

Deficit Round-Robin Based Message Ferry Routing

Ahmed Mansy, Mostafa Ammar and Ellen Zegura

School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332

{amansy, ammar, ewz}@cc.gatech.edu

Abstract—Message Ferrying is a mobility assisted scheme in which a special node, called a message ferry, is tasked with delivering data among a set of disconnected wireless nodes. One key challenge for such scheme is to design the ferry route in a way that improves certain network characteristics such as average data delivery delay or data loss ratio. Previous work has optimized ferry travel time, rather than directly optimizing message delivery performance metrics. In this paper, we revisit the basic ferry route design problem for stationary nodes with the goal of providing a framework for optimizing delay performance. We start with a Markovian Decision Problem (MDP) formulation which produces the optimal ferry route that minimizes the average data delivery delay. While this formulation, in principle, enables optimal ferry route design, it is numerically intractable for moderate to large size problems. Solutions to small problems, however, yield insight into the properties of optimal ferry routes. These insights, in turn, lead us to propose a ferry route design algorithm that takes advantage of the similarity between our problem and the link scheduling problem in traditional networks. Our algorithm is inspired by the Deficit Round Robin (DRR) algorithm which has provable properties for delay optimization when applied to link scheduling. Using simulations we show that our DRR-based algorithm produces ferry routes that are close to optimal when compared to the MDP-derived solutions for small problems. Our results also show that our algorithm produces ferry routes for moderate to large problems that significantly out perform existing solutions.

I. INTRODUCTION

Message ferrying is a networking paradigm that was developed to perform data routing in intermittently connected or completely disconnected wireless networks [1]–[3]. In this scheme, a set of mobile nodes called *message ferries* take responsibility for carrying messages among disconnected nodes. This approach is very attractive especially in cases when nodes are sparsely deployed with limited contact opportunities. It is also the only feasible connectivity approach when nodes are stationary and sparsely deployed.

A key challenge when using the Message Ferrying scheme is the design of the route of the *message ferries* in a manner that achieves desirable performance characteristics such as minimizing the average end-to-end delay or the packet drop rate. Previous work proposed solutions to this problem in cases where single or multiple ferries were used, and also in cases where nodes were mobile or stationary. A special case of the ferry route design problem was studied in [4]. In this case all nodes were assumed to be stationary and only one *message ferry* was used. The authors in [4] proved that the ferry route design for this case is NP-hard by reducing the Euclidean Traveling Salesman Problem (TSP) to the ferry route design problem. They also proposed a two-phase solution

to the problem. In the first phase they generated an initial route by using an algorithm from the well studied TSP literature. In order to meet certain bandwidth requirements, they applied a second phase where they extended the amount of time the ferry spends at each node by solving a linear optimization problem. This solution and other solutions like [5], [6] are commonly based on graph traversal algorithms and mainly approach the problem as a variation of the TSP. Although this approach is appealing because of the richness of the TSP literature, it does not actually allow for optimizing specific network characteristics such as the average data delivery delay. This is mainly because this approach does not take into consideration the non-uniformity of distances and traffic rates between different nodes in the network. For this reason, we see the need for developing a new ferry route design algorithm that has the capability of optimizing given performance metrics.

To this end, we revisit in this paper the basic problem of ferry route design for stationary sparsely deployed wireless nodes which was originally studied in [4]. Our goal is to develop a scheduling algorithm geared towards minimizing end-to-end message delays. First, we formulate the problem as a Markovian Decision Problem (MDP), and then solve the MDP to get a route for the ferry that minimizes the average end-to-end delay. One major problem we encountered with using MDPs was the exponential growth of the state space with the size of the network. Because of this we were able to get optimal ferry route solution only for small size networks. However, studying these solutions provided us with insights regarding the properties of the structure of optimal ferry routes. Guided by these insights, we next propose a ferry route design algorithm that, takes advantage of the similarity between our problem and the link scheduling problem in traditional networks. As a preview, the ferry acts as the shared resource needed for delivery, just as the link is the shared resource in link scheduling. Our algorithm is based on the Deficit Round Robin (DRR) algorithm which has provable desirable throughput and delay properties in traditional networks.

