

Poster Abstract: Throughput-Delay Tradeoff in Small and Sparse Mobile Ad hoc Networks

Swapnil Bhatia, Radim Bartoš, and Kevin Ma

Department of Computer Science

University of New Hampshire

Email: {sbhatia, rbartos, kjma}@cs.unh.edu

Abstract—The asymptotic throughput-delay tradeoff has been extensively studied for dense wireless mobile ad hoc networks (MANETs) as a function of increasing density. However, many practical applications exist for *sparse* MANETs where mobile node density goes to zero and the number of nodes is small. Asymptotic throughput-delay laws discovered for dense networks need not identically hold in this sparse regime. This paper takes the first step in characterizing the throughput-delay tradeoff for such sparse MANETs. We find that as the MANET becomes sparser, throughput decreases and delay increases, as expected. If relaying is disabled then the throughput and delay depend only on the size of the area of operation. While relaying does increase throughput, the single packet relaying strategy worsens the delay for small MANETs. Greedy relaying overcomes this worsening without trading throughput, but only for rapidly mixing mobility. Unlike in dense networks, local broadcasting does not provide any significant benefit. Packet repetition does decrease delay, but only at the expense of reduced throughput, for small MANETs. For slowly mixing mobility, relaying worsens delay for small sparse MANETs and none of the above techniques help. Our results are useful in practical underwater MANETs where n is typically small and the MANET is sparse.

I. MOTIVATION

The characterization of the throughput-delay tradeoff in wireless ad hoc networks has been the subject of study in a number of papers in recent years [1]–[7]. Most of the previous work in this area, except that of Spyropoulos et al. [2], has focused on dense wireless networks with the tradeoff being studied as the number of network nodes n goes to infinity. A fundamental assumption in such work is that the wireless network under study is sufficiently dense with the performance under increasing density in fact being the subject of study. A motivating example justifying this assumption is the ad hoc sensor network where a dense deployment of sensor nodes is desirable. Such a characterization of the throughput-delay tradeoff as a function of the number of sensors provides valuable insight into the scaling behavior of sensor networks that are large and dense.

In contrast, the practical deployment scenario for many wireless ad hoc networks, particularly those involving mobile nodes, is such that while a dense deployment is desirable, it is rarely feasible. Consider a MANET of autonomous underwater vehicles (AUVs) deployed for bathymetry or underwater surveillance. Even for such basic underwater missions, the oceanic region involved is far too vast to be amenable to sensing and measurement by a dense MANET. As a result,

practical AUV MANETs tend to be small and sparse for which extant capacity results studying scaling behavior as a function of increasing density provide little insight into the tradeoffs involved in such sparse networks. A fundamental differentiating characteristic of a sparse MANET is the high probability with which a mobile node may be outside the transmission range of any other node. Whereas the interference among concurrent transmissions plays a deciding role in the throughput-delay tradeoff in dense networks, it is clear that such interference is extremely rare in sparse networks. For example, in an underwater environment, acoustic transmissions are quickly attenuated and the distances involved are relatively large. Thus, an increased density, for small values of density, can actually improve performance whereas such an increase is detrimental in high density MANETs. In what other respects might sparse MANETs be different from dense ones? This is the motivating question of our work and this short paper takes the first step towards answering it.

II. MODEL

We model the spatial region in which the mobile nodes of a sparse MANET move as a discrete undirected graph with loops. We experiment with two graphs: the complete graph on m^2 vertices and the $m \times m$ two-dimensional torus. The mobile nodes of our MANET move in a random manner by performing a random walk on the underlying graph at each discrete time step. That is, at any time step, a mobile node situated at a vertex v either remains at v or moves to a neighboring vertex, each with probability d^{-1} , where d is the regular degree of any vertex (including the loop) in the underlying graph. We consider $n \geq 2$ mobile nodes walking randomly on the underlying graph. Each mobile node i produces data packets destined for exactly one other node denoted $\text{dest}(i)$. This source-destination mapping is fixed and is chosen by selecting a random derangement of $\{1, \dots, n\}$.

If a set M_v of (more than one) mobile nodes meet at any vertex v , then data transmission occurs according to the following rules:

- 1) Every mobile node $i \in M_v$ such that $\text{dest}(i) \in M_v$, transmits a single packet to $\text{dest}(i)$. We call this a *direct delivery*.
- 2) Every node $i \in M_v$ that could not perform a direct delivery chooses at random a $j \in M_v$ for which it carries

one or more packets delegated to it by j 's source. It then delivers at most p such packets to j , where the particular packets transmitted, if more than p are available, are also chosen at random. We call this a *relayed delivery*.

- 3) Every node $i \in M_v$ that could not perform a direct or relayed delivery transmits exactly one packet to another randomly chosen node $j \in M_v$ requesting j to deliver the packet to $\text{dest}(i)$. We refer to this as *packet delegation*.

If $|M_v| = 1$, no transmissions occur at v . When $|M_v| > 1$, each node either makes a direct or relayed delivery or delegates a packet. Each node possesses infinite space for storing delegated packets that it has accepted. Transmissions occur in a round-robin manner and are coordinated through some TDMA scheme at each vertex v . Transmission of a single packet takes a constant amount of time and the total time spent in communication at each vertex is negligible in comparison to the inter-vertex travel time. Since the network is sparse, transmissions occur concurrently at all vertices v without interfering with each other.

