Proceedings of the 45th IEEE Conference on Decision & Control
Manchester Grand Hyatt Hotel
San Diego, CA, USA, December 13-15, 2006

Minimal Delay in Controlled Mobile Relay Networks

Holly A. Waisanen
Department of EECS
Massachusetts Institute of Technology
Email: waisanh@mit.edu

Devavrat Shah
Department of EECS & ESD
Massachusetts Institute of Technology
Email: devavrat@mit.edu

Munther A. Dahleh
Department of EECS
Massachusetts Institute of Technology
Email: dahleh@mit.edu

***Abstract*— We consider a network in which a set of vehicles is responsible for picking up and delivering messages that arrive according to Poisson process with pickup and delivery location distributed in certain region uniformly at random. The vehicles are required to pickup and deliver the messages so that the average delay is minimized.**

In our previous work, we examined the benefits of vehicle control on the message delay. Specifically, we obtained optimal delay scaling where each message was required to be picked up and delivered by the same vehicle. Motivated by application to wireless networks, in this paper we remove this restriction and allow the vehicles to relay message between them. Specifically, we consider two relay methods. The first requires vehicles to relay messages directly to each other using a synchronous rendezvous schedule while the other utilizes an infinite capacity depot to store relayed messages. Under both relay models, we characterize the minimal delay scaling which demonstrates that relaying helps in reducing delay further. Surprisingly, the optimal delay scaling is achieved with only one relay per message.

We note that our results naturally apply to the classical vehicle routing setup as well as to a wireless communication network. Specifically, our results suggest that the delay reduction can be very significant in a controlled relay network .

I. INTRODUCTION

We first present a model for a message passing network, and describe the Dynamic Pickup and Delivery Problem with relays (DPDP with relays) for serving messages in this network. We will then briefly motivate this problem in the context of vehicle routing and wireless communication.

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 $\mathcal{A} = [0, \sqrt{A}]^2$. Each vehicle may move in any direction at any time with a velocity of magnitude $\leq v$. For simplicity, assume $v = 1$.

Messages are generated according to a Poisson process with time intensity $\lambda(n)$. Associated with each message j are source and destination locations denoted by $s(j) \in \mathcal{A}$ and $d(j) \in \mathcal{A}$ respectively. Both the source and destination locations are independently and identically distributed (IID) uniformly in $\mathcal{A}$. The messages need to be picked up from their source locations and delivered to their destination locations by the vehicles. Additionally, the messages may be relayed between vehicles. At the source, relay, and destination locations, the vehicle must spend a fixed onsite service time $\bar{s}(n)$ to pickup or dropoff each 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.

2) Control Policies: A control policy is a set of decision making rules that decides the pickup, relaying, and delivery schedule of arriving messages. We assume that the vehicles have no knowledge of individual messages before they arrive although the overall message arrival process and source and destination distributions are known. Once a message arrives, both its source and also its destination location are known by all vehicles within some region of interest. Vehicles then follow predetermined protocols to service these messages without any real-time communication between vehicles to coordinate their protocols dynamically.

Control policies allow messages to be relayed between vehicles such that the vehicle that initially picks up a message may not be the one to deliver it to its destination. Two types of relaying protocols are examined in this paper. In both policies, we assume that each message is handled by at most two vehicles, one that handles its pickup and the other that makes the final delivery. No other intermediate relays are allowed. First, we consider the *No-Depot* problem in which vehicles meet at some rendezvous points to pass messages, taking $\bar{s}(n)$ time to complete each transmission. Next, we consider the case in which there exists one infinite capacity depot at the center of the region, called the *1-Depot* problem. Vehicles may drop an unlimited number of messages at the depot. Messages remain at the depot until they are picked up at the depot for delivery by other vehicles. Although an unlimited number of messages may be stored at the depot, vehicles are still required to spend $\bar{s}(n)$ time at the depot to pickup or dropoff each message. It is assumed that multiple vehicles may transmit simultaneously to and from the depot while at the depot location. While the *No-Depot* setup is more realistic, we study the *1-Depot* setup to contrast the difference in performance when such Depot is available (we will find that performance does not change in

1-4244-0171-2/06/$20.00 ©2006 IEEE.

terms of scaling).

3) Performance Metrics: The delay of message m, denoted $W(m)$, 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

$$W = \limsup_{m\to\infty} \mathbb{E}[W(m)]. \tag{1}$$

If $W(\cdot)$ has a 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 $W < \infty$. A necessary condition for the existence of a stable policy is $4\rho = 4\frac{\lambda(n)\bar{s}(n)}{n} < 1$ where the constant 4 reflects the fact that each message is transmitted 4 times, once for pickup, twice for relaying, and one final time for delivery.