To illustrate the difference between our approach and previous solutions we present the following example. Figure 1 shows two different geographical configurations for a four-node network. In network Figure 1(a), the nodes form a square whose side is 1KM long. Each node sends data to each other node at a rate of 10Kbps, except for traffic going between nodes 1 and 3 which is 1Mbps in both directions. In this case, intuitively we think that in order to minimize the average data delivery delay, the optimal route for the ferry should visit nodes one and three more often than nodes two and four.

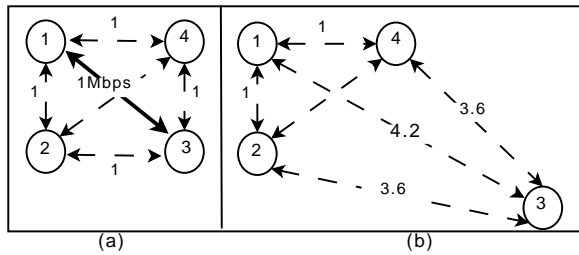


Fig. 1. Two different networks, distance in kilometer

However, a TSP based solution such as [2] will produce a cyclic route of the form $\{1, 2, 3, 4\}$. Since the new technique developed here generally does not produce cyclic routes, we use relative frequencies of node visits to describe the ferry route. In this example, our technique will produce a route that will have the following relative node visit frequencies $\{2 : 1 : 2 : 1\}$, which means that on average node 1 gets visited twice as much as node 2. This route gives about 55% improvement in the mean data delivery time over the TSP-based route.

Another case is illustrated in Figure 1(b), where three of the four nodes have the same placement as the previous example, while the fourth node (node 3) is placed significantly farther away, see distances in figure. The traffic among all nodes is equal to 100Kbps. Again one should expect that the optimal route in this case should visit node 3 less often than the other three nodes because as the ferry travels to node 3 it increases the delay for other traffic. A TSP based solution will again produce a cyclic route of the form $\{1, 2, 3, 4\}$. Our technique, however, produces a route with the following relative visit frequencies $\{3 : 2 : 1 : 2\}$ which gives about 76% improvement in average data delivery delay over the TSP solution when tested through simulations.

Message ferrying was first introduced in [1], [4] as a network paradigm for data delivery in sparse MANETs. The same authors proposed an algorithm for controlling the mobility of multiple message ferries in [2]. An algorithm for ferry route design with mobile nodes in sparse DTNs was proposed in [3]. Different than previous ferry algorithms, this algorithm assumed no coordination between the nodes and the ferry. As mentioned previously, the most relevant work to ours is [2], in which the authors presented a solution to the same problem of ferry route design with stationary nodes and a single ferry.

The rest of the paper is organized as follows. Properties of routes learned from MDP formulation are presented in section II. In section III we present our DRR-based route design framework. We present some simulation results in section IV and then we conclude the paper in section V.

II. MDP FORMULATION

Due to space limits we omit the details of the MDP formulation here and the reader is referred to [7] for the details. Below we present the properties of ferry routes we learned from solving the MDP.

We were able to solve the MDP and get the optimal ferry route for small size networks (3 or 4 nodes) with limited size buffers. An optimal route can be thought of as the set of pairs (*state*, *decision*), where *state* represents the state of the system and *decision* is the next node to visit. Basically, this pair means that if the system is in state *state*, the ferry should move to node *decision*. When applying optimal routes to ferry movement, it does not necessarily result in cyclic routes. Hence, we choose to characterize optimal ferry routes with relative frequencies of node visits.

We obtained the optimal ferry routes for many small networks that have different distance and traffic configurations (including networks in Figure 1). Studying relative frequencies of node visits for these routes, we came to the following observations:

- Nodes that have higher traffic rates (in, out, or both) tend to have higher relative frequency than nodes with lower traffic rates. This observation means that a node that sends or receives more data gets to be visited more frequently than other nodes.
- Optimal routes tend to favor travelling shorter distances, meaning that nodes that are located relatively far from other nodes usually have lower relative frequencies. On the other hand, nodes that are closer to other nodes have higher relative frequencies.

These observations about the relationship between relative frequencies of node visits and both distances and traffic rates were reminiscent of the fair link scheduling problem. In this problem, it is required that packets from different flows that share the same link be scheduled in a way that achieves fairness among flows. Deficit Round-Robin (DRR) [8] is an algorithm that produces fair link schedules. With DRR scheduling, flows with large packets tend to be serviced less frequently than flows with small packets. We explore this analogy a little further in the next section where we use it to develop the DRR-based ferry scheduling heuristic.

