Modeling the Age of Information in Emulated Ad Hoc Networks

Clement Kam*, Sastry Kompella*, Gam D. Nguyen*, Jeffrey E. Wieselthier† and Anthony Ephremides‡

*Information Technology Division, Naval Research Laboratory, Washington, DC
†Wieselthier Research, Silver Spring, MD
‡Electrical and Computer Engineering Department, University of Maryland, College Park, MD

***Abstract*—The age of information is a measure of the freshness of continually updated information, which has been studied theoretically for a variety of queueing models. We conduct an experimental evaluation of the age of information in various emulated network settings to see how well the queue models approximate a network for the purpose of age. We use CORE and EMANE to deploy a one-dimensional multi-hop network, and a two-dimensional grid multi-hop network, and we compare to the theoretical age under the D/M/1 and D/D/1 queue models. We observe that for these regular static networks, the D/D/1 models the average age performance better than the D/M/1 in lightly loaded networks with low loss. When losses increase, the age is greater than predicted by the D/D/1 model.**

I. Introduction

In many tactical networking applications, it is crucial to be able to report information updates across the network in a timely manner. However, there is a lack of fundamental understanding of how to conduct such reporting, in areas such as the frequency of update generation, queue management strategies, and feedback control algorithms. The age of information, first formally studied in [1], is a metric that is appropriate for evaluating these types of systems. Thus far, the work has focused on analyzing the age of information for simple queueing models, but there have been no studies into whether the theoretical insights for simple models can be applied to realistic wireless ad hoc networks.

The authors in [1] analyzed single server queues with a first come, first served queueing discipline. They considered memoryless and deterministic arrival and service processes, specifically M/M/1, M/D/1, and D/M/1 queues, with the D/M/1 yielding the lowest age. This was extended to multiserver queues (M/M/c) in [2]–[4], where the servers represented multiple paths in a dynamic network. It was shown that more servers yielded a shorter age, but this results in having packets arrive out of order at the monitor, such that a large percentage of traffic utilizing the system consists of obsolete packets. Other methods of controlling the age are considered in the literature, such as queueing discipline [5], [6], queue management [7]–[9], and update generation [10].

The age of information was studied in a realistic network setting in [11] using the Common Open Research Emulator (CORE) [12] and the Extendable Mobile Ad-hoc Network Emulator (EMANE) [13]. CORE and EMANE are tools developed in part by the Naval Research Laboratory (NRL) for building virtual networks, with CORE serving as the emulator for layer 3 and above of the protocol stack (network, transport, session, application), and EMANE provides a high fidelity layer 1 and 2 emulator (physical and data link). In [11], theoretical results for various system models (e.g., arrival/departure processes, multiple flows) were validated using a simple two node wireless link using the RF Pipe EMANE physical layer model and artificial path loss, delay, and jitter.

We extend this work on CORE/EMANE to an ad hoc wireless network, in which we more closely emulate a realistic network setting by using EMANE's 802.11b MAC with the OSPF-MDR ad hoc routing protocol. One- and two-dimensional multi-hop networks of varying sizes are considered and the age metric is computed for each scenario. Experiments are also conducted in the presence of varying levels of background data traffic for the two-dimensional network. We compare with theoretical results on single queue models and identify the utility and limitation of applying these models to real networks.

The paper is organized as follows. Section II provides a description of the setup of the emulated testbed. In Section III, we describe and derive the theoretical models we use for comparison. Section IV presents the results on one-dimensional networks, and section V presents the results on two-dimensional grid networks, and we summarize our findings in Section VI.

II. Testbed Setup

The system used for testing was a MacBook Pro notebook computer with a dual-core 3 GHz Intel Core i7 (Turbo Boost up to 3.5GHz) CPU with 256 KB L2 and 4MB L3 caches, and 16 GB of 1600 MHz DDR3 RAM and a 250 GB solid state disk. This notebook was running Mac OS X El Capitan (10.11). The CORE/EMANE environment was run on a virtual machine using VirtualBox 5.1.18 to run 64-bit Ubuntu 14.04.5 LTS on a 3.16.0-77-generic Linux kernel. The CORE 4.8 and EMANE 1.0.1 versions were installed. A fresh reboot of the virtual machine prior to running the emulation test provided more consistent results. In the future, we will utilize a more powerful, dedicated Ubuntu physical machine.