4) Problem Statement: Design a valid control policy for each vehicle that decides the pickup and delivery schedule of arriving messages such that: (a) The messages are delivered at rate $\lambda(n)$ by the n vehicles collectively, (b) each message is served by at most two vehicles, one for pickup and one for delivery, and (c) the average message delay is minimized, where delay of a message is the time it takes to reach destination location from time of its arrival. We will call the above defined control problem as the Dynamic Pickup and Delivery Problem with relays (DPDP with relays).

B. Relation to Previous Work

The DPDP is naturally related to the problem of dynamic vehicle routing (VR). For example, messages may be considered to be passengers which need to be taxied from one location to another, or packages that need to be couriered around town. The significant theoretical results on VR were first obtained by Bertsimas and Van Ryzin in a series of papers [1], [2] where they considered the special case of VR problem known as Dynamic Travel Repairperson Problem (DTRP). In DTRP, requests arrive at various locations in a region and there are multiple vehicles (or repairperson) available for servicing them. The goal is to service them so as to minimize the average delay. The DTRP problem is certainly a special case of question considered in this paper as we require vehicles to both pickup and then deliver.

In VR literature, special cases of DPDP problem has been studied under names such as Online Dial-a-Ride (OLDARP) and the single vehicle DPDP [3], [4]. These works examine systems with only a single vehicle, whereas we will consider the multi-vehicle case here. We have also previously studied the DPDP for the case in which messages were not allowed to be relayed between vehicles [5], [6]. Comparisons to this work will be make in Section IV. We will also discuss relation to the wireless network application of our results in Section IV.

C. Main Results

We obtain delay (order) optimal control policies for both *No-Depot* and *1-Depot* relay model. The policies are based on clever batching and their analysis is based on deriving bounds on worst-case length of Traveling Salesman Problem as well as use of bounds on average queue-sizes from classical queuing theory. We state the main results, which are proved in Sections II and III respectively [1].

Theorem 1 (No-Depot): If no depots exist and messages may be transferred only between co-located vehicles, there exists a policy such that the average delay scales as

$$W_{ND} = O\left(\frac{\lambda(n)A}{(1-(4+\varepsilon)\rho)^2 n^2}\right) + O\left(\frac{n\sqrt{A}}{1-(4+\varepsilon)\rho}\right)$$

for arbitrarily small $\varepsilon > 0$.

Theorem 2 (1-Depot): In the case that there exists at least one infinite capacity depot, there exists a policy such that the average delay scales as

$$W_{1D} = O\left(\frac{\lambda(n)A}{(1-(4+\varepsilon)\rho)^2 n^2}\right) + O\left(\frac{\sqrt{A}}{1-(4+\varepsilon)\rho}\right),$$

for arbitrarily small $\varepsilon > 0$.

The optimality of above stated Theorems 1 and 2 follows from results of Bertsimas and Van Ryzin [1], [2] on DTRP. To see this, note that in the setup of the DTRP each message needs to be only serviced at its arrival location. Since our DPDP requires both pickup and delivery, the lower bound on average delay for DTRP naturally implies the lower bound on the delay for DPDP problem. Under the DTRP setup, the average delay of demands in the system is $W_{DTRP} = \Omega\left(\frac{\lambda A}{n^2}\right)$ [2]. This establishes optimality of Theorems 1 and 2.

D. Organization

The rest of this paper is organized as follows. In Section II, we then prove Theorem 1 for the No-Depot policy. In Section III, we prove Theorem 2 for the 1-Depot policy. In Section IV, we contrast our results with the known results and our previous work as well as discuss applicability of our work in the context of wireless networks.

II. No-Depot Policy

In this section, we upper bound the average delay experienced by a message in the No-Depot policy, thus proving Theorem 1. The lack of any relaying infrastructure, such as a Depot,

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

requires vehicles to schedule rendezvous with each other in order to relay messages in a distributed manner. This presents a significant challenge in designing simple and delay-optimal schemes.

No-Depot Policy. The region is divided into a $\sqrt{n} \times \sqrt{n}$ grid of cells, each of area A/n and each vehicle is assigned to a distinct cell. Each vehicle is responsible for all of the pickups and deliveries of messages originating in or destined to that cell. Each vehicle maintains n source-destination queues of messages, one for each of the n cells in which the destination locations of arriving messages may occur. When the vehicle is ready to begin a batch service, it takes up to the first N_n messages from each of the n queues to form a batch. The total number of messages to be serviced in each batch is then at most $N_T = nN_n$.