III. DRR-BASED FERRY ROUTE SCHEDULING

One key observation in this paper is that the ferry route design problem is analogous with the link scheduling problem in traditional networks. This, in principle, allows us to use any link scheduling algorithm to solve our problem [9]–[11]. We choose to use the Deficit Round-Robin [8] algorithm because of its simplicity and desirable provable throughput and delay properties. In this section we describe the ferry route design process. We start by briefly describing the DRR operation in traditional networks, then we show the analogy between the link scheduling problem and the ferry route problem, and finally we describe our algorithm in detail.

A. Deficit Round-Robin Scheduling and Message Ferry Routing

Deficit Round Robin (DRR) is a scheme that allows fair scheduling of multiple flows over a single shared link. It is designed to deal specifically with the case of flows with

unequal and unpredictable packet lengths and does not require a-priori knowledge of average packet sizes for flows.

The details of DRR can be found in [8]. A router that implements DRR has a separate queue for each incoming flow. For each queue there is a deficit counter that keeps track of the amount of service granted to that flow. Before the DRR algorithm starts, all deficit counters are initialized to zero, and in each round a flow is allocated a quantum Q_i worth of bits. In the simplest case $Q_i = Q_j$ for all flows i, j , which will result in equal shares of the link bandwidth for all flows. However, each flow can be assigned a different relative bandwidth by having different values of Q_i . Queue i 's allocation is accumulated in a deficit counter DC_i . When a queue is visited one or more packets are serviced with the condition that their total size is less than the accumulated deficit counter. If B_i bytes are serviced then the deficit counter is decremented by B_i bytes and the scheduler moves on to the next queue.

It is straightforward to see an analogy between link scheduling using DRR and message ferry route design. Observe that the ferry itself represents a resource (analogous to a link) shared among a set of stationary nodes (analogous to flows). The distance between the current location of the ferry and a particular node is analogous to the size of the packet waiting in a flow's queue – the distance should be traversed only if the node has accumulated enough credit. A ferry visit to a node can be considered in this analogy equivalent to serving a packet from a certain flow.

To illustrate this analogy, we apply it to the network in figure 1(a). In this case we have four “flows” labeled 1 through 4, one representing each node. If we assume the ferry to be currently at node 1, it means that “flow” 1 has just been served and the DRR algorithm will have to decide which “flow” to serve next. That also means that packet lengths coming from flows (nodes) 2, 3, 4 are represented by their distances from node 1 which are 1, $\sqrt{2}$, 1 respectively.

We emphasize that this is only an analogy that inspires our heuristic and is not a formal mapping between the link scheduling problem and the message ferry route problem. It is, nevertheless, useful as it provides guidance for the construction of an efficient heuristic. Furthermore, there are enough degrees of freedom in the DRR scheduling framework that allows us to incorporate ferry route design features that were learned from our MDP formulation.

B. Ferry Route Scheduling Heuristic

Figure 2 shows the high level pseudo-code for the DRR-based route scheduling process. The details of the subroutines are explained in text. The function `getNextLocation` is called by the ferry each time it concludes a visit to a node and needs to find its next destination. The operation of our algorithm can be explained as follows. As in traditional DRR, we maintain a service quantum counter for each stationary node i , we call it $credit[i]$. Initially, $credit[i]$ is set to zero for all i . Each time the ferry needs to make a decision it computes a set of nodes that can potentially be the next

destination, we call this set the *candidate_list*. Function `getDRRCandidates` computes this list. It simply returns all nodes i that satisfy the condition $credit[i] \geq d(curr_loc, i)$, where the latter is the distance between node i and the current position of the ferry. The ferry then applies a tie breaking algorithm to select a destination from the *candidate_list*. After that, the credit value of the selected node i , $credit[i]$, is reset to zero. The tie breaking algorithm will be explained in detail later.

