A Dynamic Pickup and Delivery Problem in Mobile Networks Under Information Constraints

Holly A. Waisanen, Devavrat Shah, and Munther A. Dahleh, *Fellow, IEEE*

Abstract—This paper considers a network in which a set of vehicles is responsible for picking up and delivering messages that arrive according to a Poisson process. Message pickup and delivery locations are uniformly distributed in a convex region. The vehicles are required to pickup and deliver the messages so that the average delay is minimized. It is required that the vehicle that picks up a message must be the one to deliver it. This problem is called the dynamic pickup and delivery problem (DPDP) and has applications in the context of autonomous vehicles and wireless ad hoc networks.

The control policies considered are separable into two parts: an assignment policy used by a centralized controller to assign arriving messages to the vehicles for service and a service policy used by each vehicle to determine the service routes through its assigned messages.

Lower bounds are provided on the delay achievable by separable control policies that depend on the information constraints in place. It is proved that the optimal average delay scaling can be reduced when message destination information is available to the centralized controller in addition to the message source information. The paper also provides policies that achieve these scaling bounds, proving that these bounds are tight.

Index Terms—Dial-a-ride problem (DARP), dynamic pickup and delivery problem (DPDP), dynamic traveling repair-person problem (DTRP), less-than-truckload (LTL), unmanned aerial vehicle (UAV).

I. INTRODUCTION

WE consider the dynamic pickup and delivery problem (DPDP), in which a set of vehicles is responsible for picking up and delivering messages that arrive at different pickup locations at different times with different delivery locations. The goal is to find control policies that determine the routes of the vehicles such that the average time to first pickup and then deliver a message is minimized. In this paper, we will require that a single message must be picked up and delivered by the same vehicle. Because information about the pickup location, delivery location and arrival time of a particular message is not available until after it has arrived, this vehicle routing problem is dynamic. The control policies produce routes for vehicles as function of the arrived demands and their

H. A. Waisanen is with the Boston Consulting Group, Dallas, TX 75201 USA (e-mail: waisanh@alum.mit.edu).

D. Shah and M. A. Dahleh are with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: devavrat@mit.edu; dahleh@mit.edu).

pickup and delivery locations with the objective to minimize the average message time in system over an infinite horizon.

A classical problem in dynamic vehicle routing is the dynamic traveling repair-person problem (DTRP), which was extensively studied by Bertsimas and van Ryzin [1]–[3]. Unlike the DPDP, each demand in the DTRP corresponds to a single location to be visited. The DTRP can be thought of as a special case of DPDP by identifying the delivery location to be the same as the pick-up location. Though the DTRP is relatively simpler, we find the analysis methods developed in the above mentioned papers to be quite useful in obtaining our results. For this reason, the known results on DTRP will be described in more detail in Section II-B. We note that recent surveys [4] and [5] provide state-of-the-art in the context of dynamic vehicle routing problem.

The DPDP problem arises in many important applications. For example, consider a scenario where people telephone a cab-service exchange to request a ride. The cab-service exchange is to decide which cab picks up (and delivers) which person at what time. This problem is also known as dial-a-ride problem (DARP). Other applications include courier services, manufacturing and inventory routing, less-than-truckload (LTL) trucking, emergency services, mobile sensor networks, and unmanned aerial vehicle (UAV) routing. Most of the previous work on the pickup and delivery problem deal with static setup or periodic re-optimization over a receding horizon. We refer an interested reader to surveys [5], [6] and [7] for a detailed account of previous work. In the case that demands are packets of bits and wireless communication capability is added to the vehicles, the DPDP may also be applied to suggest control methods for mobile, multi-agent wireless networks. This will be a subject of our future work.

In this paper, our interest is two-fold: a) Obtain a lower bound on the performance of any control policy under a general stochastic setup; and b) Obtain a control policy that provides such optimal performance. Of particular interest is the quantification of the performance of the network as a function of several scaling parameters, including the number of vehicles, the total arrival rate of messages, the required service times, and to a lesser extent, the vehicle velocity and network area. Similar analysis exists for the single-stage Dynamic Traveling Repair-person Problem and a two-stage unit-capacity Dynamic Pick-up and Delivery Problem [8]. However, the methods used to derive these results are not sufficient for the multi-stage, multi-vehicle problems we are interested in. Our results for the infinite-capacity two-stage DPDP are the first of their kind. Besides analyzing system performance as a function of the scaling parameters, we also examine the impact of several other system qualities, including information structure and service type. Our re-

sults provide general methods which are different than those in the existing literature. A salient feature of our optimal control policy is the non-intuitive partitioning of space for pickup and delivery that comes out of the lower bound analysis.

A. Model

1) Vehicles and Messages: Let there be n vehicles in a geographic area $\mathcal{A} \subset \mathbb{R}^2$, which is a convex, compact set with volume A. For simplicity, we consider $\mathcal{A} = [0, \sqrt{A}]^2$, with the understanding that these results may be extended to other convex environments with the same area. Throughout this paper, regions will be labeled in calligraphic script, i.e., $\mathcal{A}$, and areas will be labeled with italics, i.e., A. Each vehicle may move in any direction at any time with a velocity of magnitude $\leq v$.

Messages are generated according to a Poisson process with time intensity $\lambda(n)$. We will express $\lambda = \Omega(f(n))$[1] for some given scaling function $f(n)$. The precise required scaling of $\lambda(n)$ will be stated in the theorems. For ease of notation, at various places in the paper we will use λ for $\lambda(n)$.

Associated with each message j are source and destination locations denoted by $s(j) \in \mathcal{A}$ and $d(j) \in \mathcal{A}$ respectively. Source locations are independently and identically distributed (IID) in $\mathcal{A}$ according to the distribution density $\phi_S : \mathcal{A} \to \mathbb{R}^+$. Similarly, destination locations are IID with density $\phi_D : \mathcal{A} \to \mathbb{R}^+$. In this paper, we assume that both source and destination locations have uniform distribution on $[0, \sqrt{A}]^2$, that is $\phi_S(\zeta) = \phi_D(\zeta) = 1/A, \forall \zeta \in \mathcal{A}$.

The messages need to be picked up from their source locations and delivered to their destination locations by the vehicles. A message is picked up (delivered) when a vehicle spends a fixed on-site service time of $\bar{s}(n)$ at the source (delivery) location to pick up (deliver) the message.

Note that $\bar{s}(n)$ is a fixed constant, but is expressed as a function of n to emphasize the connection between the arrival rate $\lambda(n)$, the number of servers n, and the maximum onsite service time that may be supported in a stable system. To see this, compare the system to an M/D/n queue, where arrival process is Poisson of rate $\lambda(n)$, service time is $2\bar{s}(n)$ (i.e., only the time spent in onsite service) and n unit rate servers that can serve the packets in the queue. The average utilization for this system is $\rho = \lambda(n)2\bar{s}(n)/n$, the product of the arrival rate and the service time per message divided by the number of vehicles that serve these messages. A necessary condition for the stability of this system, by classical queuing theory is $\rho < 1$ which is equivalent to $\bar{s}(n) < n/2\lambda(n)$. Therefore, the maximum onsite service time $\bar{s}(n)$ supportable by a stable system is implicitly a function of n and $\lambda(n)$. As stated in our results, in fact $\rho < 1$ is sufficient for the system to be stable (or have finite number of unserved requests in the system on average). Note that this sufficiency in independent of v as long as $v > 0$.

We make the requirement that the vehicle that picks up a message must be the one that delivers it. That is, messages may not be transferred between vehicles after they have been picked up.

---

[1]Recall the following notation: 1)$f(n) = O(g(n))$ means that $\exists$ a constant c and integer N such that $f(n) \leq cg(n), \forall n \geq N$. 2) $f(n) = \Omega(g(n))$ if $g(n) = O(f(n))$. 3) $f(n) = \Theta(g(n))$ means that $f(n) = O(g(n))$ and $g(n) = O(f(n))$.

Furthermore, we assume that each vehicle can carry an unlimited number of messages at any time.

2) Control Policies: A control policy, π, is a set of decision making rules that decides the pickup and delivery schedule of arriving messages, based on a set of constraints on the information available to the vehicle. In this paper, we consider policies $\pi = (\pi_A, \pi_S)$ that can be decomposed into two components, assignment and service. An assignment policy, π_A, describes how a centralized controller assigns arriving messages to vehicles on a real-time basis. A service policy, π_S, describes how each vehicle performs the pickup and delivery of its assigned messages. In this paper we will focus on characterizing optimal assignment policies. Given an assignment policy, we then draw on existing tools in vehicle routing to lower bound the delay incurred under any service policy with the given message assignment policy. We assume that neither the vehicles nor the centralized assignment controller have any knowledge of individual messages before they arrive although the overall message arrival process and source and destination distributions are known.

In particular, we limit our attention to time-invariant and spatially-based assignment policies where π_A is described by a collection of scaled densities $\{p_i(x, y)\}_{i=1}^n$ with the following property:

$$\sum_{i=1}^n p_i(x, y) = \phi_S(x)\phi_D(y) = \frac{1}{A^2}, \forall x, y \in \mathcal{A}. \qquad (1)$$

Let $p_{i,s}(x) = \int_{\mathcal{A}} p_i(x, y)dy$ and $p_{i,d}(y) = \int_{\mathcal{A}} p_i(x, y)dx$. The precise operational meaning of $\{p_i(x, y)\}_{i=1}^n$ is defined below.

We further restrict the set of assignment policies according to the information available to the controller in making message assignments. In particular, we consider two types of information structure:

Source Only Information: When a message arrives, its source location is known to the centralized controller, but vehicles do not know the destination of messages until they pick them up. When a message arrives at location x, the centralized controller randomly assigns the message to one of the vehicles, with each assignment occuring with probability $P(\text{message is assigned to vehicle } i) = p_{i,s}(x)/\phi_S(x)$. Each assignment is made independently of all previous assignments.

Because destination information may not be exploited in making message assignments, the density of destination locations served by each vehicle must be the same as the overall density of destinations, that is, $p_{i,d}(y) = 1/A, \forall y, \forall i$. Since the source and destination locations are independent, $p_i(x, y)$ has the form

$$p_i(x, y) = p_{i,s}\phi_D(y) = p_{i,s}\frac{1}{A}, \forall x, y \in A, \forall i. \qquad (2)$$

Let Π_{SO} denote the set of all policies that satisfy the assignment properties in (1) and (2) above and use Source Only information in making message assignments.

Source-Destination Information: When a message arrives, both its source and also its destination location are known to the centralized controller. These densities are used to make the message assignments in the following way. When a message arrives at location x that is destined for location y, the centralized

controller randomly assigns the message to one of the vehicles, with each assignment occuring with the following probability:

$$P(\text{message is assigned to vehicle } i) = \frac{p_i(x,y)}{\phi_S(x)\phi_D(y)}.$$

Each assignment is made independently of all previous assignments. Under the Source and Destination information structure, destination information may be used to shape the destination density and therefore (1) remains the only restriction on the assignment policy.

Let Π_{SD} denote the set of all policies satisfying this assignment property and using only information available in the Source and Destination information structure.

Our results will show that that the performance of optimal control policies is significantly affected by the particular information structure in place.