The batch service begins with a TSP tour through the pickup locations of the messages in the batch. The vehicle serves as many messages along the TSP tour as it can in time T_{TSP}. By choosing T_{TSP} appropriately (see Lemma 1), all N_T may be picked up within this interval. Since there is no depot at which to deposit messages for pickup by their delivery vehicles, vehicles must meet each other directly to perform the hand-off. For this, a pre-determined synchronous schedule is used such that each vehicle meets up with every other vehicle during each batch service time to hand off the appropriate messages (see Lemma 2). The rendezvous points at which the vehicles meet are predetermined and are distributed throughout the region. Once these inter-vehicle meetings are complete and all messages to be delivered have been received, another TSP tour constrained by T_{TSP} is performed through the destination locations of the set of messages just received from other vehicles. This completes the batch of time-length T (chosen appropriately so as to allows for all the above events to happen in that time).

Proof: [Theorem 1] To establish validity of the above described policy as well as analyze its performance, we need the following two Lemmas. The first lemma establishes existence of a synchronous schedule for rendevous between vehicles and the second lemma establishes a worst case bound on TSP tour.

Lemma 1: Given n vehicles, there exists a schedule of length n such that each vehicle visits all other $n-1$ vehicles at least once.

Proof: Consider a complete bipartite graph of $2n$ nodes, where each vehicle is represented by one node on the left and one on the right. An edge between node i on left and node j on right represents the requirement that vehicle i must meet vehicle j.

Now color the edges of this graph such that no two edges connected to the same node have the same color. By assigning color k to the edge between vehicle i on the left and vehicle $(i+k) \mod n$ on the right, this may be accomplished using n colors. The schedule is then constructed by letting each color represent a time slot in which the two vehicles are assigned to meet and transfer messages. The rendezvous point for the vehicle pairing represented by each edge may be taken to the be the center of the cell assigned to the vehicle on the left hand side in the graph. ∎

Lemma 2: Given N locations arbitrarily located in a square region of area B, there exists a tour through these points of length at most $2\sqrt{2NB}$.

Proof: First note that if $N = 1$, the total time to visit the location and then return to the starting point, starting from anywhere in the region, is at most $2\sqrt{2B}$, so the bound in the theorem holds. The following is for $N \geq 2$.

Divide the region into N cells of area B/N. Consider a tour that begins at an arbitary location, then travels directly to the center of the upper-leftmost cell. The tour then travels between the centers of all the cells in a row-by-row manner, working across and then down through the region. Once all cell centers have been visited, the vehicle returns to the starting point to complete the tour. Such a tour through the cell centers takes at most time $N\sqrt{\frac{B}{N}} + \sqrt{2B} = \sqrt{(N+2)B} \leq \sqrt{2NB}$ for $N \geq 2$.

The tour through the N arbitrarily located points is performed by following the cell tour above, but stopping in each cell to visit all of the required locations that are located within that cell. To visit each location, the vehicle travels from the cell center to the the location and then back to the cell center. Each of these visits takes at most $\sqrt{2B/N}$. Since there are N locations to visit in this way, the location visits take a total of at most $\sqrt{2BN}$ in addition to the cell tour.

Combining this with the cell tour length above, the total tour through the N locations takes at most $2\sqrt{2NB}$.

∎

Next, we use these lemmas to first obtain appropriate values of T, N_T so that all the arriving messages are eventually delivered to their destinations and we will evaluate the induced delay. Given any $\varepsilon > 0$, let

$$N_T = \frac{(1+\varepsilon)\lambda(n)}{n} T.$$

Let T_{TSP} be the worst-case travel time it takes to tour-through pickup or delivery locations of N_T messages in cell of area A/n. Then by Lemma 2

$$T_{TSP} \leq \sqrt{\frac{8N_T A}{n}} = \sqrt{\frac{8(1+\varepsilon)\lambda(n)A}{n^2}}.$$

We will assume that message requires $2\bar{s}(n)$ amount of time to be transferred from one vehicle to the other vehicle during relaying operation. Now, the total travel-time for synchronous schedule created by Lemma 1 is $n\sqrt{A}$. The total time to exchange the messages during the rendevous of vehicles is $N_T\bar{s}(n)$. Hence, the total batch time T can be bounded above as

$$\begin{aligned} T &\leq 2T_{TSP} + 4N_T\bar{s}(n) + n\sqrt{A} \\ &= \sqrt{\frac{32(1+\varepsilon)\lambda(n)A}{n^2}} + 4(1+\varepsilon)\rho T + n\sqrt{A}. \quad (2) \end{aligned}$$

