

Predicting the Performance of Geographic Delay-Tolerant Routing

Erik Kuiper
Saab Aeronautics
Linköping, Sweden

Simin Nadjm-Tehrani
Dept. of Computer and Information Science
Linköping University
Linköping, Sweden

Abstract— Performing large scale physical experiments to evaluate the packet delivery ratio, latency, and resource utilization in mobile ad hoc networks is a costly and difficult exercise. A more cost-effective method of performing trade-off studies in connection with protocol evaluations is simulation. However, simulations are often time consuming and require a detailed description of a given scenario’s mobility and topology. A more general approach for scalability and performance studies is the use of abstract models of behavior in networks of mobile nodes. This paper illustrates the power of such a model by illustrating the prediction capability of the Forward-Wait framework with respect to the routing performance of geographic delay-tolerant routing. The Forward-Wait mathematical framework describes packet movements as a sequence of alternating forwarding and waiting phases. Key input parameters to the Forward-Wait framework are distributions describing the forwarding and waiting phases, and the source destination distance. The paper shows how these distributions can be derived and used to predict the routing performance in different scenarios and a variety of network sizes.

Keywords- *Geographic routing analysis, routing performance, delay-tolerant network, opportunistic routing*

I. INTRODUCTION

When planning an operation it is of vital importance to be able to predict the communication capabilities, and thus to be able to determine if additional communication assets need to be deployed. In this paper we illustrate the use of a mathematical framework, the Forward-Wait framework [10], to efficiently predict the packet delivery performance in an intermittently connected system where geographic delay-tolerant routing is used.

In the modern battlefield high power long range but low data rate radios can be supplemented by low power short range high data rate radio systems that can be carried by all actors in the battlefield, including individual soldiers. Since we cannot assume the availability of a high-speed backbone in the battlefield, these high data rate communication assets should dynamically cooperate in forwarding data by using ad hoc networking technologies. In this paper we will denote these radio equipped actors as nodes. Due to unpredictable movement and communication interference the normal network state is that full network connectivity is never available. A node can only expect to have low latency contact via peer forwarding to a subset of nodes. To communicate with nodes

outside a connected group a delay-tolerant store-carry-forward networking technology is needed, as practically shown within the DARPA DTN Program [16].

One example of a scenario where regular disconnections could occur is when a swarm of small and cheap unmanned aerial vehicles (UAVs) shall perform area reconnaissance. Due to the complexity of coordinating large numbers of vehicles, and due to mission requirements we cannot expect them to continuously form a fully connected network. We have previously shown that by using distributed pheromones a swarm of UAVs can efficiently perform area reconnaissance without requiring centralized control [10]. The lack of centralized control, combined with a low node density, indeed produces subsets of connected nodes.

The field of delay-tolerant networking (DTN) deals with systems where the time to transfer information from source to destination is long. In this paper we will use the Forward-Wait framework to predict the routing performance of geographic routing when the delays are caused by the absence of communication paths due to node distances being too large to permit communication. The availability of communication paths are however not static since node movement constantly repositions the nodes and provides new and changed communication paths. By temporarily storing a packet in a node a delay-tolerant routing protocol can use the movement of the nodes to establish contacts for forwarding packets in cases where contemporaneous paths do not exist. In a geographic routing protocol the knowledge of the location of nodes and the (expected) position of the destination is used to guide a packet towards the destination. For most geographic routing protocols this means handing over the packet to the closest node towards the destination that the custodian (the current node holding the packet) can find.

Predicting network performance in a DTN is typically done in simulation exercises since massive physical test bed exercises emulating a realistic scenario is a daunting and expensive proposition. When overall timing/performance questions are to be answered quickly, more cost-efficient methods are required. In this paper we illustrate the use of the Forward-Wait framework to estimate the packet delivery ratio. With known or predicted statistical properties regarding packet forwarding and knowledge about acceptable delivery latencies, the Forward-Wait framework can answer questions like “What % of transmitted packets in this network will reach the

destination with an average delay of X ?. As an example of a routing protocol that is suitable for this type of modeling, we use the Location-Aware Routing for Delay-tolerant networks (LAROD) protocol [11].