When the function `getDRRCandidates` returns an empty candidate set, a new round of the algorithm operation begins. At this point, the function `updateCredits` is called to update the values of $credit[i]$. The approach we follow to update credit counters will be explained shortly.

getNextLocation (credits):

```

if getDRRCandidates(curr_loc)  $\neq$   $\emptyset$  then
    candidate_list = getDRRCandidates(curr_loc);
    next_location = tieBreak(candidate_list);
    credit[next_location] = 0;
    return next_location;
else
    /* Start a new round */
    updateCredits(credits);
    getNextLocation(credits);
end if

```

Fig. 2. `getNextLocation` function

getDRRCandidates (curr_loc):

```

candidate_list = {}
for all node  $i \neq$  curr_loc do
    if credit[i]  $\geq$  distance(curr_loc, i) then
        candidate_list += i;
    end if
end for

```

Fig. 3. `getDRRCandidates`, computes the *candidate_list*

It is clear that the algorithm operation can be controlled by the two functions **updateCredits** and **tieBreak**. Using different approaches in these two functions can change the performance of the algorithm significantly. We have considered many possible implementations for these two functions. Below we present the functions that we finally settled on for our heuristic, and the rationale behind choosing them.

updateCredits. This function is called at the beginning of each new round k . It generates the quantum service values Q_i^k for each node i and updates each node's credit counter $credit[i]$ by adding Q_i^k to the counter's old value. We compute the values Q_i^k as follows. Let r_{ij} be the traffic rate going from node i to node j , and s_i be the total aggregated traffic rate originating from and destined to node i , then s_i can be written as

$$s_i = \sum_{j=1}^N (r_{ij} + r_{ji})$$

We also define $s = \min_i(s_i)$ as the minimum aggregated data rate for all the nodes. Now, let d_{ij} be the distance between nodes i and j , and let $d_i = \sum_{j=1}^N d_{ij}/(N-1)$ be the average distance between node i and all other nodes. If we assume that the ferry was at node i at the beginning of round k , then we compute the quantum value Q_j^k for node j using the following formula

$$Q_j^k = d_i \times \sqrt{\frac{s_j}{s}}$$

The reason behind making the credit quantum directly proportional to the average distance is to ensure a candidate set that is big enough to give the ferry flexibility in choosing its next destination. This makes the credit increase consistent whether a node is close or far while favoring nodes whose distance from the current node is below average. Due to space limits we refer the reader to [7] to see the rationale behind the square root in the above equation.

tieBreak. This function is called by the ferry to select one of the candidate nodes as the next destination. To make this decision, we go through multiple levels of filtering, reducing the size of the candidates set after each level. In the first filter, we select the w least recently visited nodes by the ferry. The ferry can easily do that by keeping track of the last time it visited each node. In our implementation we chose $w = N_{candidates}/2$ where $N_{candidates}$ is the size of the candidates set. In the second filter we select half of the nodes that are closest to the current position of the ferry. In the third stage, if needed, we select the node with the maximum out traffic rate r_i where $r_i = \sum_{j=1}^N r_{ij}$. If multiple nodes have the same maximum value R_i we select the one that has more messages in the ferry buffer, and finally if we still have a tie we randomly choose one of the nodes.

Using these levels of filtering gives us significant flexibility in choosing the node from the candidate set that satisfies multiple properties at the same time. For example, insight gained from our MDP solution indicates preference for the selection of least recently visited nodes, and this is represented by the first filter. The same applies for the next filters where we chose the nodes that are closer to the ferry than the nodes with higher traffic rates, and so on. We have tried to change the order of these filters and even tried other filters, but the above described set of filters gave us the best results when tested through simulations.

IV. SIMULATION RESULTS

In this section we evaluate the performance of our DRR-based ferry route design algorithm and compare it to the performance of the TSP-based solution in [2], [4]. We use the mean of data delivery delay as our performance metric. We perform the evaluation under different node placement and traffic patterns.

A. Simulation setup

We used the ONE simulator [12] as our simulation environment. We use a network of 20 stationary nodes distributed

in a $4km \times 4km$ area. There is one additional mobile node which acts as a message ferry. Every node has a wireless interface with a radio range of 250 meters and a bandwidth of $10Mbps$. The ferry is assumed to move in straight lines and at constant speed of 15 m/sec. All stationary nodes are strictly configured to communicate only with the ferry, which means if two stationary nodes fall in radio range of each other they do not communicate or exchange data. Each point in the plots in this section was averaged over 36 simulation runs.

Nodes can be either uniformly placed in the area or clustered. Random uniform allocation is simply done by selecting both x and y coordinates of any node as a uniform random variable in the interval $(0, 4000)$. Clusters are formed in the following manner. Assume we have M clusters, then we first select M random cluster centers uniformly distributed over the area. We do that in the same way we did with the uniform random nodes. Then, for each of the N nodes, we first assign a node randomly to any of the clusters and then compute its x and y coordinates using the formula $x = X_c + Z * 4000 / \sqrt{M * N}$ and $y = Y_c + Z * 4000 / \sqrt{M * N}$ where $Z \sim N(0, 1)$, and (X_c, Y_c) are the coordinates of the cluster center.