From (2) and some manipulation will lead to the conclusion that it is sufficient to have T such that

$$T = O\left(\frac{\lambda(n)A}{(1-4(1+\varepsilon)\rho)^2 n^2}\right) + O\left(\frac{n\sqrt{A}}{(1-4(1+\varepsilon)\rho)}\right). (3)$$

Note that in time T, in a given cell $\lambda(n)T/n^2$ messages arrive that are destined for any other cell. In the above described scheme and the selection of T (as in 3), each vehicle serves upto $N_n = (1+\varepsilon)\lambda(n)T/n^2$ messages for a given pair of cells. Thus, we have a service rate higher than the arrival rate and hence by standard queueing argument, it must be stable. Next, we compute the average delay per message in this scheme.

Lemma 3: The delay experienced by a message under the above described policy and selection of T, N_T is $W_{ND} = O(T)$.

Proof: To this end, note that each message has the following types of delays: (a) waiting to get serviced in a cell after arrival and (b) the batch time T. Now, the T is bounded above as (3). To bound (a), note that messages are queued separately depending on their destination cells. Consider one particular queue for a destination cell. The arrivals to this queue happen at rate $\lambda(n)/n^2$ while every T units of time, $N_n = N_T/n$ of them get served. Delay through this queue can be upper bounded by T plus the delay through an M/D/1 queue with arrival rate $\lambda(n)/n^2$ and deterministic service requirement of $T/N_n = \frac{n^2}{(1+\varepsilon)\lambda(n)}$. It is well-known that average delay of such an M/D/1 queue is upper bounded as $O(1/\varepsilon)$, which is a constant. Thus, effective delay for (a) is $O(T)$. That is, the average delay experienced by message between arrival and delivery is $O(T)$. ■

■

III. 1-Depot Policy

In this section, we present a policy utilizing one depot and upper bound the average delay experienced by a message under this policy, thus proving Theorem 2.

1-Depot [2] The region is divided into a $\sqrt{n} \times \sqrt{n}$ grid of cells, each of area A/n and each vehicle is assigned to a distinct cell. Each vehicle is responsible for all of the pickups and deliveries of messages originating in or destined to that cell. Each vehicle maintains n source-destination queues of messages, one for each of the n cells in which the destination locations of arriving messages may occur. When the vehicle is ready to begin a batch service, it takes up to the first N_n messages from each of the n queues to form a batch. The total number of messages to be serviced in each batch is then at most $N_T = nN_n$.

The batch service begins with a TSP tour through the pickup locations of the messages in the batch. The vehicle serves as many messages along the TSP tour as it can in time T_{TSP}. By choosing T_{TSP} appropriately (see Lemma 1), all N_T can be picked up within this interval. These picked up messages are then dropped off at the Depot and vehicle picks up upto N_T messages from Depot that are destined to its cell. The vehicle performs TSP tour along the destination locations of these N_T messages in its cell in time at most T_{TSP}. This completes the batch of time-length T (chosen appropriately so as to allows for all the above events to happen in that time).

Proof: [Theorem 2] The proof of Theorem 2 is identical to that of proof of Theorem 1 with the only difference that the travel time for rendvous in No-Depot policy was $n\sqrt{A}$ while that for the 1-Depot case is $2\sqrt{A}$. Replacing the $n\sqrt{A}$ by $2\sqrt{A}$ in all the analysis, we obtain that it is sufficient to have T such that

$$T = O\left(\frac{\lambda(n)A}{(1-4(1+\varepsilon)\rho)^2 n^2}\right) + O\left(\frac{\sqrt{A}}{(1-4(1+\varepsilon)\rho)}\right). (4)$$

Further, under such batch policy the per message delay is $O(T)$. This completes the proof of Theorem 2. ■

IV. Discussion

We have presented two policies for the DPDP problem with relays. The policies are based on clever batching scheme and distributed synchronous scheduling for rendevous of vehicles. The analysis involved obtained using worst-case bounds on Euclidian TSP. The results of Bertsimas and Van Ryzin imply the optimality of these policies. It is worth noting that essentially our schemes, with the help of single relaying, decomposes the problem of DPDP into constant number of DTRP problems. This is the reason of our policies being (order) optimal.