The Forward-Wait framework models the movement of a packet as a sequence of forwarding and waiting phases. A packet is forwarded towards the destination as much as the topology allows within a partition (connected group). When the packet reaches the edge of the partition, it waits until node movements have restructured the topology (i.e. changed the partitions) in such a way that forwarding becomes possible in a newly formed partition. In order to deal with a large multiplicity of network scenarios and mobility factors, we describe each phase using probability distributions. In the forwarding phase, the distributions describe the probability that a packet moves a certain distance by multiple hops, whereas in the waiting phase the distributions describe the probability of waiting for a certain time before the next forwarding opportunity arises. The characteristics of both the forwarding and waiting distributions depend on both node mobility, communication interference and the routing protocol. By adding a distribution describing the distance from the source to the destination for a given scenario and a maximum packet delivery time, the packet delivery ratio can be computed.

The paper is organized as follows. The next section presents DTN routing, routing modeling and the Forward-Wait framework. This is followed in Section III by a description of the scenarios used to illustrate the use of the Forward-Wait framework and forwarding and waiting distributions for these scenarios. The performance estimations for the scenarios are provided in Section IV and finally the paper ends with some concluding remarks in Section V.

II. BACKGROUND AND RELATED WORK

In this section we provide an overview of routing in DTNs and some different DTN routing modeling approaches including a description of the Forward-Wait framework.

A. Routing

Routing in disconnected networks is challenging due to node mobility and the lack of stability in the wireless links [6]. The constant movement of the nodes means that the routing protocol has to handle a constant changing, and often unknown, node topology. If the mobility is unknown or uncontrollable then the duration of a contact between two nodes becomes unpredictable which adds to the challenge.

To meet these challenges several routing protocols have been proposed and evaluated. The simplest routing protocols have no knowledge of the network structure and use probabilistic or flooding based techniques to distribute one or more copies of a packet. Examples of this type of routing protocols are Randomized Routing [14] and Spray and Wait [15]. By gathering information on earlier node contacts more informed decisions regarding packet forwarding can be made. Here we find routing protocols such as MEED [8] and MaxProp [3].

If the nodes are location-aware and the (approximate) position of the destination is known then this information can be used in the forwarding decisions. One such protocol is our earlier proposed LAROD-LoDiS [11]. In LAROD packets are forwarded towards the destination, and when this is not possible forwarding is paused until node movement allows the packet to be moved further towards the destination. To obtain the position of the destination LAROD uses the location service LoDiS. Another geographic delay-tolerant protocol is GeoDTN+Nav [5]. This protocol is designed for routing in a network of streets and it has three forwarding modes. When possible it takes greedy forwarding decisions at junctions and forwards a packet along the road between junctions. When greedy forwarding is no longer possible then it uses perimeter forwarding. In the perimeter mode a switch score is calculated, and if it is beyond a certain threshold the protocol switches over to DTN mode. In DTN mode GeoDTN+Nav uses the fact that most vehicles know where they are heading. Commuter busses have scheduled routes and taxis have known destinations where they will deliver their passengers. The paper does not address the problem of how knowledge about the location of the destination is distributed (i.e. how to implement a location service). Note that while networks using directional antennas on mobile nodes need to use geographic information to point the antennas, this does not mean that they use a geographic routing protocol to route the packets [12].