Given the Poisson arrival process and the random message assignments with the above probabilities, by the Poisson splitting property the assignment process to each vehicle is an independent Poisson process with smaller rate. Precisely, independent of service policy, the messages arrive for vehicle i under a fixed assignment policy π_A according to an independent Poisson process with rate

$$\lambda_i(\pi) = \lambda_i(\pi_A) = \lambda \int_A \int_A p_i(x,y)dxdy. \qquad (3)$$

Combining (1) and (3) implies that for any valid single vehicle assignment density $p_i(x,y)$

$$\sum_{i=1}^n \lambda_i(\pi) = \lambda. \qquad (4)$$

3) Performance Metrics: The delay of message j, denoted $W(j)$, is defined to be the elapsed time between the message's arrival to the system and its delivery to its destination location. This includes any time the message waits to be picked up, the onsite service time for pickup, travel time on the vehicle before arriving at the delivery location, and finally onsite service time for delivery. The quantity W is defined to be

$$W = \limsup_{j \to \infty} E[W(j)]. \qquad (5)$$

If $W(\cdot)$ has unique stationary distribution with finite mean, then the limsup in the above definition is replaced by lim. We say that the system is stable if $W < \infty$. Recalling the stability discussion in Section I-A1, a necessary condition for the existence of a stable policy is $\rho = \lambda(n)2\bar{s}(n)/n < 1$. If π is a stable policy, then $W(\pi)$ denote the associated asymptotic delay W as defined above. Note that $W(\pi) = W(\pi_A, \pi_S)$ is a function of both the assignment policy and the service policy.

We may also define the single vehicle equivalent of the W. First, let i_j denote the j^{th} message assigned to vehicle i. Then

$$W_i = \limsup_{j \to \infty} E[W(i_j)]. \qquad (6)$$

If no messages are served by vehicle i, then $W_i = 0$. If π is such that the system has unique invariant (stationary) distribution, all $W_i(\pi)$ are finite and the corresponding limits exist then the expected delay over all messages

$$W(\pi) = \sum_{i=1}^n \frac{\lambda_i(\pi)}{\lambda} W_i(\pi). \qquad (7)$$

Note that $\sum_{i=1}^n \lambda_i = \lambda$ for any valid assignment policy. The system is then stable if $W_i < \infty, \forall i$. For a single vehicle with message arrival rate λ_i, a necessary condition for the existence of a stable policy is $\rho_i = \lambda_i 2\bar{s}(n) < 1$. Therefore, $\rho_i < 1, \forall i$ is necessary to guarantee $W_i < \infty, \forall i$.

4) Problem Definition: We seek stable policies that minimize that average delay per message. These policies may be described by a collection of densities $\{p_i(x,y)\}_{i=1}^n$ with the information constraints described in the set description. That is, we must solve the following optimization problem. For emphasis, the dependence of the various terms on the assignment and service policies is given explicitly

$$\mathcal{OPT} : \min_{\pi \in \Pi_I} W(\pi) = \min_{(\pi_A, \pi_S) \in \Pi_I} \sum_{i=1}^n \frac{\lambda_i(\pi_A)}{\lambda} W_i(\pi_A, \pi_S).$$

We call this problem the Dynamic Pickup-Delivery Problem (DPDP).

B. Main Results

The goal of the current paper is to find the minimum average message delay achievable by any valid control policy for the Dynamic Pickup and Delivery Problem. We further divide the control policies into two categories based on the information structure in place for making the control decisions. In the Source Only structure, only message source locations are known before the message is picked up. In the Source and Destination structure, both the source and destination locations of messages are known as soon as the message arrives. First, we will prove lower bounds on the average message delays achieveable by control policies from these two groups. We will then propose policies that adhere to these information structures and will show that the order of the delay scaling demonstrated by these policies matches that of the lower bounds. Therefore these policies are order optimal and the lower bounds may be achieved. In particular, we prove the following two theorems:

Theorem 1:
(a) For any policy in Π_{SO} under the Source Only information structure, the average delay per message is finite only if $\rho = 2\lambda(n)\bar{s}(n)/n < 1$ and it is lower bounded as

$$W_{SO} \geq \max\left\{\frac{\gamma^2\lambda(n)A}{v^2(1-\rho)^2n} - \frac{n(1-2\rho)}{2\lambda(n)}, \frac{c_1\sqrt{A}}{v(1-\rho)}, 2\bar{s}(n)\right\}$$

with constants $\gamma = 2/3\sqrt{2\pi}$ and $c_1 \approx 0.52$.

(b) Further, if $\rho < 1$ then there exists a policy using Source Only information, for which the average delay is finite and is upper bounded as

$$W_{SO} = O\left(\frac{\lambda(n)A}{v^2(1-\rho)^2n}\right) + O\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + O(\bar{s}(n))$$

TABLE I
AVERAGE DELAY SCALING OVER VARIOUS RANGES OF $\lambda(n)$

Information Structure Type	$\lambda(\mathbf{n}) = \mathbf{o}(\mathbf{n})$	$\lambda(\mathbf{n}) = \mathbf{\Omega}(\mathbf{n})$ and $\lambda(\mathbf{n}) = \mathbf{o}(\mathbf{n^{3/2}})$	$\lambda(\mathbf{n}) = \mathbf{\Omega}(\mathbf{n^{3/2}})$
Source Only	$\Theta\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + \Theta(\bar{s}(n))$	$\Theta\left(\frac{\lambda A}{v^2(1-\rho)^2 n}\right)$	$\Theta\left(\frac{\lambda A}{v^2(1-\rho)^2 n}\right)$
Source and Destination	$\Theta\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + \Theta(\bar{s}(n))$	$\Theta\left(\frac{\sqrt{A}}{v(1-\rho)}\right)$	$\Theta\left(\frac{\lambda A}{v^2(1-\rho)^2 n^{3/2}}\right)$

for all $\lambda(n)$. Therefore the lower bound scaling is achievable, and $\rho < 1$ is necessary and sufficent for stability.

Theorem 2:

a) For any policy in Π_{SD} under the Source-Destination information structure, the average delay per message is finite only if $\rho = 2\lambda(n)\bar{s}(n)/n < 1$. In that case, the following lower bounds hold. If both $\lambda(n)A/v^2 n^{3/2} \to \infty$ and $\lambda(n)\sqrt{A}/vn^{5/4} \to \infty$, then

$$W_{SD} = \Omega\left(\frac{\lambda(n)A}{v^2(1-\rho)^2 n^{3/2}}\right) + \Omega\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + \Omega\left(\bar{s}(n)\right).$$

b) Further, if $\rho < 1$ then there exists a policy using Source and Destination information for which the average delay is finite and is upper bounded as

$$W_{SD} = O\left(\frac{\lambda(n)A}{v^2(1-\rho)^2 n^{3/2}}\right) + O\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + O(\bar{s}(n))$$

for all $\lambda(n)$. Therefore the lower bound scaling is achievable and $\rho < 1$ is necessary and sufficient for stability.

Theorem 1 first quantifies the achievable performance for control policies with some minimal amount of information. Theorem 2 then quantifies the effect of additional information on achievable performance. We note that even the full information case is greater than the results on the DTRP in [2] by a factor of $\sqrt{n}$.

C. Interpretation of $\lambda(n)$ Scaling

For convenience, the results of the above theorems are given in Table I. For ease of exposition, assume $\sqrt{A}/v = \Omega(1)$, that is, the time to cross the entire region is constant or increasing.

The delay of a message is made up of three components: 1) time spent traveling directly to the source and destination locations of the message itself, 2) time the vehicle to which the message is assigned spends traveling to serve other messages, and 3) onsite service times.

When $\lambda(n) = o(n)$ is very small and $1 - \rho$ is of moderate size, the onsite service time can dominate. In all other cases, the impact of the onsite service times is generally captured by the $(1-\rho)$ terms in the denominator. The scaling of $\lambda(n)$ determines which of the other two travel time terms dominates.

For $\lambda(n) = o(n)$, the arrival rate of messages per vehicle, $\lambda(n)/n$, shrinks as the number of vehicle increases. Therefore, most of the delay is accumulated during the travel associated with the message's own service. Roughly, this is captured by the $\sqrt{A}$ term, which is the average distance between uniformly

distributed source and destination locations. This lower bound will be proven in Section II-D. Note that this scaling is not a function of the number of vehicles n.

For $\lambda(n) = o(n)$ and $\lambda(n) = \Omega(n)$, the average delay per message is increasing as a function of n for the Source Only policy, but shrinks for the Source and Destination policy. Briefly, this is because the number of vehicles n can affect both the average rate of messages arriving to the vehicle, $\lambda(n)/n$, and also the area that a vehicle must cover to service those messages. The Source and Destination case performs significantly better, because as n increases, destination information may be exploited to shrink the region over which each vehicle must travel (roughly $A/\sqrt{n}$). The region shrinks faster than the rate at which the number of messages to be served per vehicle is growing ($\lambda(n)/n$). Due to this, again the travel time between the source and destination of an individual message dominates. The Source Only policies cannot exploit this information to reduce their travel time, and therefore the delay in traveling between other messages dominates.

For $\lambda(n) = \Theta(n^{3/2})$, the number of messages per vehicle is always growing faster than any valid policy can shrink the service regions for the vehicles. Therefore, in this regime, the travel time for other messages dominates.

D. Organization

The rest of the paper is organized as follows. In Section II, we provide some preliminary technical results that will be useful in the remaining analysis. Section III proves the lower bounds claimed in Theorems 1(a) and 2(a). Section IV describes and analyzes policies that achieve the claimed performance of Theorems 1(b) and 2(b). Finally, in Section V we present discussion and directions for future work.

II. PRELIMINARY TECHNICAL RESULTS

In this section, we review several technical results that will be used in the remainder of this paper.

A. Single Vehicle Assignment Distributions

Equations (1)–(4) provide characterizations of valid collections of densities $p_i(x, y)$. These equations may also provide some useful bounds on the density $p_i(x, y)$ for a single vehicle i.

First, we will adopt the following notation: for any reasonable function $g(\cdot)$

$$\mathbb{E}_\theta[g(\cdot)] \triangleq \int_{\mathcal{A}} g(\theta) \, d\theta.$$

Essentially, $\mathbb{E}_\theta$ is the standard (Lebesgue) integration. We retain the reference to the variable θ in order to sometimes differentiate the integration with respect to source and destination location. Note that $\mathbb{E}[\cdot]$ is not the same as $E[\cdot]$ as the latter has already been defined as expected value.

Assume that the arrival rate of messages to be served by vehicle i is fixed to be λ_i. Then, from (3) we have

$$\mathbb{E}_x[\mathbb{E}_y[p_i(x,y)]] = \frac{\lambda_i}{\lambda}. \tag{8}$$

That is, each $p_i(x,y)$ is a scaled probability density with scaling λ_i/λ.