In contrast to these results, in our previous work in the DPDP problem, we considered the case in which relaying between vehicles was not allowed and instead each message was required to be picked-up and delivered by the same vehicle. We had established that under such restriction, the optimal delay was shown to be $W = \Theta(\lambda A/n^{3/2})$ with $\lambda(n) = \Omega(n^{3/2})$ [6]. Thus, only one relay significantly reduces the delay performance.

A. Connections to Wireless Communication Network

An important application of the DPDP problem is in the context of wireless communication networks. In this setup, vehicles (or mobile nodes) provide communication infrastructure for messages to be transmitted from their sources to destinations. Recently, there has been an exciting development

[2] This policy was prompted by a discussion with Emilio Frazzoli.

to study the throughput capacity and delay in such wireless ad-hoc networks [7] [8]. Most of the previous work assumes some form of random mobility model for wireless nodes for one of the following reasons: (a) node mobility is not predictable when nodes are uncontrolled or (b) lack of framework for studying controlled setup.

In this paper, we have considered the setup where vehicles are deployed primarily to provide wireless communication infrastructure. In such a setup, it makes sense to control their movement so as to utilize the maximal network capacity (i.e. maximal throughput) while minimizing the delay. In wireless network, relaying is quite natural in order to minimize delay. Specifically, our policies based on single relay for DPDP problem suggests the optimal order performance in terms of delay for maximal (in terms of order) throughput for such wireless networks.

The key distinguishing feature of wireless communication is the interference caused by any transmission for other simultaneous transmissions. Any communication policy, such as our policies, needs to account for such interference. In our description of our policy, specifically the No-Depot policy that is very relevant to wireless communication setup, we have ignored the issues of interference. However, the structure of the policy can easily allow to handle the interference as follows. The Protocol model, primarily introduced by Gupta and Kumar [9] to study throughput capacity of wireless networks under interference suggests that transmission in from a node A to node B is successful if (i) node B is within transmission radius r from node A, and (ii) no other simultaneously transmitting node, say C, is within distance $(1+\delta)r$ of B for some constant $\delta > 0$. Next, we provide brief sketch of how to modify the No-Depot policy so as to account for the effect of the interference.

To this end, let wireless transmission radius of each node be set equal to the $\sqrt{\frac{2A}{n}}$, the longest distance between any two nodes in a cell under No-Depot policy. Then, any two nodes within a cell can successfully comunicate if no node among their 6 neighboring cell is transmitting according to the Protocol model (with small enough $\delta > 0$). Using a simple time-division scheme, it is possible that each cell gets to transmit fraction $1/7$ of the time. That is, if we reduce $\bar{s}(n)$ by $1/7$ to that considered in the DPDP setup, then effectively our policy for DPDP problem is the same as long as we make sure that all communications are happening within a cell only. Guaranteeing this is straightforward in the No-Depot case as all pickup and delivery happen when vehicle is inside the same cell as the pickup-delivery locations; the rendevous locations are selected such that each vehicle pair that is to exchange data is in separate cell from all other nodes. Thus, we have shows that the No-Depot policy works under Protocol model of wireless network with the same order of throughput and delay guarantees. This establishes the optimal delay scaling for controlled wireless network under Protocol model.

It would be an interesting challenge to extend such results for wireless communication model beyond the Protocol or disk model.

References

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

[2] ——, “Stochastic and dynamic vehicle routing in the euclidean plane: the multiple-server, capacitated vehicle case,” *Oper Res*, vol. 41, pp. 60–76, 1993.

[3] E. Feuerstein and L. Stougie, “On-line single-server dial-a-ride problems,” *Theo Comp Sci*, vol. 268, pp. 91–105, 2001.

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

[5] H. A. Waisanen, D. Shah, and M. A. Dahleh, “Control reduces delay in mobile wireless networks,” in *43rd Annual Allerton Conf. on Communcation, Control, and Computing*, 2005.

[6] ——, “Optimal delay in wireless networks with controlled mobility and maximal throughput,” in *17th International Symposium on Mathematical Theory of Networks and Systems*, Kyoto, Japan, 2006.

[7] M. Grossglauser and D. Tse, “Mobility increases the capacity of wireless networks,” in *Proc. IEEE INFOCOM*, Anchorage, Alaska, 2001, pp. 1360–1369.

[8] A. E. Gamal, J. Mammen, B. Prabhakar, and D. Shah, “Throughput-delay trade-off in wireless networks – part i: The fluid model,” *Special joint issue of IEEE Transaction on Information Theory and ACM Transaction on Networking*, 2006.

[9] P. Gupta and P. Kumar, “The capacity of wireless networks,” *IEEE Trans. on Info Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.