B. DTN Routing Modeling

Most models describing routing in DTNs are designed for very sparse networks where groups of connected nodes do not exist. The most common model used to estimate routing performance in sparse opportunistic DTNs is the inter contact time (ICT) model [7][13][14]. This model is appropriate for very sparse networks and unguided or random routing protocols since it is intrinsically limited to forwarding packets to encountered nodes only. The model describes the time between node encounters using various ICT metrics. The two most commonly used metrics are inter meeting time and next encounter time. The inter meeting time is the time between encounters of two specified nodes, and the next encounter time specifies when a node next encounters any other node. These encounters are often described using the exponential distribution [7][13][14]. Due to the analytically attractive properties of the exponential distribution the models can present closed form equations describing the delivery delay. The exponential distribution has been contended by Chantreau et al. [4] in a study of actual encounter data from humans carrying mobile devices. They have found that the encounter distributions exhibit a power law distribution with a coefficient less than one. The work has been continued by Karagiannis et al. [9] who have shown that the power law distribution is only valid up to a certain time after which the distribution decays exponentially. Another perspective on the ICT distributions comes from Zhang et al. [17] who have studied encounter properties in a network of scheduled buses. One important observation they made was that the delivery delays between bus pairs can differ quite significantly. When studying the waiting distributions in Section III we will see that the tails of the distributions are exponentially distributed, but not the head.

Due to the discrete nature of the ICT model, its application in large systems is computationally complex. To overcome this problem, Altman et al. [1] have proposed to describe packet spreading among nodes using a fluid model. The main limitation of their model is the assumption of monotone relaying schemes, that is, in the entire duration of the routing process being analyzed, the number of nodes that hold a copy of a message does not decrease. For a single-copy routing protocol, such as the ones we consider, the fluid model is not suitable since that model addresses multiple-copy routing schemes.

A major limitation of the previously described models is that they only consider very sparse scenarios where connected groups of nodes do not need to be taken into account. If the node movements and node density are such that nodes normally *do* have neighbors, then these models are not appropriate. One metric that accounts for connected partitions and non-uniform node mobility is the space-time graphs abstracted via delay expansion by Asplund [2]. Delay expansion enables us to determine bounds on worst-case latency for a wide class of broadcast protocols. The key idea is to describe the least number of uninformed nodes that will meet an informed node during the time period of interest. By using the number of informed nodes as a parameter in the function, non-uniform mobility can be handled. As the latency expressions in [2] are derived for broadcast, the results are not directly applicable to unicast routing.

C. The Forward-Wait Framework

To model geographic single copy routing in a DTN with connected groups of nodes we use the Forward-Wait framework [10]. The Forward-Wait framework describes how a packet is transported towards the destination by modeling two phases, forwarding and waiting. The forwarding phase is used to characterize how much closer a packet gets to the destination when it can be forwarded within a partition. The waiting phase accounts for how long time a packet has to wait after a forward (at the edge of a partition) until it can be forwarded again. While the model is assumed to be used to describe geographic routing, it is agnostic to the routing mechanism used to forward the packet.

Figure 1 provides an illustration of forwarding and waiting of two arbitrary packets from a source to a destination. In each forwarding phase, a packet travels some distance until it has to wait. While forwarding, the model ignores elapsed time which is seen in the figure by the horizontal lines. The waiting phases on the other hand ignore the node movement in space, as shown by the vertical lines. Our earlier evaluations have shown that these simplifications do not substantially decrease the usefulness of the framework. The two cases in the figure illustrate that, when a packet is generated by the source, it either has to wait if the source node is at the edge of a partition (upper track) or it can be forwarded within the current partition until it reaches the edge (lower track). Thus, a packet will travel in time and distance in a manner similar to one of the two tracks depicted in Figure 1 during its life time.

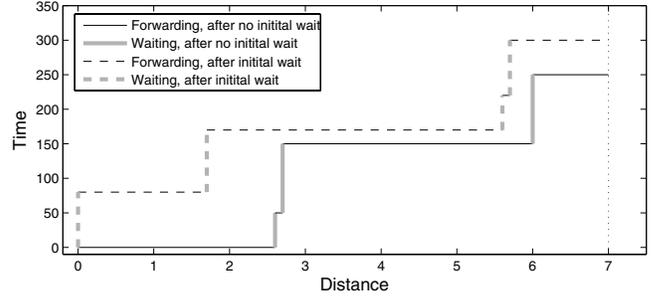


Figure 1. Time-distance illustration of packet routing.

