is much more pronounced than the effect of adding the latter.

VII. Conclusions and future work

As the scenario of multiple operators managing a fleet of UAVs become possible, we hope that this work provides preliminary guidelines towards establishing a sound basis for the design and management of such a system. We investigated how the utilization history of an operator affects its performance at cognitive supervisory tasks. We also give bounds that can be used to estimate the appropriate relative numbers of UAVs and operators.

As an immediate future plan, we plan to extend our analysis to a stochastic setting, in order to account properly for non-deterministic queueing effects. We also plan to study conditions under which the Support System can effectively use secondary tasks to derive optimal performance from human performance when the rate of generation of targets is very low. We propose to use discrete event simulation to numerically compute the system performance as a function of λ, n and m under our proposed policy and a variety of human performance curves. We will then compare these results to the theoretical bounds. This will establish the range in which the high and low load limits are applicable. This will also establish the conditions under which our lower bounds are tighter or looser. Finally, the numerical results will help illuminate the parameters that dominate performance in the intermediate cases.

An extension that we would like to explore would be to include a second type of task that is more discretionary. For instance, a task that gives the operators *context* which can have the affect of raising their situational awareness, or another task that that indirectly improves the system performance in some way but could decrease their situational awareness, for example, a maintenance task.

Overall, we believe that the analysis of systems with humans in and on the loop, based on mathematical models both of human cognitive processes and of autonomous vehicle performance will provide system designers with powerful tool enabling the *a priori* characterization of mixed-initiative systems, before requiring validation with human subjects.

Acknowledgments

This work was supported in part by the Michigan/AFRL Collaborative Center on Control Science, AFOSR grant no. FA 8650-07-2-3744. The research for the experimental data used in this paper was performed at the MIT Humans and Automation Laboratory, under the direction of Prof. M. L. Cummings, and was supported by ONR and Charles River Analytics, in the context of the Plan Understanding for Mixed-Initiative Control of Autonomous Systems (PUMA) project. We are grateful to Prof. Cummings for graciously making the data available to us, and for providing valuable feedback on the paper. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of collaborators and of the supporting organizations.

References

[1] J. J. Enright, E. Frazzoli, K. Savla, and F. Bullo, "On multiple UAV routing with stochastic targets: performance bounds and algorithms," in *AIAA Conf. on Guidance, Navigation and Control*, Aug. 2005. Electronic Proceedings.

[2] K. Savla, E. Frazzoli, and F. Bullo, "Traveling Salesperson Problems for the Dubins vehicle," *IEEE*

Transactions on Automatic Control, vol. 53, no. 9, 2008. (Submitted Jun 2006) To appear.

[3]W. Li and C. Cassandras, "A cooperative receding horizon controller for multivehicle uncertain environments," *IEEE Transactions on Automatic Control*, vol. 51, no. 2, pp. 242–257, 2006.

[4]M. L. Cummings, S. Bruni, S. Mercier, and P. J. Mitchell, "Automation architecture for single operator, multiple uav command and control," *The International C2 Journal*, vol. 1, no. 2, pp. 1–24, 2007.

[5]J. Little, "A proof of the queueing formula $l = \lambda w$," *Operations Research*, vol. 9, pp. 383–387, 1961.

[6]R. M. Yerkes and J. D. Dodson, "The relation of strength of stimulus to rapidity of habit-formation," *Journal of Comparative Neurology and Psychology*, vol. 18, pp. 459–482, 1908.

[7]C. E. Nehme, J. W. Crandall, and M. L. Cumming, "Using discrete-event simulation to model situational awareness of unmanned-vehicle operators," in *Virginia Modeling, Analysis and Simulation Center Capstone Conference*, (Norfolk, VA), 2008. To appear.

[8]B. Donmez, C. E. Nehme, M. L. Cummings, and P. de Jong, "Modeling situational awareness in a multiple unmanned vehicle supervisory control discrete event simulation," *Journal of Cognitive Engineering and Decision Making*, 2008. Submitted.

[9]P. K. Agarwal and M. Sharir, "Efficient algorithms for geometric optimization," *ACM Computing Surveys*, vol. 30, no. 4, pp. 412–458, 1998.

[10]Z. Drezner, ed., *Facility Location: A Survey of Applications and Methods*. Springer Series in Operations Research, New York: Springer Verlag, 1995.

[11]N. Megiddo and K. J. Supowit, "On the complexity of some common geometric location problems," *SIAM Journal on Computing*, vol. 13, no. 1, pp. 182–196, 1984.

[12]J. Cortés, S. Martínez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.

[13]D. J. Bertsimas and G. J. van Ryzin, "Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles," *Operations Research*, vol. 41, no. 1, pp. 60–76, 1993.