By definition, $p_i(x,y) \geq 0, \ \forall (x,y)$. Further, from the defining (1), $p_i(x,y) \leq 1/A^2$. This implies

$$p_i(x,y) \in \left[0, \frac{1}{A^2}\right]. \tag{9}$$

Finally, note that since the vehicle must both pickup and deliver each message assigned to it, exactly half of the service locations visited by a vehicle are pickup locations. Therefore, we may define $f_i(\zeta), \zeta \in \mathcal{A}$, the normalized probability density of vehicle i servicing (i.e., either picking up or delivering) a message at location ζ, as a uniform mixture of the pickup and delivery distributions. That is

$$f_i(\zeta) = \frac{1}{2}\frac{\lambda}{\lambda_i}\left[\mathbb{E}_x[p_i(\cdot,\zeta)] + \mathbb{E}_y[p_i(\zeta,\cdot)]\right]. \tag{10}$$

B. Dynamic Traveling Repairperson Problem

Before beginning our analysis of the DPDP problem, it is important to more precisely state a few results on the related Dynamic Traveling Repair-person Problem (DTRP) that were proven by Bertsimas and van Ryzin [1]–[3] and that will be used in our lower bound analysis of the DPDP. The DTRP considers the case in which demands arrive to a convex environment $\mathcal{A}$ of area A according to some arrival process with demands being randomly located in the region according to some distribution. A demand is serviced when a vehicle arrives to the demand location and spends a random amount of onsite service time, s, to service the demand. To perform these services, there are n vehicles that travel with bounded velocity $\leq v$ within $\mathcal{A}$. The average system utilization is defined in the standard queueing theory sense to be $\rho = \lambda\bar{s}/n$. The demands are to be serviced in such a way that all demands are eventually serviced and average delay between arrival and service of the demands, W, is minimized.

In the case that demands arrive according to a Poisson process with rate λ and demand locations are independently and identically uniformly distributed in $\mathcal{A}$, the average delay of message in the system is:

Theorem 3: (Theorem 2 in [2]) :

$$W \geq \gamma^2 \frac{\lambda A}{n^2 v^2(1-\rho)^2} - \frac{n(1-2\rho)}{2\lambda} \tag{11}$$

for constant $\gamma = 2/3\sqrt{2\pi}$.

The results of [3] treat the more general case of non-Poisson arrivals and nonuniform iid demand distributions. Although the DPDP presently considers Poisson arrivals, [3] shows that the Poisson assumption is easily taken care of with little change to the delay results. Further, they consider two classes of policies: spatially unbiased and spatially biased. Spatially unbiased policies require that the average expected delay of a message is the same regardless of the demand location, and spatially biased policies simply remove this restriction. Therefore, if we are not concerned about the notion of spatial biasedness, the results on spatially biased policies provide the strongest result. Below we state a slightly modified version of result in [3] on the average delay over all messages that arrive according to demand distribution $f(\zeta)$ and are served under a spatially biased policy.

Theorem 4: (Theorem 2 From [3] (Modified)): If both $\lambda\mathbb{E}[\sqrt{f}]/vn \to \infty$ and also $\lambda\mathbb{E}[\sqrt{f}]^2/v^2 n \to \infty$, then

$$W = \Omega\left(\frac{\lambda(\mathbb{E}[f^{2/3}])^3}{v^2(1-\rho)^2 n^2}\right).$$

Theorem 4 follows with a slight modification of the proof in [3], which may be found in the appendix .

C. Asymptotic Scaling of TSP Tour Length

For analysis of the specific control policies we will present, it is necessary to provide some results on the scaling behavior of solutions to the Traveling Salesperson Problem (TSP). In the case that N, the number of locations to be visited on the tour, is large, the length of the TSP tour may be bounded with the following asymptotic result originally due to Beardwood, Halton, and Hammersley [9], but more recently stated in [10]:

Theorem 5: Given N points uniformly distributed over a region of area A, and denoting the expected length of the optimal TSP tour through these points as L_N, there exists a constant $0 < \beta < \infty$ such that

$$\lim_{N\to\infty} \frac{E[L_N]}{\sqrt{N}} = \beta\sqrt{A} \tag{12}$$

with probability 1. β has been estimated through simulation to be $\beta \approx 0.72$. Furthermore, the variance of the length of the optimal tours scales as $var(L_N) = O(1)$. That is, for N large, $E[L_N] \approx \beta\sqrt{NA}$ and further $var(L_N)/N = 0$.

Theorem 5 may also be generalized for non-uniform distributions as in [11].

Theorem 6: If $X_1, \ldots, X_N$ are identically and independently distributed (i.i.d.) according to a general absolutely continuous distribution with density $f(x)$ and compact support $\mathcal{A}$, then the following limit holds:

$$\lim_{N\to\infty} \frac{E[L_N]}{\sqrt{N}} = \beta \int_\mathcal{A} f^{1/2}(x)dx.$$

When N is not large, the average length of a TSP tour through N points may be bounded by the length of a worst case tour. Consider a tour that travels row-wise through the center points of a grid of $\sqrt{A/N} \times \sqrt{A/N}$ cells. The tour pauses in each cell to visit any demand locations within that cell. Within each cell, the tour begins at the center point, travels to each of demand locations in the cell and then returns to the center of the cell before progressing to the next cell. The tour through the center points has length $N\sqrt{A/N} + 2\sqrt{A}$. Each of the intracell demand detours adds at most $\sqrt{2A/N}$ for a total of $N\sqrt{2A/N}$.

Summing these two components gives a worst case TSP tour length bounded by $2\sqrt{2}\sqrt{NA}$.

That is, even when N is not large, the scaling of the TSP tour length is bounded in terms of N and A in the same way as in Theorem 5, with the scaling constant β increased to $2\sqrt{2}$.

D. Lower Bound on Average Delay in Light Traffic

In the current setup, vehicles must pause at a service location for the duration of the message pickup and delivery service times. This restriction implies a lower bound on message delay, even in the case that the arrival rates per vehicle are small.

Before stating the main theorem of this section, we require some additional notation. Let message j arrive at time t_j, complete pickup service at time v_j, and complete delivery service and depart the system at time y_j. With this notation, the arrival process is equivalent to $\Lambda(t) = \max\{j | t_j < t\}$. The counting processes associated with the cumulative pickup and delivery services may be defined as $\mathcal{V}(t) = card\{j | v_j < t\}$ and $D(t) = card\{j | y_j < t\}$, respectively.

Let $N(t)$ be the number of the messages in the system at time t. Because each message in the system is awaiting exactly one of two kinds of service at any time t, we further define $N_1(t)$ to be the number that have arrived but have not been picked up, and let $N_2(t)$ be the number that have been picked up but not yet delivered. These three processes are defined in terms of the arrival and service counting processes as

$$N(t) = \Lambda(t) - D(t) = N_1(t) + N_2(t)$$
$$N_1(t) = \Lambda(t) - \mathcal{V}(t)$$
$$N_2(t) = \mathcal{V}(t) - D(t).$$

We may also define the limiting distributions for the number in system seen by an arrival or a departure as

$$P(\tilde{N}^- = k) \triangleq \lim_{j \to \infty} P(N(t_j) = k),$$
$$P(\tilde{N}^+ = k) \triangleq \lim_{j \to \infty} P(N(y_j) = k)$$

when the appropriate limits exist. $\tilde{N}_1^-, \tilde{N}_1^+, \tilde{N}_2^-, \tilde{N}_2^+$ and their limiting expectations are defined similarly.

Theorem 7: For any stable policy for the No Relay DPDP for which the following properties hold:

1) arrivals to vehicles are independent Poisson processes;
2) onsite message service can only occur when a vehicle is stopped at the message service location;
3) $W_i, N_1^-, N_1^+, N_2^-, N_2^+$ have limiting distributions for all i;
4) $\lambda_i \geq \lambda_{i'} \implies W_i \geq W_{i'}$;
5) $\rho_i = 2\lambda(n)\bar{s}(n)/n < 1$;

the expected message delay W is lower bounded as follows:

$$W \geq \frac{c_1\sqrt{A}}{v(1-\rho)}$$

where $c_1 \approx 0.52$.

This proof requires two main steps. First, a lemma relating the delay while the vehicle is in onsite service time to the total delay

is proven. Then, this is combined with travel delay to derive the result. First, we have the follow Lemma bounding $E[W_{O,i}]$ for a single vehicle with arrivals of rate λ_i.

Lemma 1: When each of the following distributions exist for a single vehicle i: $\tilde{W}_i, \tilde{W}_{O,i}, \tilde{N}_{1,i}^-, \tilde{N}_{1,i}^+, \tilde{N}_{2,i}^-, \tilde{N}_{2,i}^+$, the onsite service time and total service time of messages served by that vehicle are related as

$$W_{O,i} \geq \rho_i W_i$$

where $\rho_i = 2\lambda_i \bar{s}(n)$.

Proof: For this proof, we shall drop the reference to the vehicle index i and assume that the limits over the message index j are taken only for j that are served by vehicle i, that is, $j \in \{(i_1, i_2, \ldots, i_j, \ldots) : W_i(i_j) > 0\}$. First, we find the relation between $W_O(j)$ and $W(j)$ for an individual message, and then we take the appropriate limits. For a work-conserving system, these two measures are the same, that is, the system is always in onsite service while there are messages in the system waiting to be served.

$W_{O,i}(j)$ is equal to the sum of three terms: 1) the time, denoted by $R(t_j)$, to complete the service of the message (if any) in service when message j arrives, 2) the total number of complete pickups and deliveries completed in the interval $[t_j, y_j)$ of length $W(j)$, multiplied by the service time $\bar{s}(n)$, and 3) the message's own final delivery service. If $R(t_j) = 0$, that is, there is no message in service at time t_j, then in terms of the service completion processes, $\mathcal{V}(t)$ and $D(t)$, $W_O(j)$ is defined as

$$W_O(j) = \bar{s}(n) \left[(\mathcal{V}(y_j) - \mathcal{V}(t_j)) + (D(y_j) - D(t_j)) + 1 \right].$$

If $R(t_j) > 0$, then the completion of the message in service at time t_j is already included in the difference $(\mathcal{V}(y_j) - \mathcal{V}(t_j)) + (D(y_j) - D(t_j))$. To add $R(t_j)$, we must first subtract this service. Adding in the final service of the selected message itself yields

$$W_O(j) = \bar{s}(n) \left[(\mathcal{V}(y_j) - \mathcal{V}(t_j)) + (D(y_j) - D(t_j)) \right] + R(t_j).$$

In either case, because we are looking for a lower bound, we may ignore the residual terms and use the following bound:

$$W_O(j) \geq \bar{s}(n) \left[(\mathcal{V}(y_j) - \mathcal{V}(t_j)) + (D(y_j) - D(t_j)) \right]. \quad (13)$$

We may compute the number of services by relating them to the number in system processes and the arrival process

$$\mathcal{V}(y_j) - \mathcal{V}(t_j) = \Lambda(y_j) - \Lambda(t_j) - N_1(y_j) - N_1(t_j)$$
$$D(y_j) - D(t_j) = \mathcal{V}(y_j) - \mathcal{V}(t_j) - N_2(y_j) - N_2(t_j).$$

Because $N_1(t)$ and $N_2(t)$ are both unit increment/decrement processes, $\tilde{N}_1^- \overset{d}{=} \tilde{N}_1^+$ (and likewise for $\tilde{N}_2$) when these distributions exist. This is due to an extension of Burke's Theorem, which may be found in [12]. In particular, $E[\tilde{N}_1^-] = E[\tilde{N}_1^+]$ and $E[\tilde{N}_2^-] = E[\tilde{N}_2^+]$. Further, the Poisson arrival rate implies that, for each interval $[t_j, y_j]$

$$E[(\Lambda(y_j) - \Lambda(t_j))] = \lambda E[(y_j - t_j)]. \quad (14)$$

Combining these two facts, and taking limits, we have

$$
\lim_{j\to\infty} E[(V(y_j) - V(t_j))] = \lim_{j\to\infty} E[(\Lambda(y_j) - \Lambda(t_j))]
$$
$$
= \lambda W, \qquad (15)
$$
$$
\lim_{j\to\infty} E[(D(y_j) - D(t_j))] = \lim_{j\to\infty} E[(V(y_j) - V(t_j))]
$$
$$
= \lambda W. \qquad (16)
$$

Therefore, combining (15) and (16) with (13) and adding back in the i notation yields

$$
W_{O,i} \geq \bar{s}(n) \left[\lambda_i W_i + \lambda_i W_i \right] \geq \rho_i W_i.
$$

$\blacksquare$

With this Lemma, we complete the proof of Theorem 7.