The forwarding and waiting components of the framework are described using random variables. For forwarding, a random variable describes the distance traveled during this phase, and for waiting, a random variable describes the time spent in this phase. Examples of distributions describing the random variables for a specific scenario are provided in Section III. There, it will also be shown why it is important to separate the cases where a packet had to wait at the source or not before it could be forwarded.

From the random variables describing the individual forwardings, two different random variables are derived, D_n and D'_n . Both variables describe the distance traveled when a packet has gone through exactly n forwarding phases. Due to the differences in the initial forwarding distributions, i.e. whether the packet had to wait at the source or not, two separate random variables are required: One for a packet that could be immediately forwarded (D_n), and another if it was forwarded after a wait (D'_n). Analogously, a random variable that describes the time waited after n waits are derived, W_n . Together with the probability that a packet has to wait at the source before it can be forwarded, denoted by P_{wi} , the probability that a packet is delivered within time t at a distance of d is given by (1).

$$P_d(t, d) = P_{wi} \cdot \sum_{n=1}^{\infty} \left((P(D'_n \geq d) - P(D'_{n-1} \geq d)) \cdot P(W_n \leq t) \right) + (1 - P_{wi}) \cdot \sum_{n=1}^{\infty} \left((P(D_n \geq d) - P(D_{n-1} \geq d)) \cdot P_w(W_{n-1} \leq t) \right) \quad (1)$$

While $P_d(t, d)$ provides information on delivery probabilities it does not say anything about average delivery ratios. For that to be computed we need to know the distance distribution of the source-destination distances for packets to be transmitted. With a known distance distribution, described by the random variable Y , the delivery ratio can be computed according to (2). For the equation to be valid it is assumed that when a packet is at a distance of r (the nominal radio range) then delivery is guaranteed. If this is not the case then the equation has to be updated to take a probabilistic last hop distance into account.

$$P_{TTL}(t) = \int_0^r pdf(Y, x) dx + \int_r^{\infty} pdf(Y, x) \cdot P_d(t, x - r) dx \quad (2)$$

The framework has been validated in using data from ns-2 simulations [10]. The properties of the forward, wait and distance to destination random variables were extracted from the core packet forwarding data provided by the simulations. Using this data the delivery ratio computed by the framework was shown to accurately predict the delivery ratio provided by the simulations.

III. SCENARIOS AND SCENARIO DISTRIBUTIONS

A key factor determining the system wide average packet delivery ratio is node movement. The impact of the node movement is due to the fact that it to a large extent determines the temporally dependent partitions. Another key factor is the distance a packet has to travel to reach the destination. To illustrate the impact of the source-destination distance on the packet delivery ratio, we have chosen to use two different abstract scenarios to illustrate the application of the Forward-Wait framework. In both scenarios we consider independent and evenly distributed nodes in a square area with side length k . In the first scenario, called any-to-any, a node will communicate with equal probability with any other node. The distance between these two nodes are then distributed according to (3). In the second scenario, called any-to-C&C, all communication is between a mobile node and a command and control center (C&C) located at the center of one of the square edges. The distance between the mobile node and the C&C is then described by (4). The cdfs for the two distance distributions are illustrated in Figure 2. In the figure we clearly see that the communication distances are longer in the any-to-C&C scenario compared to the any-to-any scenario.