The MAC model in EMANE is the emulated 802.11b MAC (data rate of 11Mbps). The MAC and PHY settings are given in Tables I and II. For routing, we use the OSPF-MDR

TABLE I
IEEE 802.11abg EMANE Model Settings

Universal Phy Model Parameter	Value
Mode	0 (802.11b)
Promiscuous Mode	0
Unicast Rate	4 (11 Mbps)
RTS Threshold	0
Flow Control	Off
WiFi Multimedia	Off
Queue Size	0:255
Min. Contention Window	0:32
Max. Contention WIndow	0:1024
Arbitration Inter Frame Space	0:2
txop (Transmission Opportunity)	0:0
Retry Limit	3
Report Radio Metrics via R2RI	Off

TABLE II
Universal PHY EMANE Model Settings

Universal Phy Model Parameter	Value
RF Bandwidth	1.0M
Frequency	2.4G
Frequency of Interest	2.4G
subid	1
System Noise Figure (dB)	4
Transmit Power (dBm)	-36.04
Antenna Gain (dBi)	0.0
Enable Fixed Antenna Gain	On
Noise Mode	outofband
Noise Bin Size (μs)	20
Propagation Model	2ray

protocol [14], which is a MANET extension to OSPF. Traffic is generated using NRL's multi-generator (MGEN) network test tool, which tests IP network performance using TCP and/or UDP/IP traffic. We generate UDP packets with data payloads of size 128 bytes. For each emulation test, some settling time is allotted to bring up nodes and establish routes, and then data traffic (MGEN) tests are run for 10 seconds, and averaged over 10 separate runs.[1]

III. Results From Theory

To compare with results from theory, we consider the average age of two single-server queue models, a D/M/1 and D/D/1. Both models have a deterministic arrival process, but the D/M/1 has a memoryless service process (exponentially distributed service time) whereas the D/D/1 has a deterministic service process (constant service time). We use the average latency from the experimental results as the parameter to model the exponential and constant service times, from which we can compute the average age under each model.

[1]Each 10 second run gathers at least 128 samples, and over 8000 samples in the highest rate case.

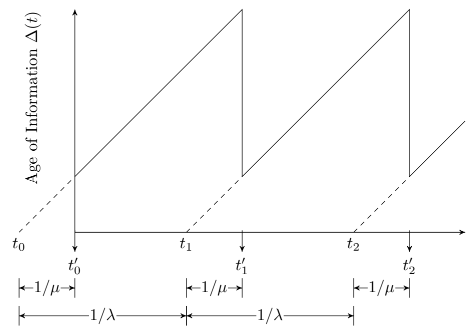

Fig. 1. Age function for D/D/1 queue.

The arrival process is deterministic with period $1/\lambda$. The average service time is given by $1/\mu$. For $\lambda < \mu$, the average age for the D/M/1 is given in [1] as

$$\Delta_{D/M/1} = \frac{1}{\mu}\left(\frac{1}{2\rho} + \frac{1}{1-\beta}\right) \quad (1)$$

where

$$\beta = -\rho\mathcal{W}(-\rho^{-1}e^{(-1/\rho)})$$

and $\mathcal{W}(\cdot)$ is the Lambert W function, and the utilization $\rho = 1/(\mu D)$.

The average age for the D/D/1 can be easily derived as follows. For $\lambda < \mu$, we have the age as a function of time in Figure 1, where t_n indicates the time of arrival and t'_n indicates the time of service for packet n. This is a periodic linear function, so it is clear that the average age is equal to the age at the midpoint of the line from t'_0 to t'_1, for example. Therefore, the age is equal to

$$\Delta_{D/D/1} = \frac{1}{\mu} + \frac{1}{2\lambda} \quad (2)$$

IV. One-Dimensional Networks

We first consider one-dimensional tandem networks, in which nodes are placed in a line, spaced 100 meters apart, 2 meters above ground level. The first node in the line transmits updates to the last node in the line. We choose the transmit power to be -36.02dB, which yields a packet reception of a 128 byte payload packet with about a 96% probability in the absence of interference when nodes are 100 meters apart.[2] There is a 0% success for nodes 200 meters apart, so there is no skipping over nodes in the route. We first run an experiment to determine the approximate per packet latency for each network size. We transmit data at a low rate of 5 packets per second for 10 seconds, and we average over 10 runs. The results are shown in Table III. The histogram is shown in Figure 2.