Proof of Theorem 7: Consider the total waiting time $W(j)$ of a randomly tagged message j. Because onsite service can occur only when the vehicle is not traveling, the waiting time may be divided into two parts: $W_T(j)$, the time that the vehicle is traveling between message locations, and $W_O(j)$, the time the vehicle spends in onsite service. $W_O(j)$ includes the onsite service time of the tagged message as well at the onsite service times of any other messages served between the tagged message's arrival and final delivery service.

Recall that $W = W_T + W_O$. The travel time may be bounded by the time to travel the expected distance between the source and destination locations of the randomly tagged message . The actual time in travel may include deviations from this straight line distance, and so this term is a lower bound on $W_T(j)$. Because sources and destinations are independently and uniformly distributed, this distance is $c_1\sqrt{A}$, where the constant $c_1 \approx 0.52$ (see [13, p. 135]). Therefore

$$
W_T \geq \frac{c_1\sqrt{A}}{v}.
$$

For the onsite waiting time, we have the following claim: $W_O \geq \rho W$. This does not follow immediately from Lemma 1 which was proven for a single vehicle only. However, taking the weighted sum of these terms for each vehicle and applying the definitions in (7) (and similarly for W_O)

$$
\sum_{i=1}^{n} \frac{\lambda_i}{\lambda} W_{O,i} \geq \sum_{i=1}^{n} \frac{\lambda_i}{\lambda} \rho_i W_i, \qquad (17)
$$
$$
W_O \geq \frac{\sum_{i=1}^{n} \rho_i}{n} \sum_{i=1}^{n} \frac{\lambda_i}{\lambda} W_i \qquad (18)
$$
$$
\geq \rho W. \qquad (19)
$$

The implication of (18) from (17) is given by the assumption that W_i and $W_{i,O}$ are both increasing functions of λ_i (and therefore of ρ_i). Combining these bounds on W_T and W_O

$$
W \geq \frac{c_1\sqrt{A}}{v} + \rho W
$$
$$
\geq \frac{c_1\sqrt{A}}{v(1-\rho)}.
$$

$\blacksquare$

III. LOWER BOUNDS ON AVERAGE DELAY

Now that we have presented some preliminaries that will be useful in our analysis, we prove the claimed lower bounds of Theorems 1(a) and 2(a) for arbitrary policies. Policies achieving these lower bounds will be described in Section IV.

A. Lower Bound: Source Only

We consider the Source Only information structure to be one of minimal information. Because destination locations are not known immediately upon message arrival, this information may not be exploited when assigning messages to vehicles.

Theorem 1(a): For any policy in Π_{SO} under the Source Only information structure, the average delay per message is finite only if $\rho = 2\lambda(n)\bar{s}(n)/n < 1$ and it is lower bounded as

$$
W_{SO} \geq \max \left\{ \frac{\gamma^2 \lambda(n) A}{v^2(1-\rho)^2 n} - \frac{n(1-2\rho)}{2\lambda(n)}, \frac{c_1\sqrt{A}}{v(1-\rho)}, 2\bar{s}(n) \right\}
$$

with constants $\gamma = 2/3\sqrt{2\pi}$ and $c_1 \approx 0.52$.

Proof: Consider a fixed stable assignment and service policy in Π_{SO}. Each message is assigned to its vehicle immediately upon arrival. Each vehicle is then treated as a queue of messages that have been assigned to it. Consider the queue at vehicle i. From (3), the arrival process to vehicle i is a Poisson process of rate λ_i.

To lower bound the average delay of messages at a single vehicle, we consider a simplified system in which the same message assignment process holds, but messages arrive directly at the vehicle according to a Poisson process of rate λ_i. That is, vehicles do not spend any time in picking up messages. Further, for consistency of the ρ notation, let the onsite service time for delivering each message be $2\bar{s}(n)$. This simplified system naturally has lower delay than the original system.

Because vehicles only have access to information about the source locations of messages, the destination locations of the messages may not be exploited by the message assignment policy. Since the distribution of destination locations is independent of the arrival locations, the distribution of the destination locations of the messages assigned to a single vehicle is the same as that of the overall destination process, irrespective of assignment policy. So for any policy in Π_{SO}, each vehicle will service messages with destination locations distributed uniformly at random in $\mathcal{A}$. Thus, it is sufficient to lower bound delay of the following simplified system: each vehicle has messages arriving according to a Poisson process of rate λ_i with uniformly distributed delivery locations and onsite service time $2\bar{s}(n)$ for delivery.

Therefore, for each vehicle, this delivery problem may be formulated as a single-vehicle Dynamic Traveling Repairperson Problem. In this formulation, $\rho_i = 2\lambda_i \bar{s}(n)$, where the factor 2 reflects the two onsite service times required for pickup and delivery in the original system. Applying the DTRP results of Theorem 3 to this formulation, we obtain the average delay for messages served by a single vehicle with message arrival rate λ_i:

$$
W_i \geq \gamma^2 \left(\frac{\lambda_i A}{v^2(1-\rho_i)^2} \right) - \frac{1-2\rho_i}{2\lambda_i}. \qquad (20)
$$

This DTRP result lower bounds the delay achievable by any service policy to serve the messages assigned to a single vehicle at a fixed rate λ_i. Using this result to focus on assignment policies only, we then have the following relaxation of the optimization $\mathcal{OPT}$:

$$\mathcal{OPT}_1 : \min_{\{p_i(x,y)\}_{i=1}^n} \sum_{i=1}^n \frac{\lambda_i(\pi_A)}{\lambda} W_i^*(\pi_A)$$

$$s.t.\ W_i^*(\pi_A) \leq W_i(\pi_A, \pi_S),\ \forall \pi_S$$

$$\sum_{i=1}^n \lambda_i(\pi_A) = \lambda$$

$$\lambda_i(\pi_A) = \lambda \int_A \int_A p_i(x,y)dxdy, \forall i$$

$$\sum_{i=1}^n p_i(x,y) = \frac{1}{A^2},\ \forall x,y \in \mathcal{A}$$

$$p_i(x,y) = \frac{1}{A}p_i(x),\ \forall y, i.$$

$\mathcal{OPT}_1$ minimizes the weighted sum of lower bounds on the average delays over all vehicles by the selection of a valid assignment policy. The lower bounds arise by bounding the delay that may be achieved by any service policy given the fixed assignment policy. The weights are given according to a joint constraint on the policies used by the individual vehicles. If the minimum is finite, each of the $W_i(\pi)$ must be finite as well and the system is stable.

To lower bound the average delay over all messages, we may then bound the solution of $\mathcal{OPT}_1$ by further optimizing over the collection of $\{\lambda_i\}_{i=1}^n$ of valid assignment policies:

$$\mathcal{OPT}_1 \geq \min_{\{\lambda_i\}_{i=1}^n} \sum_{i=1}^n \frac{\lambda_i}{\lambda}\left(\gamma^2\left(\frac{\lambda_i A}{v^2(1-\rho_i)^2}\right) - \frac{1-2\rho_i}{2\lambda_i}\right)$$

$$s.t. \sum_{i=1}^n \lambda_i = \lambda.$$

Note that the optimization over the set of all $\{p_i(x,y)\}_{i=1}^n$ has been replaced by the relaxed restriction on the sum of the λ_i. This relaxation certainly provides a lower bound to the original optimization $\mathcal{OPT}_1$.

Removing constant terms and noting that $\sum_{i=1}^n \rho_i = \sum_{i=1}^n 2\lambda_i\bar{s}(n) = n\rho$, this is equivalent to

$$\min_{\{\lambda_i\}_{i=1}^n} \sum_{i=1}^n \frac{\lambda_i^2}{(1-\rho_i)^2} = \sum_{i=1}^n \frac{\lambda_i^2}{(1-2\lambda_i\bar{s})^2}$$

$$s.t. \sum_{i=1}^n \lambda_i = \lambda.$$

This optimization is straightforward to solve, using Lagrange multipliers, for example. We find that the optimal solution is $\lambda_i = \lambda/n, \forall i$. In this case, $\rho_i = 2\lambda(n)\bar{s}(n)/n, \forall i$, and the maximum service time allowed for stability of the service policies on each of the vehicles is then $\bar{s}(n) < n/2\lambda$. Therefore, we have the following lower bound on the average delay over all vehicles:

$$W_{SO} \geq \gamma^2\left(\frac{\lambda(n)A}{v^2(1-\rho)^2 n}\right) - \frac{n(1-2\rho)}{2\lambda(n)}. \tag{21}$$

If $\lambda(n) = \Omega(n)$, the first term in (21) dominates the second. For $\lambda(n) = o(n)$, however, this delay shrinks to 0, and the scaling

bound from Theorem 7 dominates. Combining these two results, we have therefore proven Theorem 1(a). ∎

B. Lower Bound: Source-Destination

If both the Source and Destination locations are known upon message arrival, assignment policies may exploit this information to limit the area covered by each vehicle in making its pickups and deliveries. We show that this has the effect of reducing the minimum average delay of messages in the system.

Theorem 2(a): For any policy in Π_{SD} under the Source-Destination information structure, the average delay per message is finite only if $\rho = 2\lambda(n)\bar{s}(n)/n < 1$. In that case, the following lower bounds hold. If both $\lambda(n)A/v^2 n^{3/2} \to \infty$ and $\lambda(n)\sqrt{A}/vn^{5/4} \to \infty$, then

$$W_{SD} = \Omega\left(\frac{\lambda(n)A}{v^2(1-\rho)^2 n^{3/2}}\right) + \Omega\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + \Omega\left(\bar{s}(n)\right).$$

Proof: Consider a fixed stable assignment and service policy in Π_{SD} based on source-destination information. Given the source and destination assignment distributions induced by this policy, we will again construct a set of simplified single-vehicle DTRP systems, the average delay of which will lower bound the delay encountered by messages in the original DPDP system.