$$pdf(Y, d) = \begin{cases} 2d \cdot \left(\frac{\pi}{k^2} - \frac{4d}{k^3} + \frac{d^2}{k^4} \right) & 0 \leq d \leq k \\ 2d \cdot \left(\frac{2}{k^2} \tan^{-1} \left(\frac{2k^2 - d^2}{2k\sqrt{d^2 - k^2}} \right) - \frac{2}{k^2} + \frac{4\sqrt{d^2 - k^2}}{k^3} - \frac{d^2}{k^4} \right) & k < d < \sqrt{2}k \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

$$pdf(Y, d) = \begin{cases} \frac{2d\pi}{k^2} & 0 \leq d \leq k/2 \\ \frac{2d}{k^2} \tan^{-1} \left(\frac{k}{\sqrt{4d^2 - k^2}} \right) & k/2 < d < k \\ \frac{2d}{k^2} \left(\tan^{-1} \left(\frac{k}{\sqrt{4d^2 - k^2}} \right) - \tan^{-1} \left(\frac{\sqrt{d^2 - k^2}}{k} \right) \right) & k < d < \frac{\sqrt{5}}{2}k \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

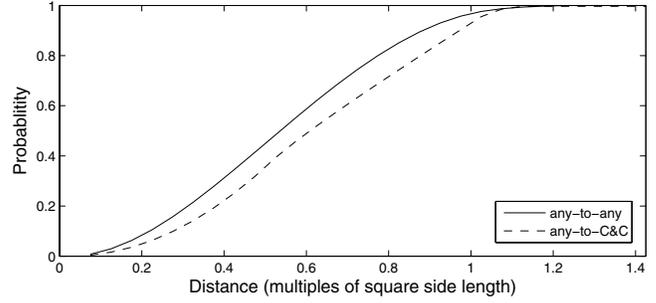


Figure 2. Cdf of distance to destination for the two scenarios.

The forwarding and waiting distributions used by the Forward-Wait framework capture how the node movement, properties of the radio communication, and the routing protocol affect forwarding. Due to the scenario-specific properties of these distributions they have to be uniquely determined for each scenario analyzed. However, the effort invested in creating these is much less that what is needed if full-blown simulations would be performed. For the illustration in this paper we have used an abstract description of the node movement and routing (see below). The reason for this choice is that abstract modeling enables us to obtain more data points, which can illustrate the power of the Forward-Wait framework more vividly. Both scenarios use the same movement and routing assumptions which means that the forwarding and waiting distributions will be the same for the two scenarios.

In the abstract models we have assumed that the nodes are Poisson distributed, move at a constant speed and that the routing is done by a geographic delay-tolerant routing protocol (e.g. LAROD). These abstractions then produce the distributions for different node densities as illustrated in Figures 3 to 5. The distributions have been generated by performing Monte Carlo simulations since the distributions are too difficult to derive analytically. For details see [10]. To make the forwarding distributions independent from radio technology we have chosen to present distance in the unit of the nominal radio range, and to treat node densities by the average node degree. The average node degree, c , is defined as $\rho\pi r^2$ where ρ is the node density and r is the nominal radio range. For Poisson distributed nodes in an infinite system, the average node degree equals the average expected number of neighbors of a node. The node densities used in this paper (node degree 2-5) represent the density range where nodes are grouped into partitions, but where the system is not yet very well connected.

Figure 3 shows the forwarding distribution from the source using the complementary cdf (ccdf) for four different node densities. The ccdfs also show the probability that a packet could not be forwarded from the source (P_{wi}) due to the fact that there were no appropriate nodes to forward the packet. As expected a higher node density means larger partitions and thus longer forwarding distances.

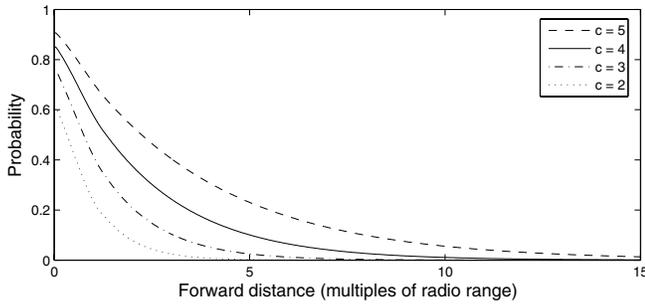


Figure 3. Probability of being forwarded at least the distance from the source with no initial wait (ccdf) for four different node densities.