Using the given latency ($1/\mu$), we run tests for packet rates λ such that we have values of utilization $\rho = \lambda/\mu$ between 0 and 1. We plot the average age vs. ρ for line networks

[2]We use the default EMANE IEEE 802.11abg packet completion rate curves.

TABLE III
AVERAGE LATENCY VS NUMBER OF RELAYS

Number of Relays	Average Latency
0	0.001278
1	0.002411
2	0.003946
3	0.005580

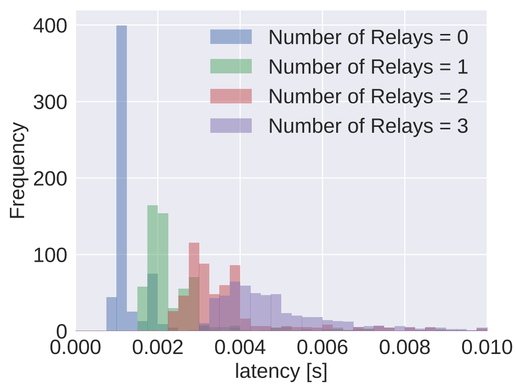

Fig. 2. Histogram of latencies for various numbers of relays.

with $0, \ldots, 3$ relays in Figures 3-6. When there is no relay (Figure 3), we note that for lower utilization ρ, both D/M/1 and D/D/1 track the average age closely. This is because the age is less sensitive to the service time when the packet rate is lower, since packets will need to wait less for other packets ahead of them. However, for higher ρ, the experimental age falls between the D/D/1 and the D/M/1 ages. This suggests that using the deterministic or memoryless service model alone is insufficient to predict the average age performance. However, for $\rho = 0.9$, there is a loss rate of 58%, so the rate is a bit higher than recommended operation. This is possibly due to some nonlinear processing or buffering delay in the system causing packet loss at high rates (704 packets/second), highlighting a shortcoming of the $\mu = 1/$latency model for the service rate.

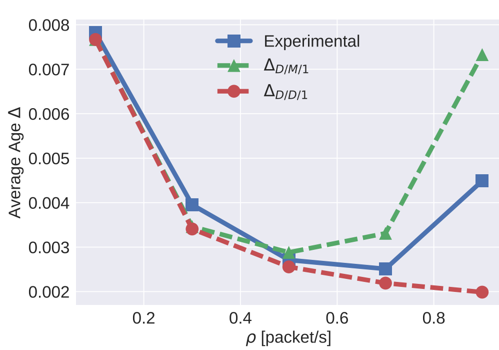

Fig. 3. Age vs. rate, no relays.

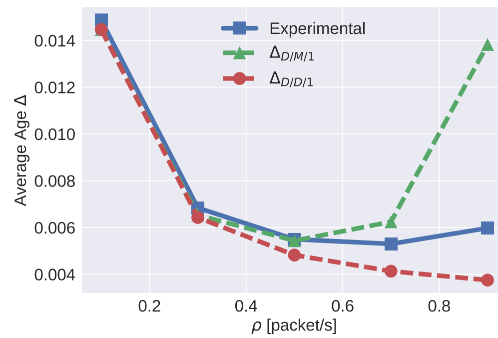

Fig. 4. Age vs. rate, 1 relay.

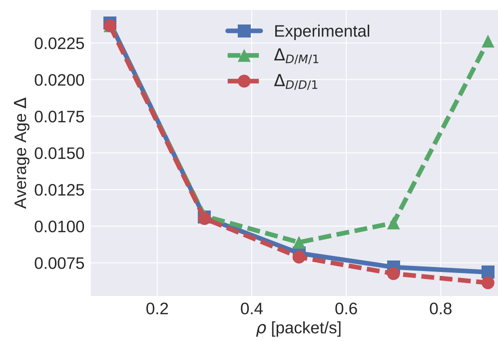

Fig. 5. Age vs. rate, 2 relays.