As in the previous section, examine the service policy of a single vehicle i and consider the queue induced by messages that are assigned to this vehicle. The DTRP demand location associated with each message is selected uniformly at random between the source $s(j)$ or the destination $d(j)$ location of the message. That is, instead of performing both pickup and delivery as in the DPDP or delivery only as in the proof of the Source Only lower bound above, this DTRP visits exactly one of the pickup and delivery locations for each message, with either location being chosen with probability 1/2. Therefore, the distribution of demand locations arriving to this DTRP queue is the same as the normalized density of vehicle i's pickup and delivery locations, i.e., $f_i(\zeta)$. Again assume that the onsite service time required to perform the message service is $2\bar{s}(n)$. Note that since the DTRP queue ignores either the pickup or delivery requirement of each message, the delay of the demands in the DTRP queue is less than that of messages in the original system.

This DTRP queue fits the framework of the single vehicle Dynamic Traveling Repairperson Problem with generalized demand distributions. Then, according to Theorem 4, we have the following bound on minimum delay for a single vehicle policy with demand distribution $f_i(\zeta)$, arrival rate $\lambda_i = \Omega(1)$, and $\rho = 2\lambda_i\bar{s}$

$$W_i \geq \gamma^2 \frac{\lambda_i(\mathbb{E}_\zeta[f_i^{2/3}])^3}{v^2(1-\rho_i)^2}. \tag{22}$$

Recall the relation of $f_i(\zeta)$ to $p_i(x,y)$ and the constraints imposed by the delivery requirement on $p_i(x,y)$ from (8)–(10). First, the $p_i(x,y)$ have the following basic constraints:

$$p_i(x,y) \in \left[0, \frac{1}{A^2}\right], \tag{23}$$

$$\mathbb{E}_x[\mathbb{E}_y[p_i]] = \frac{\lambda_i}{\lambda}. \tag{24}$$

Then (10) implies the following two lower bounds:

$$\mathbb{E}_\zeta[f_i^{2/3}] \geq \left(\frac{\lambda}{2\lambda_i}\right)^{2/3} \mathbb{E}_\zeta[\mathbb{E}_x[p_i(x,\zeta)]^{2/3}] \qquad (25)$$

$$\mathbb{E}_\zeta[f_i^{2/3}] \geq \left(\frac{\lambda}{2\lambda_i}\right)^{2/3} \mathbb{E}_\zeta[\mathbb{E}_y[p_i(\zeta,y)]^{2/3}]. \qquad (26)$$

We may take a convex combination the two lower bounds above to form the following optimization problem $\mathcal{OPT}_2$ which will then be used to lower bound the delay of a single vehicle policy with fixed arrival rate λ_i

$\mathcal{OPT}_2$:

$$\min_{p_i(x,y)} \quad \frac{1}{2}\left(\mathbb{E}_\zeta[(\mathbb{E}_x[p_i(x,\zeta)])^{2/3}] + \mathbb{E}_\zeta[(\mathbb{E}_y[p_i(\zeta,y)])^{2/3}]\right)$$

$$s.t. \quad p_i(x,y) \in \left[0,\frac{1}{A^2}\right],$$

$$\mathbb{E}_x[\mathbb{E}_y[p_i(x,y)]] = \mathbb{E}_y[\mathbb{E}_x[p_i(x,y)]] = \frac{\lambda_i}{\lambda}.$$

Consider a convex combination of two densities satisfying (23) and (24), i.e., $p_i^3(x,y) = \alpha p_i^1(x,y) + (1-\alpha)p_i^2(x,y), \forall x,y \in \mathcal{A}$. It is easy to see that the set of valid probability distributions satisfying (23) and (24) is convex. Further, by the concavity of $(\cdot)^{2/3}$, both of the lower bounds (25) and (26) are concave in $p_i(x,y)$ and so is their sum. Thus, $\mathcal{OPT}_2$ is a concave minimization over a convex set. Hence, it must attain its optima on the boundary of the feasible bounded convex set.

The boundary of the constraint set defined by (23)–(24) implies that $p_i(x,y) \in \{0, 1/A^2\}$ for all (x,y) (almost surely w.r.t. Lebesgue measure). Condition (24), along with this implication, will provide the following complete characterization of boundary:

$$p_i(x,y) = \begin{cases} \frac{1}{A^2} & \text{for all } x \in \mathcal{A}_1^i, y \in \mathcal{A}_2^i \\ 0 & \text{otherwise} \end{cases} \qquad (27)$$

with $\mathcal{A}_1^i, \mathcal{A}_2^i \subset \mathcal{A}$ of areas such that $A_1^i A_2^i = A^2 \lambda_i / \lambda$.

To minimize the cost function in $\mathcal{OPT}_2$, we must select the boundary points where the areas of A_1^i and A_2^i are equal, i.e., both are equal to $A\sqrt{\lambda_i/\lambda}$.

For any p_i satisfying the above properties we have

$$\mathbb{E}_x[p_i(x,\zeta)] = A\sqrt{\frac{\lambda_i}{\lambda}}\frac{1}{A^2} = \frac{1}{A}\sqrt{\frac{\lambda_i}{\lambda}}$$

$$\text{and } \mathbb{E}_\zeta[\mathbb{E}_x[p_i(x,\zeta)]^{2/3}] = A\sqrt{\frac{\lambda_i}{\lambda}}\left(\frac{1}{A}\sqrt{\frac{\lambda_i}{\lambda}}\right)^{2/3}$$

$$= A^{1/3}\left(\frac{\lambda_i}{\lambda}\right)^{5/6}$$

and therefore the bound (25) on $\mathbb{E}[f_i^{2/3}]$ becomes

$$\mathbb{E}[f_i^{2/3}] \geq \left(\frac{\lambda}{2\lambda_i}\right)^{2/3} \mathbb{E}_\zeta[\mathbb{E}_x[p_i(x,\zeta)]^{2/3}]$$

$$= \frac{1}{2^{2/3}}A^{1/3}\left(\frac{\lambda_i}{\lambda}\right)^{1/6}. \qquad (28)$$

Cubing (28) and then substituting this bound into (22), we thus have the following bound on minimum delay for messages served by vehicle i:

$$W_i \geq \frac{\gamma^2}{4}\frac{\lambda_i A}{v^2(1-\rho_i)^2}\sqrt{\frac{\lambda_i}{\lambda}}. \qquad (29)$$

This result lower bounds the delay achievable by any service policy for a single vehicle serving messages at rate λ_i with an assignment density $p_i(x,y)$ that is valid for a single vehicle in the original DPDP system. We may again lower bound the solution of $\mathcal{OPT}_1$ by further optimizing over the collection of $\{\lambda_i\}_{i=1}^n$. Repeating the analysis that led to the optimization problem (21), the corresponding optimization here is

$$\mathcal{OPT}_1 \geq \min_{\{\lambda_i\}_{i=1}^n} \sum_{i=1}^n \frac{\lambda_i^{5/2}}{(1-2\lambda_i\bar{s})^2} \qquad (30)$$

$$s.t. \sum_{i=1}^n \lambda_i = \lambda. \qquad (31)$$

As above, this average delay is minimized with all λ_i equal to λ/n and again $\rho_i = 2\lambda(n)\bar{s}(n)/n, \forall i$. Therefore, we have the following lower bound on the average delay with Source and Destination information:

$$W_{SD} \geq \frac{\gamma^2}{4}\frac{\lambda(n)A}{v^2(1-\rho)^2 n^{3/2}}. \qquad (32)$$

Since $\lambda_i = \lambda/n = \Omega(1)$ was required for the application of the DTRP theorem with generalized demand distributions, this bound may only be valid for $\lambda = \Omega(n)$. However, for $\lambda(n) = o(n^{3/2})$, this delay shrinks to 0, and the scaling bound from Theorem 7 again dominates. We have therefore proven Theorem 2(a). ∎

IV. POLICIES

In this section, we describe policies that achieve the delay performance claimed in Theorems 1 and 2 for Source Only and Source Destination information respectively. These policies provide additional insight into the effect of information structure on achievable delay. Furthermore, both policies achieve the lower bounds presented in the previous section.

A. Source Only Policy

In the Source Only information structure, vehicles do not know the destination of messages before they are picked up, thus this information may not be used by vehicles in deciding which messages to pick up. In fact, in the source only policy described below, each message is assigned to any of the vehicles at random. We note that "smarter" message assignments are possible to minimize the vehicles' time spent in picking up messages. For example, a vehicle could be assigned all messages that arrive in a given limited area. However, since the vehicles must still traverse the whole region to deliver messages, regardless of assignment policy, the vehicle deliveries will dominate the delay and no message assignment process with only source information can improve the order of the performance for large

arrival rates. A more complete description of the policy is given below.

a) *Message Assignment.* Upon arrival, each message is assigned to one of the vehicles uniformly at random. The message is not immediately picked up, but the vehicle is notified of the message assignment. Since the message assignment is a uniform splitting of the Poisson arrival process, the assignment of messages to each vehicle is Poisson with an arrival rate of λ/n. All messages assigned to a single vehicle that arrive in the interval $[kT, (k+1)T)$ form a batch, where T, the batch time interval, is a parameter to be determined. Each batch is deposited into a queue for its assigned vehicle upon formation at time $(k+1)T$.

b) *Message Service.* Batches for each vehicle are served in First Come, First Serve order from the vehicle's batch queue. Pickups are performed along a TSP tour through the source locations which is computed at the beginning of the interval. Once pickups are complete and destination information is collected, a TSP tour through the delivery locations is computed and the deliveries are performed accordingly. To perform each service, the vehicle stops at the source (destination) location for $\bar{s}(n)$ time to pickup (deliver) the associated message.

Theorem 1(b): Further, if $\rho < 1$ then there exists a policy using Source Only information, for which the average delay is finite and is upper bounded as

$$W_{SO} = O\left(\frac{\lambda(n)A}{v^2(1-\rho)^2 n}\right) + O\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + O(\bar{s}(n))$$

for all $\lambda(n)$. Therefore the lower bound scaling is achievable, and $\rho < 1$ is necessary and sufficent for stability.

Proof of Theorem 1(b): Consider the queue of batches assigned to a randomly selected vehicle i. Note that by the symmetry of the vehicle policies, the average delay of messages at a single vehicle is the same as the average delay over all vehicles. Since the batch interarrival time is fixed at T, the batches form a D/G/1 queue. This batching protocol is stable if and only if the expected time to service each batch of messages, T_B, is less than T, the expected time between batch arrivals. The first part of the proof bounds T in terms of the system parameters so that this stability condition is met.

The batch service time requires two TSP tours, one for pickup and one for delivery, plus the associated onsite service times to perform each service. Let N_T be the number of messages arriving in $[kT, (k+1)T)$ that are assigned to vehicle i. Therefore, using Theorem 5 to bound the travel time required for each of the shortest paths (pickup and delivery), the total expected service time required to service the messages accumulated in $[kT, (k+1)T)$ is

$$E[T_B] \leq 2E[\frac{1}{v}E[L_{N_T}|N_T] + N_T\bar{s}(n)] \tag{33}$$

$$= 2E[\beta\frac{1}{v}\sqrt{A}\sqrt{N_T} + N_T\bar{s}(n)] \tag{34}$$

$$\leq 2\beta\frac{1}{v}\sqrt{A}\sqrt{E[N_T]} + 2E[N_T]\bar{s}(n) \tag{35}$$

$$= 2\beta\frac{1}{v}\sqrt{A}\sqrt{\frac{\lambda T}{n}} + 2\frac{\lambda T}{n}\bar{s}(n) \tag{36}$$

where (34) is by Theorem 5, (35) is by concavity of $\sqrt{\cdot}$, and (36) is given by the Poisson distribution of N_T.