In Figure 4 we see the forwarding distribution after a wait using the complementary cdf (ccdf) for four different node densities. From the figure we see that after a wait there is a significant probability that the packet is forwarded only a short distance. For details regarding the causes see [10].

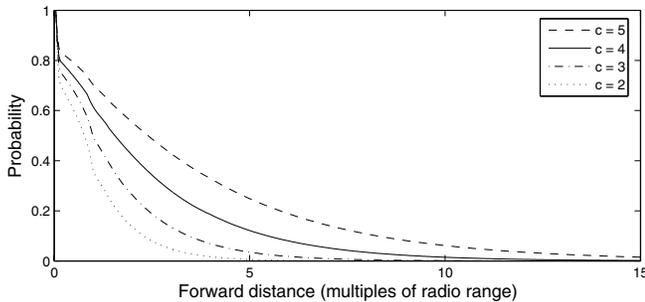


Figure 4. Probability of being forwarded at least the distance after a wait (ccdf) for four different node densities.

When a packet has reached the edge of a partition we have assumed that the routing protocol makes repeated randomized forwarding attempts. This means that the routing protocol waits for some random time before the next forward attempt is made. If the attempt is not successful the process is repeated. This results in the waiting cdf presented in Figure 5. The impact of the node density is clear, and higher node densities means shorter waits until the packet can be forwarded again.

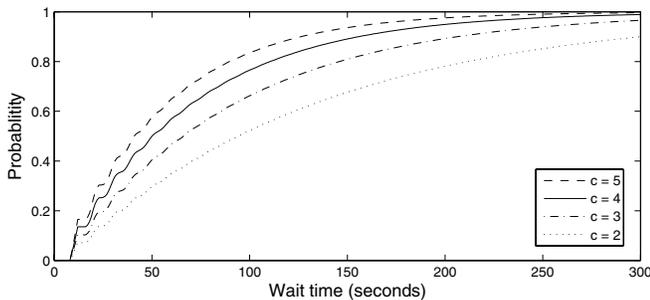


Figure 5. Probability of waiting at most a certain time for four different node densities.

IV. PERFORMANCE RESULTS

Based on the scenario-unique data presented in the previous section the Forward-Wait framework can predict the delivery ratio at a very low computational cost, even for large scenarios.

The capability to predict the delivery ratio for large scenarios is illustrated in Figure 6 and Figure 7. The two figures show the delivery ratio for the two scenarios at a fixed node density of 4 (average node degree) for different number of nodes (100, 1000 and 10000 respectively). The impact of the increased number of nodes at a constant node density is that transmission distances increase, which is seen in the figures by the need for a longer delivery time to achieve the same delivery ratio. Due to the slightly longer distance to destination the any-to-C&C scenario has a slightly lower delivery ratio compared to the any-to-any scenario at the same packet life time.

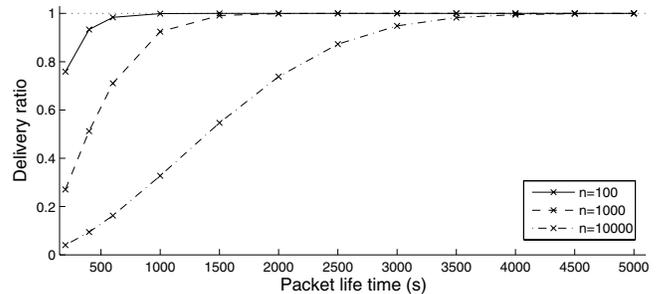


Figure 6. Delivery ratio for different sizes of the any-to-any scenario with a constant average node degree.