As the number of relays in the network increases, we observe that for smaller ρ D/M/1 and especially the D/D/1 track the experimental results closely, and for larger ρ, the D/D/1 models the average age well while the D/M/1 does not. This suggests that the deterministic service model can provide a good prediction of the average age performance for a larger number of relays in the network. We have plotted histograms of the latencies for the no relay case and 3 relay case in Figures 7 and 8, respectively. We see that the latency is higher in the 3 relay case than in the no relay case, and the distribution of latencies is relatively smooth in the 3 relay case, whereas there are multiple peaks in the no relay case. It seems that more relays leads to a smoothing of the distribution of latencies, and the similarity of the latencies makes the D/D/1 a better fit for predicting the average age. When the distribution of latencies is more irregular, a mixture of the D/D/1 and D/M/1 theoretical results may provide a reasonable prediction.

V. TWO-DIMENSIONAL GRID NETWORKS

We extend from a one-dimensional line network to a two-dimensional grid network, where the source node is located at the origin and the destination node is at $((N-1)100, (N-1)100)$ meters for an $N \times N$ grid of nodes spaced 100 meters apart. Choosing the transmit power to be -36.02dB allows adjacent horizontal and vertical nodes (100 meters apart) to receive with 96% success rate, whereas adjacent diagonal nodes (141.4 meters apart) have a 0% success rate. Again,

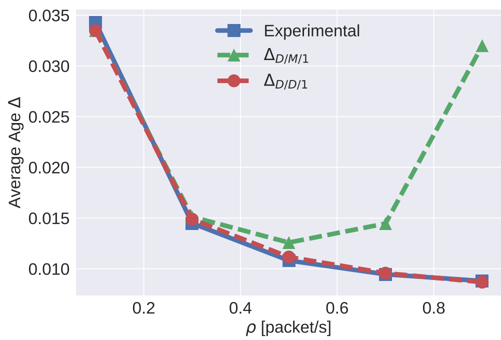

Fig. 6. Age vs. rate, 3 relays.

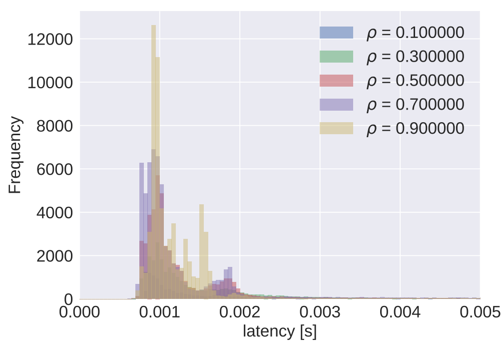

Fig. 7. Latency histogram, no relays.

we start with an experiment to determine the approximate per packet latency for each network size, transmitting data at a low rate of 5 packets per second for 10 seconds, averaged over 10 runs. The results are shown in Table IV, and the histogram is shown in Figure 9.

A. No Background Traffic

Using the given latency ($1/\mu$), we again run tests for utilization $\rho = \lambda/\mu$ between 0 and 1. We first consider the case where there is no competing data traffic in the background. We plot the average age vs. ρ for network sizes of $2\times 2, 3\times 3$, and 4×4 in Figures 10-12. For a 2×2 network (Figure 10), we again observe that both D/M/1 and D/D/1 track the average age closely for lower utilization ρ, but as utilization increases, the experimental average age is close to the D/D/1 prediction, but a little higher. There is one relay in the route for the 2×2 grid, and this is similar to the results seen in the line network for one relay (Figure 4). For the 3×3 grid, there are 3 relays to get to the destination, and we observe results similar to Figure 6. For the 4×4 grid, there are 5 relays, and we observe that the D/D/1 model gives a good prediction for all ρ. We plot the histogram of latencies in Figures 13 and 14 for the 2×2 and 4×4 cases, respectively. We see the distribution in the 2×2 case is relatively smooth and similar except for $\rho \geq 0.7$. In the 4×4 case, we see that the distributions are very smooth, which corresponds to a good match with the D/D/1 model.