Therefore the following bound on T is sufficient for stability:

$$T \geq 2\beta\sqrt{\frac{\lambda A}{v^2 n}}\sqrt{T} + \rho T \geq E[T_B] \tag{37}$$

$$\implies T \geq \frac{4\beta^2 \lambda A}{v^2(1-\rho)^2 n}. \tag{38}$$

With $\lambda(n) = \Omega(n)$, the number of messages served by each vehicle is large and by Theorem 5, $\beta \approx 0.78$.

The second part of the proof uses the batch interval time T to compute the average message delay. For the remainder of the proof, arbitrarily select some $\kappa > 1$ and fix T to be

$$T = \kappa\frac{4\beta^2 \lambda A}{v^2(1-\rho)^2 n}. \tag{39}$$

Message delay has four components: 1) time waiting for batch to form, 2) time batch spends in queue, 3) time waiting for service of other vehicles in batch, and 4) time of own service. Since batch interarrival time $= T$, each message waits at most T for its batch to form, bounding 1). Letting T_Q denote the expected amount of time the batch spends in queue, 2) may be bounded using the following lemma:

Lemma 2: For the policy in Theorem 1(b) with batch time $T = \kappa(4\beta^2 \lambda A/v^2(1-\rho)^2 n)$ for some $\kappa > 1$, the delay of the batch in the queue is bounded by

$$T_Q = O(T). \tag{40}$$

The proof, found in Appendix B, uses Kingman's Bound from queueing theory [14], and is largely a matter of algebra.

Delay components 3) and 4) may be bounded by bounding the expected total batch service time for the batch in which an arbitrary message arrives. We compute the expected batch service time by first conditioning on the size of the batch in which a message arrives and then taking the expectation over this batch size. The upper bound on batch service time will be computed separately for two cases: $\lambda T/n > 1$ and $\lambda T/n \leq 1$. Both cases begin the same way as below.

If a message arrives in a batch of size B, according to the worst case TSP tour discussion in Section II-C, the total service time of the batch may be bounded by

$$S_B \leq 2\left(\frac{2\sqrt{2}\sqrt{A}\sqrt{B}}{v} + \bar{s}B\right). \tag{41}$$

By the law of random incidence, a randomly selected message arrives in a batch of size B with probability

$$P\{\text{message in batch of size } B\} = \frac{BP\{\text{batch has size } B\}}{E[B]}$$

where the batch sizes are Poisson with parameter $\lambda T/n$. Therefore, the expected batch service time is

$$
\begin{aligned}
E[S_B] &= \frac{4\sqrt{2}\sqrt{A}}{v} \sum_{k=1}^{\infty} \frac{k^{3/2} \left(\frac{\lambda T}{n}\right)^k e^{-\lambda T/n}}{\frac{\lambda T}{n} k!} + \\
&\quad 2\bar{s}(n) \sum_{k=1}^{\infty} \frac{k^2 \left(\frac{\lambda T}{n}\right)^k e^{-\lambda T/n}}{\frac{\lambda T}{n} k!} \\
&= \frac{4\sqrt{2}\sqrt{A}}{v} E[(B+1)^{1/2}] + 2\bar{s}(n) E[B+1] \\
&\leq \frac{4\sqrt{2}\sqrt{A}}{v} \sqrt{\left(\frac{\lambda T}{n}+1\right)} + 2\bar{s}(n)\left(\frac{\lambda T}{n}+1\right) \\
&= \frac{4\sqrt{2}\sqrt{A}}{v} \sqrt{\left(\frac{\lambda T}{n}+1\right)} + \rho T + 2\bar{s}(n). \quad (42)
\end{aligned}
$$

For $\lambda T/n > 1$, $\lambda T/n + 1$ can be bounded by $2\lambda T/n$. In that case, (42) can be bounded by

$$
\begin{aligned}
E[S_B] &\leq 8\sqrt{\frac{A\lambda T}{nv^2}} + \rho T + 2\bar{s}(n) \\
&= 16\sqrt{\kappa}\beta \frac{\lambda A}{v^2(1-\rho)n} + \rho T + 2\bar{s}(n) \\
&= O\left((1-\rho)T\right) + \rho T + 2\bar{s}(n) \\
&= O(T) + O(\bar{s}(n)).
\end{aligned}
$$

Combining this with delay components 1) and 2) above, for $\lambda T/n > 1$

$$
W_{SO} = O(T) + O(\bar{s}(n)). \quad (43)
$$

For $\lambda T/n \leq 1$, $\lambda T/n + 1$ is bounded above by 2. In that case, (42) is bounded by

$$
E[S] \leq 8\frac{A}{v} + \rho T + 2\bar{s}(n). \quad (44)
$$

Since $\rho < 1$, $\sqrt{A}/v < \sqrt{A}/v(1-\rho)$. Note also that $\lambda T/n \leq 1$ implies that for some constants c_1 and c_2

$$
\begin{aligned}
\frac{\lambda^2 A}{v^2(1-\rho)^2 n^2} &\leq c_1^2 \Rightarrow \frac{\lambda}{n} \leq c_1 \frac{v(1-\rho)}{\sqrt{A}} \\
&\Rightarrow T \leq c_2 \frac{v(1-\rho)}{\sqrt{A}} \frac{A}{v^2(1-\rho)^2} = O\left(\frac{\sqrt{A}}{v(1-\rho)}\right).
\end{aligned}
$$

Therefore, combining with 1) and 2) above, for $\lambda T/n \leq 1$

$$
W_{SO} = O\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + O(\bar{s}(n)). \quad (45)
$$

Therefore

$$
W_{SO} = O\left(\frac{\lambda(n)A}{v^2(1-\rho)^2 n}\right) + O\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + O(\bar{s}(n))
$$

for all $\lambda(n)$, and Theorem 1(b) is proven. ∎

B. Source-Destination Policy

In the Source-Destination information structure, destination information may be used by vehicles in deciding which messages to pick up. By exploiting this information, vehicles need not traverse the entire geographical region when servicing messages, but may instead only pick up messages that have both source and destination locations in a limited area. In the source destination policy described below, each vehicle is assigned a pickup region and a delivery region. Messages are not assigned to a random vehicle as above, but are instead assigned to the vehicle that has the message's source location in its pickup region and the message's destination location in its delivery region. Even though the message service policy is similar to that used in the Source Only policy above, Theorem 2(b) shows that the change in assignment policy made possible by using both source and destination information has a significant effect on message delay. A more complete description of the policy is given below.

a) *Message Assignment.* Divide the geographical region into an $\sqrt{A/\sqrt{n}} \times \sqrt{A/\sqrt{n}}$ grid of subregions, each of area $A/\sqrt{n}$. To each of the n ordered pairs of subregions, assign exactly one vehicle to service that pair. Each vehicle is assigned to pickup all messages that originate in the first subregion of its assigned ordered pair that have a destination location in second assigned subregion. As before, all messages assigned to a single vehicle that arrive in the interval $[kT, (k+1)T)$ form a batch, where T, the batch time interval, is a parameter to be determined. Each batch is deposited into a queue for its assigned vehicle upon formation at time $(k+1)T$ for appropriate k.

b) *Message Service.* As before, batches for each vehicle are served in First Come, First Serve order from the vehicle's batch queue. Batch pickups and deliveries are performed in the same way as in the policy with Source only information with the notable addition of possible interregion travel time between source region and destination region.

Theorem 2(b): Further, if $\rho < 1$ then there exists a policy using Source and Destination information for which the average delay is finite and is upper bounded as

$$
W_{SD} = O\left(\frac{\lambda(n)A}{v^2(1-\rho)^2 n^{3/2}}\right) + O\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + O(\bar{s}(n))
$$

for all $\lambda(n)$. Therefore the lower bound scaling is achievable and $\rho < 1$ is necessary and sufficient for stabilty.

Proof: Theorem 2(b): Service of assigned messages is the same as in the Source only policy described above except that the TSP tours are performed over possibly distinct subregions of the environment. Each TSP tour now ranges over a subset of the geographical region $\mathcal{A}$ with area $A/\sqrt{n}$. Travel time between subregions must also be included in the batch service time analysis. Since the total geographical region is a square of area A, this interregion travel time may be upper bounded by $1/v\sqrt{2A}$.

Therefore, as before, the total expected service time required to service the messages accumulated in $[kT, (k+1)T]$ is

$$
\begin{aligned}
E[T_B] &\leq 2E[\frac{1}{v}E[L_{N_T}|N_T] + N_T\bar{s}(n) + 2\frac{1}{v}\sqrt{2A}] \\
&= 2\beta\frac{1}{v}\sqrt{\frac{A}{\sqrt{n}}}\sqrt{\frac{\lambda T}{n}} + 2\frac{\lambda T}{n}\bar{s}(n) + \frac{2\sqrt{2}\sqrt{A}}{v}.
\end{aligned}
$$

Again, $\beta \approx 0.78$ if $\lambda(n) = \Omega(n)$, else a worst case tour may be used to prove the above for $\beta = 2\sqrt{2}$.

Therefore the following bound on T is sufficient for stability:

$$T \geq 2\beta \frac{1}{v}\sqrt{\frac{\lambda A}{n^{3/2}}}\sqrt{T} + \rho T + \frac{2\sqrt{2}\sqrt{A}}{v} \geq E[T_B]. \quad (46)$$

This equation is quadratic in $\sqrt{T}$ and may be easily solved for T. Specifically, the following T is sufficient for stability:

$$T \geq \frac{4\beta^2 \lambda A}{v^2(1-\rho)^2 n^{3/2}} + \frac{2\sqrt{A}}{v(1-\rho)}. \quad (47)$$

As before, the total message delay as a function of the batch time T may be bounded by fixing a batch scaling constant κ and then using Kingman's bound with $A = A/\sqrt{n}$ and $T = \kappa\left(4\beta^2\lambda A/v^2(1-\rho)^2 n^{3/2} + 2\sqrt{A}/1-\rho\right)$. Therefore, similar to the Source Only case with the altered scaling of the area A and the addition of the constant interregion travel time

$$W_{SD} = O\left(\frac{\lambda A}{v^2(1-\rho)^2 n^{3/2}}\right) + O\left(\frac{2\sqrt{A}}{v(1-\rho)}\right) + O(\bar{s}).$$

∎

V. CONCLUSION

In this paper, we have presented a dynamic vehicle routing problem, the DPDP, and obtained lower and upper bounds on the scaling of the average message delay. These results are a significant extension of the existing results on the DPDP. Furthermore, the information that is available in making assignment decisions has a significant effect on the delay scaling. In the case that Source Only information is available, the average delay scales as $\Theta(\lambda(n)A/v^2 n)$. In the case that both Source and Destination information is available, the average delay scales as $\Theta(\lambda(n)A/v^2 n^{3/2})$ which is an additional $O(\sqrt{n})$ improvement over the case where only source information is available. From a system design standpoint, these scalings quantify the performance improvements achievable by adding additional information gathering capabilities to the vehicles.