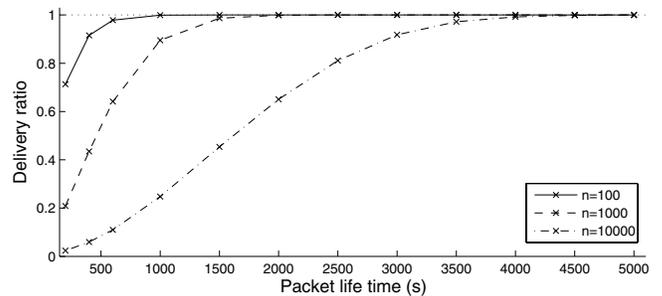


Figure 7. Delivery ratio for different sizes of the any-to-C&C scenario with a constant average node degree.

In a practical setting the scenario area is probably fixed and the question that needs to be answered is the packet TTL and/or the number of radio equipped units (nodes) required to provide an acceptable delivery ratio. In Figure 8 and Figure 9 we provide the delivery ratio provided by the Forward-Wait framework for the two scenarios for different node densities at a constant area size. The area was set to 20 by 20 radio radii which means that an average node degree of 2 corresponds to an average of 255 nodes, and a node degree of 5 leads to 637 nodes. As indicated by the distributions in Section III a higher density means longer forwarding distances and shorter delays which is here seen by a higher delivery ratio at a set packet time to live (TTL).

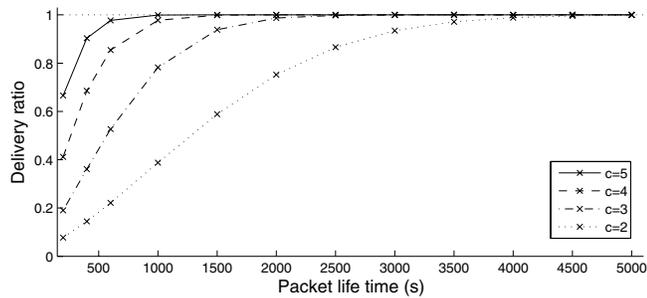


Figure 8. Delivery ratio for different node densities of the any-to-any scenario with a constant area size.

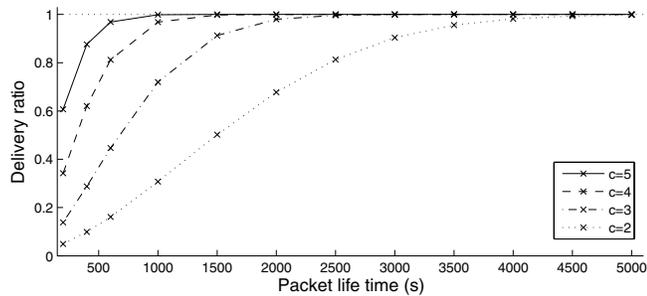


Figure 9. Delivery ratio for different node densities of the any-to-C&C scenario with a constant area size.

V. CONCLUSIONS

In this paper we have shown how the Forward-Wait framework can be used to quickly estimate the packet delivery ratio in different scenarios. The framework can be integrated into a radio network planning tool that determines a suitable TTL used in the configuration of the mobile devices for a given scenario. If an adequate delivery ratio cannot be achieved at the maximum acceptable TTL then additional communication assets might be required or different instructions regarding how the soldiers and vehicles may move can be given. If changes are made that affect the earlier assumptions, then the delivery ratio computations have to be redone and the results reevaluated. In the choice of the TTL a careful balance have to be struck between communication overhead and delivery ratio. A lower TTL reduces the communication overhead but also decreases the delivery ratio.

The key problem in the process is to determine accurate distributions for forwarding, waiting and distance to destination. Although we have derived them from mathematically grounded descriptions in this paper, we have in earlier work shown that the distributions can be derived from measurements from simulations and then re-used with a fast turn-around in the Forward-Wait framework predictions. In a practical environment the distributions should be based on earlier obtained empirical data that is modified based on known

differences between the current operation and the conditions under which the data was collected.