TABLE IV
AVERAGE LATENCY VS GRID SIZE

Grid Size	Average Latency
2×2	0.002462
3×3	0.005452
4×4	0.007798

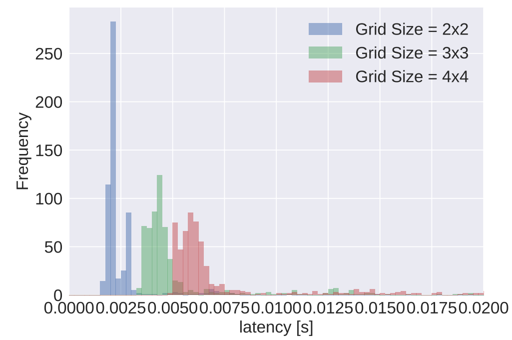

Fig. 9. Histogram of latencies for various grid sizes.

B. With Background Traffic

We now consider the case where there is some traffic in the background. Each node excluding the source and destination

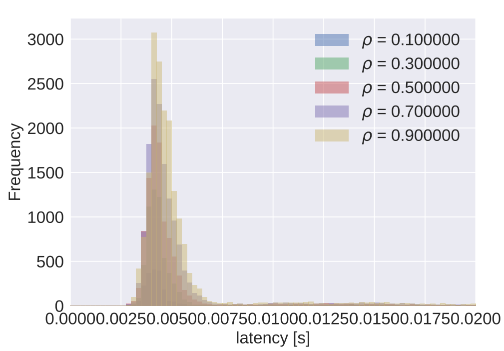

Fig. 8. Latency histogram, 3 relays.

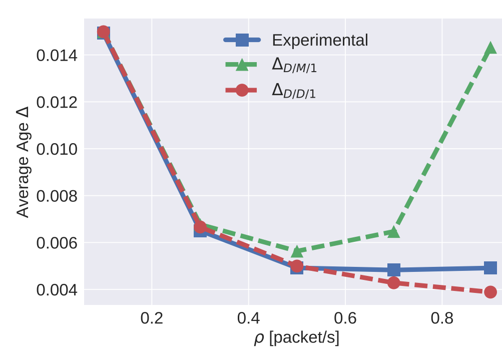

Fig. 10. Age vs. rate, 2×2.

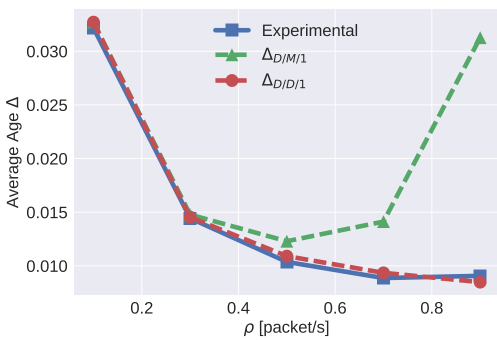

Fig. 11. Age vs. rate, 3×3.

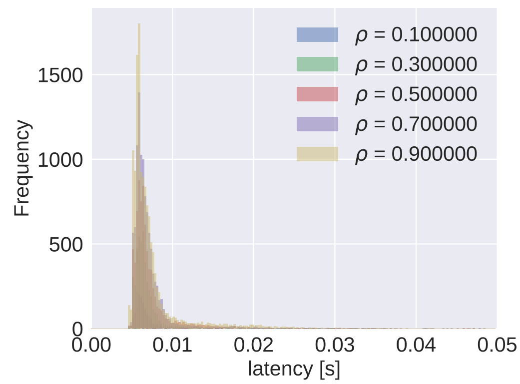

Fig. 14. Latency histogram, 4×4, no background traffic.

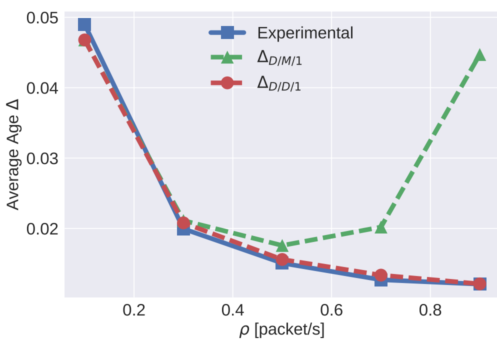

Fig. 12. Age vs. rate, 4×4.