The DTRP results in [2], [3] bound the average delay for the each of the pickup and delivery problems as $\Omega(\lambda(n)A/v^2 n^2)$. This is an additional $O(\sqrt{n})$ improvement over the full information case we have examined here. We note that as long as vehicles are required to perform physical pickups and deliveries at the source and destination locations, the DTRP lower bound serves as a lower bound on the DPDP problem. It is conjectured that this delay bound can be achieved by removing the restriction that the same vehicle that picks up a message is the one that delivers it. In a followup paper [15], we show that this delay bound can be indeed achieved for the pickup and delivery problem in relay networks.

Note that the set of policies under consideration is somewhat restrictive. We consider only policies with separable assignment and service policies. Assignments are made by a centralized controller independent of the current service requirements associated with each of the vehicles. Further, the use of the $\{p_i(x,y)\}_{i=1}^n$ for the assignment policy fails to include any policies in which batches of requests are collected into a $G/G/n$ queue at a centralized depot and served in FCFS order. Such policies were proposed for the DTRP in [1]. The assignments of

consecutive messages are likely to be correlated due to their collection into a single batch, and therefore the independent assignment property fails to hold. Comparing the delay of the $G/G/n$ to the average delay of a collection of n $G/G/1$ queues as in Section IV, the $G/G/n$ assignment and service policy actually has a lower average delay than the n $G/G/1$ queues. This difference disappears as the traffic increases and the probability of vehicle idleness approaches 0. Therefore, in the limit, our policy restriction does not seem to hurt us in terms of finding the minimum delay scaling, at least for $G/G/n$ policies.

We also note that the centralized assignment policies presented in Section IV may be decentralized given appropriate assumptions on inter-vehicle communication. In the Source and Destination case, message assignments are based only on the locations associated with each message. After a centralized initialization period in which vehicles are assigned to pickup and delivery regions, no centralized decision making or inter-vehicle communication is required. A similar message assignment policy based on source locations could also be used to decentralize the Source Only policy. We note that these types of decentralized policies are not robust to vehicle dropout without the addition of some recovery mechanism to address vehicle failures occuring after the initialization period.

APPENDIX A
PROOF OF THEOREM 4

In this appendix, we prove Theorem 4(b) which is a modified version of Theorem 2 in [3].

The Dynamic Traveling Repairperson (DTRP) problem refers to the following setup: demands arrive to a closed and bounded region $\mathcal{A}$ of area A according to a stationary renewal process. Demands are independently and identically distributed according to the demand distribution $f(x)$. There are n vehicles traveling in the region with bounded velocity v to service these demands. A demand is serviced when a vehicle arrives at the demand location and spends a random service time s at that location. The goal is to service the demands with the minimum average delay W between message arrival and service.

Before proving lower bounds on this average delay, [3] provides a few additonal definitions and assumptions. Let $W(x) = [W(j)|x(j) = x]$ be the waiting time conditioned on message service location where it exists. Then the normalized waiting time function is then defined as

$$\Psi(x) = \frac{W(x)}{W}.$$

The queue occupancy density is defined to be

$$\phi(x) = f(x)\Psi(x). \quad (48)$$

Theorem 4 holds when $W(x)$, $\Psi(x)$ and $\phi(x)$ are well-defined. The proof follows the same sequence of lemmas as in the proof of Theorem 2 in [3]. Two of these lemmas are stated here without proof. We will provide the full proof beginning with the modification of Lemma 5 from [3].

The total service time associated with a demand is defined to be the onsite service time s plus the incremental travel time between the demand and the next demand to be serviced. Denoting

the distance to be traveled after the jth demand as d_j, the total service time associated with demand j is then $d_j/v + s_j$.

Lemma 3: The average interdemand travel time is related to $E[Z^*]$, the expected minimum distance between any two active demands, according to

$$E[Z^*] \triangleq \lim_{j \to \infty} E[Z^*(j)] \leq \lim_{j \to \infty} E[d_j] \triangleq \bar{d}.$$

Lemma 4: $E[Z^*]$, the expected minimum distance between active demands, is related to the system parameters as follows:

$$\lim_{N \to \infty} \sqrt{N} E[Z^*] \geq \gamma \int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx$$

where $\gamma \geq 2/3\sqrt{\pi}$.

In [3], Lemma 5 is stated as follows:

Lemma 5 From [3]:

$$\lim_{\rho \to 1} W(1 - \rho)^2 \geq \gamma^2 \frac{\lambda[\int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx]^2}{v^2 n^2}.$$

As we are interested in the case with scaling of parameters other than ρ, we instead prove the following result:

Lemma 5: (Lemma 5 From [3] (Modified)): If both $\lambda \mathbb{E}[\sqrt{f}]/vn \to \infty$ and also $\lambda \mathbb{E}[\sqrt{f}]^2/v^2 n \to \infty$, then

$$W \geq \Omega \left(\frac{\lambda[\int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx]^2}{v^2 (1 - \rho)^2 n^2} \right).$$

Proof: Consider the following necessary condition for stability

$$\bar{s} + \frac{\bar{d}}{v} \leq \frac{n}{\lambda}. \tag{49}$$

Using the fact that $E[Z^*] \leq \bar{d}$ from Lemma 3, multiplying the second term on the left hand side above by $\sqrt{N}/\sqrt{N}$ and rearranging implies

$$\sqrt{N} \geq \frac{\lambda \sqrt{N} E[Z^*]}{v(1 - \rho)n}. \tag{50}$$

We show in Lemma 7 below that $N \to \infty$ as both $\lambda \mathbb{E}[\sqrt{f}]/vn \to \infty$ and also $\lambda \mathbb{E}[\sqrt{f}]^2/v^2 n \to \infty$. Therefore, with this scaling, we apply Lemma 4 and

$$\sqrt{N} = \Omega \left(\frac{\lambda \int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx}{v(1 - \rho)n} \right). \tag{51}$$

Squaring both sides of (51) and applying Little's Theorem, $N = \lambda W$, we then have

$$W = \Omega \left(\frac{\lambda[\int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx]^2}{v^2 (1 - \rho)^2 n^2} \right) \tag{52}$$

and the modified lemma is proven. ∎

To complete the proof of Theorem 4, we use the proof of Theorem 2 in [3] as originally written. This proof solves for $\min[\int_{\mathcal{A}} \phi^{-1/2}(x) f(x) dx]^2$ as a function of $f(x)$. Theorem 4 here differs from Theorem 2 in [3] only in the restatement of the limiting terms as in the modified lemma.

To complete our modifed proof, we must show that $N \to \infty$ when both $\lambda \mathbb{E}[\sqrt{f}]/vn \to \infty$ and also $\lambda \mathbb{E}[\sqrt{f}]^2/v^2 n \to \infty$ in the DTRP system. We first prove a preliminary lemma on the scaling of the system workload in a DTRP queue where

workload is defined as in the standard definition of workload in the context of networks:

Definition (Workload): The workload in the system at time t, $V(t)$, is the amount of time it takes the n vehicles to serve all of the messages currently in the system at time t.

To show that the average work in system goes to ∞ as both $\lambda \mathbb{E}[\sqrt{f}]/vn \to \infty$ and also $\lambda \mathbb{E}[\sqrt{f}]^2/v^2 n \to \infty$, we have the following lemma.

Lemma 6: For $\lambda \mathbb{E}[\sqrt{f}]/vn \to \infty$, the average workload in the system V scales as

$$V = \Omega \left(\frac{\lambda \mathbb{E}[\sqrt{f}]^2}{v^2 n} \right).$$

Proof: Assume the vehicle started serving at time $-\infty$. Now consider any time, say 0. Let $V(0)$ denote the amount of workload in the system at time 0. Since time 0 is arbitrary, $V(0)$ is distributed like the stationary distribution of workload. Let $A(s)$ denote the minimal amount of time it takes to serve messages arriving in interval $[-t, 0]$. Then, it is easy to see that

$$V(0) \geq (A(t) - nt)^+. \tag{53}$$

That is, the work in system is greater than difference between the amount of arrived work in an interval of length t and the maximum possible work completed by the n vehicles in the interval. The (53) is true for all t. Further, the time 0 is a randomly chosen time and hence represents the stationary time. Hence, we obtain the time average of workload in the system, $E[V]$, is lower bounded as

$$E[V] \geq E[A(t)] - nt, \, \forall t \geq 0. \tag{54}$$

Thus, to compute lower bound on average workload V, we need to compute $E[A(t)]$. That is, we need to compute the average minimal time required to serve messages arriving to the system in an interval of length t. Let $\Lambda(t)$ be random number of arrivals happening in time interval of length t. Then, $A(t)$ can be lower bounded by the length of shortest path connecting all source and destination locations of these $\Lambda(t)$ messages. The length of a shortest path through a set of locations is no longer than twice the length of the shortest cycle through these points, the TSP tour. Hence, to obtain lower bound $A(t)$, it is sufficient to consider the length of TSP tour through the source and destination location of $\Lambda(t)$ points.

Recall the BHH Theorem of Theorem 6 which bounds the length of a TSP tour. Let L_N denote the length of the TSP tour through N points independently and identically distributed according to probability density $f(\cdot)$. Then, for any $\epsilon > 0$, there exists a $\bar{N}$ such that

$$E[L_N] \geq \beta_{TSP} \sqrt{N} \mathbb{E}[\sqrt{f}] - \epsilon \sqrt{N} \tag{55}$$

where β_{TSP} is a finite positive constant. In particular, choose $\epsilon = (1/2)\beta_{TSP} \mathbb{E}[\sqrt{f}]$. Then, Theorem 5 implies that there exists a $\bar{N}$ such that for all $N \geq \bar{N}$, the following holds:

$$E[L_N] \geq \frac{\beta_{TSP}}{2} \sqrt{N} \mathbb{E}[\sqrt{f}]. \tag{56}$$

We would like to apply (56) to $N = \Lambda(t)$ for t sufficiently large. Note that $E[\Lambda(t)] = \lambda t$. Due to the Poisson property of the

arrival process, $\Lambda(t) \geq \lambda t/2$ with probability at least 1/2 for large enough λ. Therefore, $P(\sqrt{\Lambda(t)} \geq \sqrt{\lambda t/2}) \geq 1/2$ and

$$E[\sqrt{\Lambda(t)}] \geq \frac{1}{2}\sqrt{\frac{\lambda t}{2}}. \tag{57}$$

Assume that t is sufficiently large so that $\lambda t \geq \bar{N}$ and (56) holds. Substituting in (57), we may lower bound A as

$$E[A(t)] \geq \hat{\beta}\frac{\sqrt{\lambda t}E[\sqrt{f}]}{v} \tag{58}$$

where $\hat{\beta} = \beta_{TSP}/8\sqrt{2}$. From (54) and (58), we obtain

$$E[V] \geq \hat{\beta}\frac{\sqrt{\lambda t}\mathbb{E}[\sqrt{f}]}{v} - nt. \tag{59}$$

Consider $t^* = \lambda\hat{\beta}^2\mathbb{E}[\sqrt{f_i}]^2/2v^2n^2$. Note that the condition of the lemma that $\lambda\mathbb{E}[\sqrt{f}]/vn \to \infty$ implies that $\lambda t^* \to \infty$ as required for (56) to hold. Then, from (59) we obtain

$$E[V] \geq \frac{\sqrt{2}-1}{2}\frac{\hat{\beta}^2\lambda\mathbb{E}[\sqrt{f}]^2}{v^2n} \tag{60}$$

for $\lambda\mathbb{E}[\sqrt{f_i}]/vn$ sufficiently large, and Lemma 6 is proven. ∎

Lemma 7: If both $\lambda\mathbb{E}[\sqrt{f}]/vn \to \infty$ and also $\lambda\mathbb{E}[\sqrt{f}]^2/v^2n \to \infty$, the average number in queue $N \to \infty$ as well.