REFERENCES

- [1] E. Altman, T. Basar, F. De Pellegrini. Optimal monotone forwarding policies in delay tolerant mobile ad-hoc networks. *Performance Evaluation*. 2010, Volume 67, Issue 4. Elsevier
- [2] M. Asplund. *Disconnected Discoveries: Availability Studies in Partitioned Networks*. PhD Thesis. 2010. Linköping University Electronic Press. <http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-60553>
- [3] J. Burgess, B. Gallagher, D. Jensen, B. N. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. *Proceedings 25th IEEE International Conference on Computer Communications (INFOCOM)*. 2006. IEEE
- [4] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, J. Scott. Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms. *Proceedings 25th IEEE International Conference on Computer Communications (INFOCOM)*. 2006. IEEE
- [5] P.-C. Cheng, K. C. Lee, M. Gerla, J. Härrri. GeoDTN+Nav: Geographic DTN Routing with Navigator Prediction for Urban Vehicular Environments. *Mobile Networks and Applications*. Volume 15, Number 1, 2010. Springer
- [6] E. M. Daly, M. Haahr. The challenges of disconnected delay-tolerant MANETs. *Ad Hoc Networks*. Volume 8, Issue 2, 2010. Elsevier
- [7] R. Groenevelt, P. Nain, G. Koole. The message delay in mobile ad hoc networks. *Performance Evaluation*. 2005, volume 62, issues 1-4. Elsevier
- [8] E. P. C. Jones, L. Li, J. K. Schmidtke, P. A. S. Ward. Practical Routing in Delay-Tolerant Networks. *IEEE Transactions on Mobile Computing*. 2007, Volue 6, Number 8. IEEE
- [9] T. Karagiannis, J.-Y. Le Boudec, M. Vojnovic. Power Law and Exponential Decay of Inter Contact Times between Mobile Devices. *Proceedings of the 13th annual ACM international conference on Mobile computing and networking (MobiCom)*. 2007. ACM
- [10] E. Kuiper. *Geographic Routing in Intermittently Connected Mobile Ad Hoc Networks – Algorithms and Performance Models*. Forthcoming PhD Thesis, 2011
- [11] E. Kuiper, S. Nadjm-Tehrani. Geographical Routing with Location Service in Intermittently Connected MANETs. *IEEE Transactions on Vehicular Technology*, 2011, volume 60, number 2, IEEE
- [12] R.A. Nichols, A.R. Hammons. DTN-based free-space optical and directional RF networks. *Proceedings of 2008 Military Communications Conference (MILCOM)*. 2008. IEEE
- [13] T. Small, Z. J. Haas. The Shared Wireless Infostation Model - A New Ad Hoc Networking Paradigm (or Where there is a Whale, there is a Way). *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing (MobiHoc)*. 2003. ACM
- [14] T. Spyropoulos, K. Psounis, C. S. Raghavendra. Efficient Routing in Intermittently Connected Mobile Networks: The Single-Copy Case. *IEEE/ACM Transactions on Networking*, 2008, volume 16, number 1. IEEE/ACM
- [15] T. Spyropoulos, K. Psounis, C. S. Raghavendra. Efficient Routing in Intermittently Connected Mobile Networks: The Multiple-Copy Case. *IEEE/ACM Transactions on Networking*, 2008, volume 16, number 1. IEEE/ACM
- [16] J.L. Torgerson. DARPA DTN Phase 3 Core Engineering Support Final Report. JPL Publication D-65838
- [17] X. Zhang, J. Kurose, N. Levine, D. Towsley, H. Zhang. Study of a Bus-based Disruption-Tolerant Network: Mobility Modeling and Impact on Routing. *Proceedings of the 13th annual ACM international conference on Mobile computing and networking (MobiCom)*. 2007. ACM