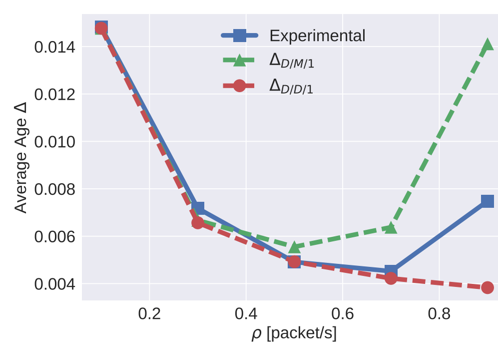

Fig. 15. Age vs. rate, 2×2, $\rho = 0.3$ per node of background traffic.

generates traffic at a rate of $\rho = 0.3$, and transmits to a randomly chosen non-source, non-destination node. We average over 10 runs, and plot the results in Figures 15-17. For the 2×2 case, the main difference is that the age is between the D/M/1 and D/D/1 models for $\rho = 0.9$. The effect of the background traffic is that the update packets lead to packet loss and delay in the case of high ρ, yielding a higher average age.

The average age for the 3×3 case is shown in Figure 16. In this case, there is packet loss even for lower ρ, due to the longer route with the same amount of competing traffic per node. This results in a slightly higher average age relative to that of the $D/D/1$ model. This effect is amplified in the 4×4 network, particularly for higher ρ, where the average age experimentally is drastically higher than that of the D/M/1 and D/D/1 models. We plot the loss for all 3 network sizes in Figure 18, and we see that the losses correspond to the increase in average age over the D/D/1. Even if we adjust our D/D/1 model to account for an effectively lower ρ due to packet loss, we still do not get a good approximation of the experimental age. The background traffic has a significant

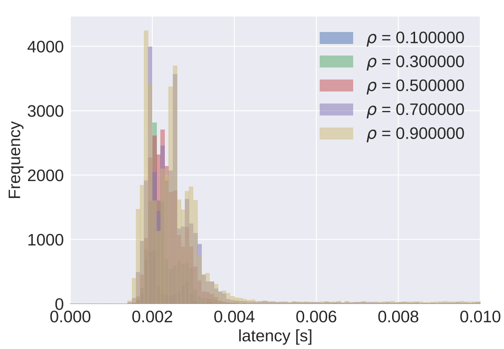

Fig. 13. Latency histogram, 2×2, no background traffic.

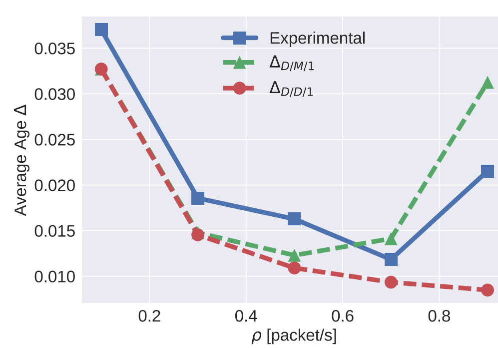

Fig. 16. Age vs. rate, 3×3, $\rho = 0.3$ per node of background traffic.

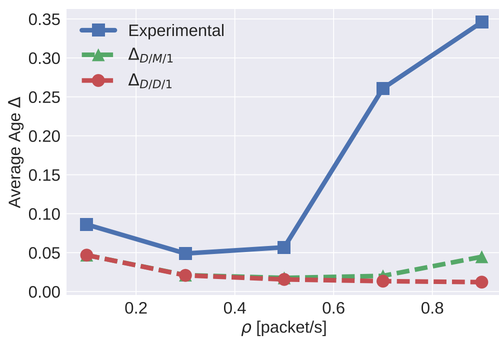

Fig. 17. Age vs. rate, 4×4, $\rho = 0.3$ per node of background traffic.