Message arrivals follow a Poisson process in our simulations. If a flow sends traffic at a rate of w bps then the corresponding Poisson arrivals flow will have a rate of w/msg_size , where msg_size is the message size in bits. We assume that the ferry and all nodes have infinite buffers and that messages have an infinite TTL (Time-To-Live).

We use both uniform and non-uniform traffic distributions. In uniform traffic, every node sends traffic to every other node with the same data rate. In non-uniform traffic distribution, certain flows have higher data rates than others. We randomly select 10% of the flows as high data rate flows and keep the rest of the flows sending at rates similar to the uniform traffic case. In all the experiments below, high data rate flows have rate of 500Kbps. Note that a flow is defined here by a $(source, destination)$ pair.

B. Uniform node allocation

We compare the performance of the DRR-based algorithm with the TSP-based solution in [2]. We perform this study using uniform and non-uniform traffic distributions.

Uniform traffic. Figure 4 shows the mean of data delivery delay versus node traffic rate. Our first observation is that our solution and the TSP-based solution have very similar performance at low data rates (upto 100 kbps), with the TSP-based solution being slightly better. When node traffic increases to 150 Kbps, we observe a significant improvement in our scheme over the TSP solution. The improvement is approximately 80% in the mean delivery delay. Our explanation for that behavior is as following. In the case of uniform traffic, queues at nodes fill up at the same rate with messages waiting to be sent. The TSP-based solution serves these queues at the same rate because the ferry visits a node only once in a cycle. Our solution, on the other hand, takes advantage of the non-uniformity in distances between node pairs and favors nodes

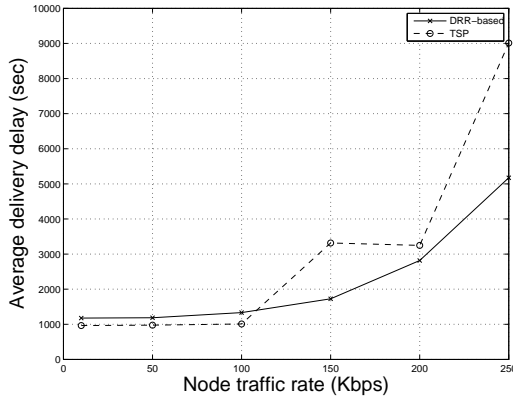


Fig. 4. Uniformly placed nodes, uniform Poisson traffic

that are relatively closer to other nodes.

At node traffic rate of 200 Kbps, the TSP-based solution performance improves significantly. In order to explain the reason for that, recall that the solution in [2] is a 2-phase solution. In the first one a TSP route is obtained, and in the second phase the periods of time the ferry spends at each node can get extended to meet certain bandwidth requirements. Before 200 Kbps, phase-2 of the TSP solution gives zero extended periods, however, at 200 Kbps or above it gives positive extended periods which improves the performance of the TSP solution. However, at 250 Kbps despite having positive extended periods, our solution performs about 80% better than the TSP solution.

Non-uniform traffic. Figure 5 shows the average delivery delay plotted against node traffic rate for this case. Note that the x-axis represents the traffic rate of a node after excluding the high traffic rate flows. We can see that our scheme always performs about 30% better than the TSP-based scheme. The only exception is at 100 Kbps where the two schemes perform very close to each other. The reason is that at 100 Kbps the TSP-based scheme starts getting positive extended periods of time for the ferry which improves the mean delay.

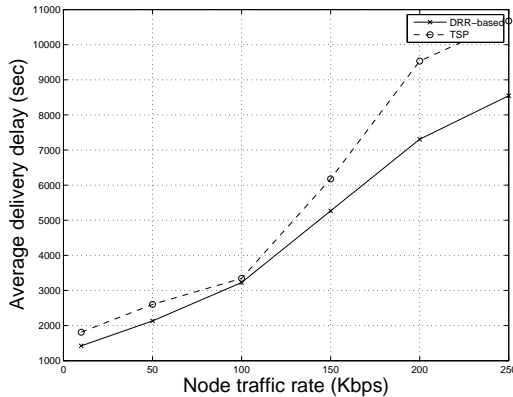


Fig. 5. Uniformly placed nodes, non-uniform Poisson traffic

C. Clustered node allocation

In this section we evaluate our algorithm with the nodes being clustered into 4 clusters. Clusters are randomly created as explained earlier in the section. Evaluation is performed for both uniform and non-uniform traffic.