Proof: The first condition of the Lemma 7, $\lambda\mathbb{E}[\sqrt{f}]/vn \to \infty$, implies that Lemma 6 holds. With this lemma, the second condition, $\lambda\mathbb{E}[\sqrt{f}]^2/v^2n \to \infty$, implies $E[V] \to \infty$ as well.

The work associated with each message is upper bounded by the diameter of the region plus the onsite service time, $\sqrt{2}\sqrt{A} + 2\bar{s}(n)$. Therefore, because the average work in the system is going to ∞ and the work associated with each message is finite, the average number of messages in the system, N, must be going to ∞ as well. ∎

APPENDIX B
PROOF OF BATCH QUEUEING TIME

In this appendix, we prove the following lemma, bounding T_Q, the time a batch spends in queue for the Source Only Policy given in Section IV-A.

Lemma 1: For the policy in Theorem 2(b) with batch time $T = \kappa(4\beta^2\lambda A/v^2(1-\rho)^2n)$ for some $\kappa > 1$, the delay of the batch in the queue is bounded by

$$T_Q = O(T). \tag{61}$$

Proof: T_Q may be bounded by using Kingman's bound for the delay in a G/G/1 queue. That is

$$T_Q \leq \frac{\lambda_B(\sigma_A^2 + \sigma_S^2)}{2(1-\rho_B)} \tag{62}$$

where σ_A^2 is the variance of the interarrival times and σ_S^2 is the variance of the service times. In this context, the batch interarrival times are deterministic so $\sigma_A^2 = 0$ and $\lambda_B = 1/T$, the arrival rate of batches. With T fixed as in (39), $\rho_B \leq 1/\kappa$.

Bounding σ_S^2 requires some additional effort. First note that $\sigma_{T_B}^2 = E[T_B^2] - E[T_B]^2 \leq E[T_B^2]$. T_B has two parts: 1)

$2L_{N_T}$ = the interdemand travel times for the pickup and delivery tours and 2) $2N_T\bar{s}(n)$ = the onsite service times for pickup and delivery. Then

$$\begin{aligned} E[T_B^2] &= E[(2L_{N_T} + 2N_T\bar{s}(n))^2] \\ &= 4E[L_{N_T}^2] + 4\bar{s}(n)E[L_{N_T}N_T] \\ &\quad + 4\bar{s}(n)^2E[N_T^2]. \end{aligned} \tag{63}$$

Compute the terms of (63) individually. First

$$\begin{aligned} E[L_{N_T}^2] &= var(L_{N_T}) + E[L_{N_T}]^2 &\text{(64)} \\ &= var(E[L_{N_T}|N_T]) + E[var(L_{N_T}|N_T)] \\ &\quad + E[E[L_{N_T}|N_T]]^2 &\text{(65)} \\ &\leq \beta^2 A var(\sqrt{N_T}) + O(1) + \beta^2 AE[N_T] &\text{(66)} \\ &\leq 2\beta^2 AE[N_T] + O(1) &\text{(67)} \\ &= 2\beta^2 A\frac{\lambda T}{n} + O(1) &\text{(68)} \end{aligned}$$

where (65) uses the formulas for iterated variance and iterated expectation, (66) is by the BHH theorem, Theorem 5, and (67) is by concavity of $\sqrt{\cdot}$.

Next, the second term in (63) is

$$\begin{aligned} E[L_{N_T}N_T] &= E[E[L_{N_T}N_T|N_T]] &\text{(69)} \\ &= \beta\sqrt{A}E[N_T^{3/2}] &\text{(70)} \\ &\leq \beta\sqrt{A}\left(E[N_T^2]\right)^{3/4} &\text{(71)} \\ &= \beta\sqrt{A}\left(\frac{\lambda T}{n} + \left(\frac{\lambda T}{n}\right)^2\right)^{3/4} &\text{(72)} \\ &\leq 2\beta\sqrt{A}\left(\frac{\lambda T}{n}\right)^{3/2} &\text{(73)} \end{aligned}$$

where (71) is by concavity and (73) assumes that $\lambda T/n \geq 1 \Rightarrow \lambda = \Omega(n)$. Note that if $\lambda T/n < 1$, the system is very lightly loaded and a policy based on the worst case TSP may be used to again bound $T_Q = O(T)$ in a similar way, without the variance terms.

The last term in (63) is just the second moment of a Poisson variable

$$E[N_T^2] = \frac{\lambda T}{n} + \left(\frac{\lambda T}{n}\right)^2 \leq 2\left(\frac{\lambda T}{n}\right)^2. \tag{74}$$

Finally, put all of these terms together

$$E[T_B^2] = 8\frac{\beta^2 A\lambda}{n}T + 4\rho\sqrt{\frac{\beta^2 A\lambda}{n}}T^{3/2} + 4\rho^2 T^2 + O(1).$$

Substituting this into (62) above

$$\begin{aligned} T_Q &\leq \left(\frac{8}{(1-\frac{1}{\kappa})^2} + \frac{4\rho\sqrt{\kappa}}{(1-\frac{1}{\kappa})^2} + \frac{4\rho^2\kappa}{(1-\frac{1}{\kappa})^2}\right)\frac{\beta^2\lambda A}{n} + O(1) \\ &= O\left(\frac{\beta^2\lambda A}{n}\right). \end{aligned}$$

Therefore, given the batch time $T = \kappa(4\beta^2\lambda A/v^2(1-\rho)^2n)$, we see that $T_Q = O(T)$. ∎

References

[1] D. Bertsimas and G. van Ryzin, "A stochastic and dynamic vehicle routing problem in the Euclidean place," *Operat. Res.*, vol. 39, no. 4, pp. 601–615, 1991.

[2] D. Bertsimas and G. van Ryzin, "Stochastic and dynamic vehicle routing in the Euclidean plane: The multiple-server, capacitated vehicle case," *Operat. Res.*, vol. 41, pp. 60–76, 1993.

[3] D. Bertsimas and G. van Ryzin, "Stochastic and dynamic vehicle routing with general demand and interarrival time distributions," *Adv. Appl. Prob.*, vol. 20, no. 4, pp. 947–978, 1993.

[4] W. Powell, P. Jaillet, and A. Odoni, "Stochastic and dynamic networks and routing," in *Network Routing*, M. Ball, T. Magnanti, C. Monma, and G. Nemhauser, Eds. Amsterdam, The Netherlands: North-Holland, 1995, pp. 141–295.

[5] M. Gendreau and J.-Y. Potvin, "Dynamic vehicle routing and dispatching," in *Fleet Management and Logistics*, T. Crainic and G. Laporte, Eds. Boston, MA: Kluwer, 1998, pp. 115–126.

[6] M. Savelsbergh and M. Sol, "The general pickup and delivery problem," *Transport. Sci.*, vol. 29, no. 1, pp. 17–29, 1995.

[7] G. Desaulniers, J. Desrosiers, A. Erdmann, M. Solomon, and F. Soumis, "Vrp with pickup and delivery," in *The Vehicle Routing Problem*, P. Toth and D. Vigo, Eds. Singapore: SIAM, 2002, pp. 225–242.

[8] M. Swihart and J. Papastavrou, "A stochastic and dynamic model for the single-vehicle pick-up and delivery problem," *Eur. J. Operat. Res.*, vol. 114, pp. 447–464, 1999.

[9] J. Beardwood, J. Halton, and J. Hammersley, "The shortest path through many points," *Proc. Cambridge Phil. Soc.*, vol. 55, pp. 299–327, 1959.

[10] R. Karp and J. Steele, "Probabilistic analysis of heuristics," in *The Traveling Salesman Problem*, E. Lawler, J. Lenstra, A. R. Kan, and D. Shmoys, Eds. New York: Wiley, 1985.

[11] J. Steele, "Subadditive Euclidean functionals and nonlinear growth in geometric probability," *Annal. Prob.*, vol. 9, pp. 365–276, 1981.

[12] X. Papaconstantinou and D. Bertsimas, "Relations between the pre-arrival and post-departures state probabilities and the FCFS waiting-time distribution for the $e_k/g/s$ queue," *Naval Res. Log. Quart.*, vol. 37, pp. 135–149, 1990.

[13] R. Larson and A. Odoni, *Urban Operations Research*. Englewood Cliffs, NJ: Prentice-Hall, 1981.

[14] R. Wolff, *Stochastic Modeling and the Theory of Queues*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

[15] H. A. Waisanen, D. Shah, and M. A. Dahleh, "Fundamental performance limits for multi-stage vehicle routing problems," *Operat. Res.* pp. 1–24, Aug. 2007 [Online]. Available: http://www.mit.edu/~devavrat/multi-stage.pdf

Holly A. Waisanen received the B.S. degree in applied mathematics and the M.S. degree in systems and control engineering from Case Western Reserve University, Cleveland, OH, in 2003 and the Ph.D. degree from the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (MIT), Cambridge, in 2007.

She was a member of the Laboratory for Information and Decision Systems, MIT. She is currently a consultant with the Boston Consulting Group, Dallas, TX.

Dr. Waisanen received both the NSF Graduate Research Fellowship and the National Defense Science and Engineering Graduate (NDSEG) Fellowship in 2004.

Devavrat Shah received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, Bombay, in 1999 and the Ph.D. degree from the Computer Science Department, Stanford University, Stanford, CA, in 2004.

He is currently an Assistant Professor with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (MIT), Cambridge. He is a member of the Laboratory for Information and Decision Systems (LIDS). He is also a member of the Operations Research Center (ORC). His research interests are in the network algorithms, stochastic analysis of networks, and large scale statistical inference.

Dr. Shah received the President of India Gold Medal in 2004, the IEEE INFOCOM Best Paper Award in 2004, the ACM SIGMETRIC/Performance Best Paper Award in 2006, the 2005 George B. Dantzig Best Dissertation Award from the INFORMS, and the NSF CAREER Award in 2006.

Munther A. Dahleh (F'00) was born in 1962. He received the B.S. degree from Texas A & M university, College Station, in 1983 and the Ph.D. degree from Rice University, Houston, TX, in 1987, all in electrical engineering.

Since then, he has been with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (MIT), Cambridge, where he is now a full Professor. He is interested in problems at the interface of robust control, filtering, information theory, and computation which include control problems with communication constraints and distributed mobile agents with local decision capabilities.

Dr. Dahleh received the NSF Presidential Young Investigator Award (1991), the George Axelby Outstanding Paper Award in 2004, and the Hugo Schuck Award for Theory.