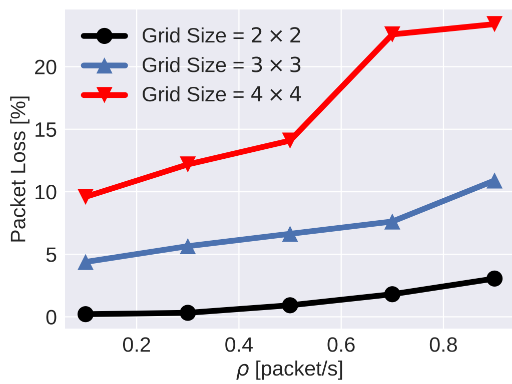

Fig. 18. Loss vs. ρ, $\rho = 0.3$ per node of background traffic.

effect on the latency, leading to a much higher age.

VI. Conclusion

In this work, we have experimentally evaluated the age of information in a wireless ad hoc network setting and compared it to results from theory, which focus on single queue models. We use a network emulation approach with CORE and EMANE to test an information updating system operating over static one- and two-dimensional multi-hop networks of varying sizes, with structured topologies. We show that the D/D/1 does a better job of modeling the average age than the D/M/1, specifically in low loss cases where the distribution of latencies for various ρ are smooth and similar. In some cases a mixture of the two models may provide a good prediction when the latencies are less smooth. Future work includes studying larger network sizes on more capable emulation hardware, random network deployment, and the impact of mobility on the average age.

Acknowledgment

This work was supported by the Office of Naval Research.

References

[1] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. IEEE INFOCOM*, Orlando, FL, Mar. 2012, pp. 2731–2735.

[2] C. Kam, S. Kompella, and A. Ephremides, "Age of information under random updates," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, Istanbul, Turkey, Jul. 2013, pp. 66–70.

[3] ——, "Effect of message transmission diversity on status age," in *Information Theory (ISIT), 2014 IEEE International Symposium on*, June 2014, pp. 2411–2415.

[4] C. Kam, S. Kompella, G. D. Nguyen, and A. Ephremides, "Effect of message transmission path diversity on status age," *IEEE Transactions on Information Theory*, vol. 62, no. 3, pp. 1360–1374, March 2016.

[5] S. Kaul, R. Yates, and M. Gruteser, "Status updates through queues," in *Conference on Information Sciences and Systems (CISS)*, Princeton, NJ, Mar. 2012, pp. 1–6.

[6] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Optimizing data freshness, throughput, and delay in multi-server information-update systems," in *2016 IEEE International Symposium on Information Theory (ISIT)*, July 2016, pp. 2569–2573.

[7] M. Costa, M. Codreanu, and A. Ephremides, "On the age of information in status update systems with packet management," *IEEE Transactions on Information Theory*, vol. 62, no. 4, pp. 1897–1910, April 2016.

[8] C. Kam, S. Kompella, G. D. Nguyen, J. E. Wieselthier, and A. Ephremides, "Age of information with a packet deadline," in *Information Theory (ISIT), 2016 IEEE International Symposium on*, July 2016.

[9] ——, "Controlling the age of information: Buffer size, deadline, and packet replacement," in *2016 Military Communications Conference (MILCOM 2016)*, Nov. 2016.

[10] Y. Sun, E. Uysal-Biyikoglu, R. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," in *IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications*, April 2016, pp. 1–9.

[11] C. Kam, S. Kompella, and A. Ephremides, "Experimental evaluation of the age of information via emulation," in *MILCOM 2015 - 2015 IEEE Military Communications Conference*, Oct 2015, pp. 1070–1075.

[12] (2015) Common Open Research Emulator (CORE) website. [Online]. Available: http://www.nrl.navy.mil/itd/ncs/products/core

[13] (2015) Extendable Mobile Ad-hoc Network Emulator (EMANE) website. [Online]. Available: http://www.nrl.navy.mil/itd/ncs/products/emane,https://github.com/adjacentlink/emane

[14] R. Ogier and P. Spagnolo, "Mobile Ad Hoc Network (MANET) Extension of OSPF Using Connected Dominating Set (CDS) Flooding," RFC 5614, Aug. 2009. [Online]. Available: https://rfc-editor.org/rfc/rfc5614.txt