Uniform traffic. Figure 6 shows the results for the 4-clusters case with uniform traffic. We can observe that at lower data rates (10 – 50 Kbps), similar to the uniform node allocation case, our algorithm and the TSP-based solution have similar performance. After that, our algorithm improves significantly over the TSP solution. Looking closer at the ferry routes with clustered nodes, we observe two things. First, the ferry goes more frequently to clusters with more nodes. In such clusters, the ferry will collect and deliver more data specially with uniform traffic. Second, the ferry goes less frequently to clusters that are relatively far from the other clusters.

Non-uniform traffic. In this case, we do not differentiate between intra-cluster and inter-cluster flows when choosing the high rate ones, instead they are selected randomly from all flows in the network.

Figure 7 shows the results for the 4-clusters case with non-uniform traffic. Our algorithm always outperforms the TSP-based solution. For low to moderate traffic rates (10 – 150Kbps), our algorithm achieves about 80% – 150% lower average delay than the TSP-based approach. At higher rates (200 – 250Kbps) the improvement in delay becomes lower.

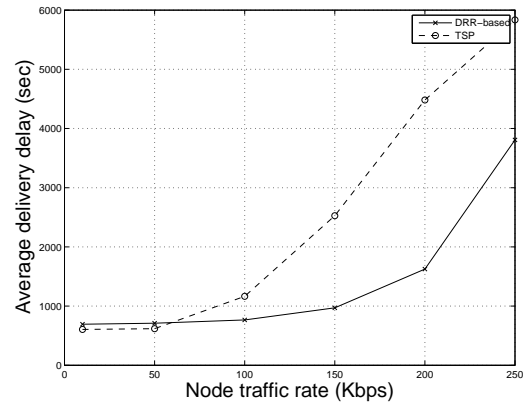


Fig. 6. Clustered nodes (4 clusters), uniform Poisson traffic

D. Weighted DRR-based ferry routing

In this section we explore the possibility of using our DRR-based routing algorithm to impose specific user preferences on certain flows in the network. Basically, we ask the following question: can we use our algorithm to give preference to reducing the mean delivery delay of certain flows in the network more than others? This can be very beneficial if certain flows contain important information that needs to be delivered faster than data coming from other flows of lower priorities.

In order to give preference to reducing the mean delay of any flow over others, the end nodes of that flow should be

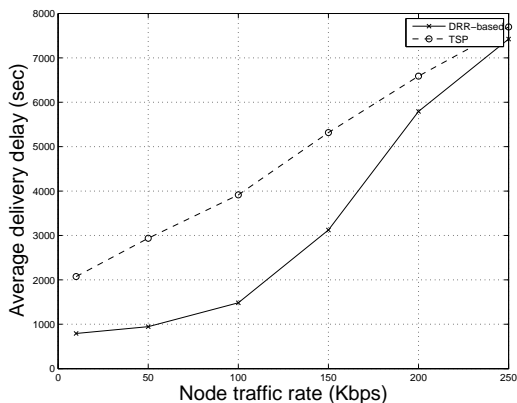


Fig. 7. Clustered nodes (4 clusters), non-uniform traffic

TABLE I
DRR VS WEIGHTED-DRR RESULTS

Scheme	Mean delay	mean {1 – 3} delay	1 : 2 : 3 : 4
DRR	556	556	1 : 1 : 1 : 1
Weighted-DRR	722	473	2 : 1 : 2 : 1

visited more often than with the normal algorithm application. As mentioned earlier in section III-B, the operation of our algorithm is governed by the functions `updateCredits` and `tieBreak`. In order to increase the probability of choosing either of the two ends of the flow of preference, we may need to modify both functions. Increasing the amount of service quantum granted to nodes that are members (senders or receivers) of preferred flows will increase the likelihood they will be selected as potential candidates when the ferry is deciding about its next destination. However, without changing the tie breaking logic, this is not enough to guarantee a higher chance of getting actually selected as the next destination. To achieve that, the `tieBreak` function should be more sensitive to these preferred nodes such that if the `candidate_list` contains any nodes that belong to preferred flows, the tie breaking logic should be biased towards selecting one of these nodes.

We give a simple example to illustrate the effect of using these changes on both the resulting ferry route and the mean delivery delay. Consider the node configuration in Figure 1(a), but assume all traffic flows are equal to 100 Kbps. We select the flow between nodes {1, 3} as the preferred one. In this example, we modified the functions `updateCredits` and `tieBreak` as follows. In `updateCredits`, at the beginning of a new round, we double the service quantum granted to nodes 1, 3. In `tieBreak`, after applying the first tie breaking filter (see section III-B), if we find either of nodes 1, 3 or both in the `candidate_list`, we randomly select one of them as the next destination. Table I summarizes the results with and without using the weighted DRR approach. Column 4 represents the relative frequencies of visiting nodes 1, 2, 3, 4 in the resulting ferry route. We observe that using the weighted DRR scheme improved the mean delivery delay for the traffic

between nodes 1, 3 by about 17.5%. However, in the same time, the mean delay time for other flows increased about 30%.

We believe that improving preferred flows mean delay can be done by the weighted DRR scheme. However, this problem needs a more careful and detailed study to develop more sophisticated methods that can be generally applied with more confidence in their results. We leave this as future work.

V. CONCLUDING REMARKS

In this paper we revisit the problem of message ferry route design in the case of stationary nodes and a single ferry. We start by formulating the problem as a Markovian Decision Problem (MDP) that minimizes the average data delivery delay. However, due to the exponential growth of the MDP state space, we were able to solve the MDP and get the optimal routes only for small problems.

Driven by both the insights we learned from optimal routes of small problems and the similarity between our problem and the link scheduling problem in traditional networks, we propose a new ferry route design algorithm that is inspired by the Deficit Round-Robin link scheduling algorithm. Simulations results show that our algorithm gives better mean data delivery delay than other existing approaches.

Future work includes developing the weighted version of our DRR-based ferry route design algorithm. We will also consider performing more simulations in different network conditions to improve the algorithm performance.

REFERENCES

- [1] W. Zhao, M. Ammar, and E. Zegura, "A message ferrying approach for data delivery in sparse mobile ad hoc networks," in *MobiHoc '04: Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing*, 2004, pp. 187–198.
- [2] —, "Controlling the mobility of multiple data transport ferries in a delay tolerant network," in *IEEE INFOCOM*, 2005, pp. 1407–1418.
- [3] M. M. mBin Tariq, M. Ammar, and E. Zegura, "Message ferry route design for sparse ad hoc networks with mobile nodes," in *MobiHoc '06: Proceedings of the 7th ACM international symposium on Mobile ad hoc networking and computing*, 2006, pp. 37–48.
- [4] W. Zhao and M. Ammar, "Message ferrying: Proactive routing in highly-partitioned wireless ad hoc networks," in *IEEE Workshop on Future Trends in Distributed Computing Systems*, 2003.
- [5] D. Jea, A. Somasundara, and M. Srivastava, "Multiple controlled mobile elements (data mules) for data collection in sensor networks," in *DCOSS*, 2005.
- [6] R. Sugihara and R. Gupta, "Improving the data delivery latency in sensor networks with controlled mobility," in *DCOSS*, 2008.
- [7] A. Mansy, M. Ammar, and E. Zegura, "Deficit round-robin based message ferry routing," Georgia Tech, Tech. Rep. GT-CS-11-03, 2011.
- [8] M. Shreedhar and G. Varghese, "Efficient fair queueing using deficit round-robin," in *SIGCOMM*, 1995, pp. 231–242.
- [9] S. K. A. Demers and S. Shenker, "Analysis and simulation of a fair queueing algorithm," in *Sigcomm*, 1989.
- [10] J. Nagle, "On packet switches with infinite storage," *IEEE Transactions on Communication*, vol. COM-35, 1987.
- [11] L. Zhang, "Virtual clock: A new traffic control algorithm for packet switched networks," *ACM Transactions on Computing*, vol. 9, 1991.
- [12] A. Keränen, J. Ott, and T. Kärkkäinen, "The one simulator for dtn protocol evaluation," in *Simutools '09: Proceedings of the 2nd International Conference on Simulation Tools and Techniques*, 2009, pp. 1–10.