The above rules are similar to those used in Grossglauser et al. [7]. In addition, we also study the following variants. When *delegation* is disabled, a node is only capable of direct delivery via method (1) above, i.e., method (3) is not available. When delegation is enabled and $p = \infty$ in method (2), we call this *greedy relaying*. If *local broadcasting* is enabled, then a packet transmitted by a node i via method (3) is broadcasted to all nodes $j \in M_v$, i.e., i delegates the packet to all other nodes present at v through a single transmission [4]. If *packet repetition* is enabled with parameter r , then every packet produced by a node i is delegated r times by i or until it is delivered directly, whichever occurs earlier.

III. SIMULATION RESULTS

We simulated our model of a sparse MANET on a complete graph and a torus and measured the average throughput and delay. We adopt a natural definition for sparsity; it is the difference in the orders of magnitude of the number of mobile nodes (n) and the number of vertices in the underlying graph (m^2). The scaling behavior with a fixed $n = 10$ and increasing m is shown in Fig. 1. Delay increases and throughput decreases with increasing sparsity because mobile nodes must walk a larger average number of random steps for a meeting.

Fig. 2 shows the throughput-delay tradeoff on a complete graph for a fixed $m = 100$ with n varied between 2 and 500. For any node i , the inter-meeting time with $\text{dest}(i)$ and with any other node j is key in determining the delay and throughput in our model. When relaying is disabled in a sparse MANET, direct delivery is the only method of communication. The measured delay without delegation shown in Fig. 2 thus corresponds to the inter-meeting time of the complete graph, which is m^2 for a complete graph with loops on m^2 vertices. Thus, the delay is independent of the number of nodes n . Of course, this is not the case in dense MANETs where n decides the interference [7]. Thus, for sparse MANETs, throughput and

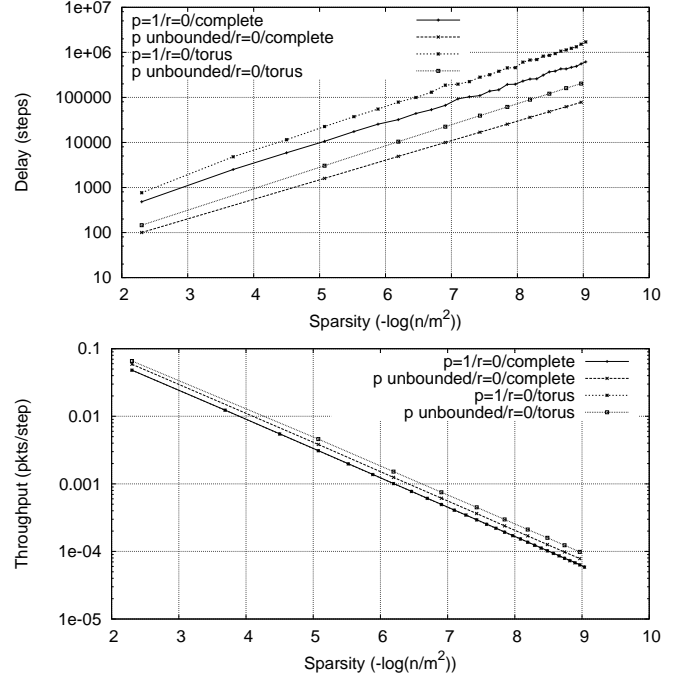


Fig. 1. Delay (top) and throughput when m is varied with $n = 10$ constant.

delay remain constant as a function of sparsity when packet delegation is disallowed.

When relaying is enabled, throughput improves as expected [3], [7]. This trend of improvement in throughput continues with decreasing sparsity as long as the MANET remains sparse enough to resist inter-transmission interference. On the other hand, relaying increases the average delay by two orders of magnitude over the baseline delay when packet delegation is disallowed. For $n > 2$ but smaller than a critical value (about a 150 in this case), relaying significantly worsens delay, while for larger n , it improves delay. Thus, for small sparse MANETs, plain relaying ($p = 1, r = 0$) is not the best strategy. However, we find that a greedy relaying strategy is indeed effective in lowering delay in small MANETs. Thus, the increased delay in small sparse MANETs can be attributed to the restriction on the number of packets p that any node i can deliver in a single meeting with any other node j via method (2). If p is finite, then i may have to meet j multiple times before it can deliver all the packets that it has accumulated from j 's source, leading to a significantly higher delay. A large value of p on the other hand requires that a longer time be spent at each vertex—this tradeoff depends upon the transmission bandwidth and the speed of movement. Delay can also be decreased through packet repetition ($r > 0$), but only at the cost of substantially decreased throughput. Local broadcast has no discernible effect on delay or throughput because $|M_v| > 2$ necessary for local broadcast to occur is a rare event in sparse MANETs.

The delay is significantly higher for a MANET on the torus (Fig. 3) because a walk on a torus mixes slower than that on a complete graph (a mixing time of $\Theta(1)$ versus $\Theta(m^2)$

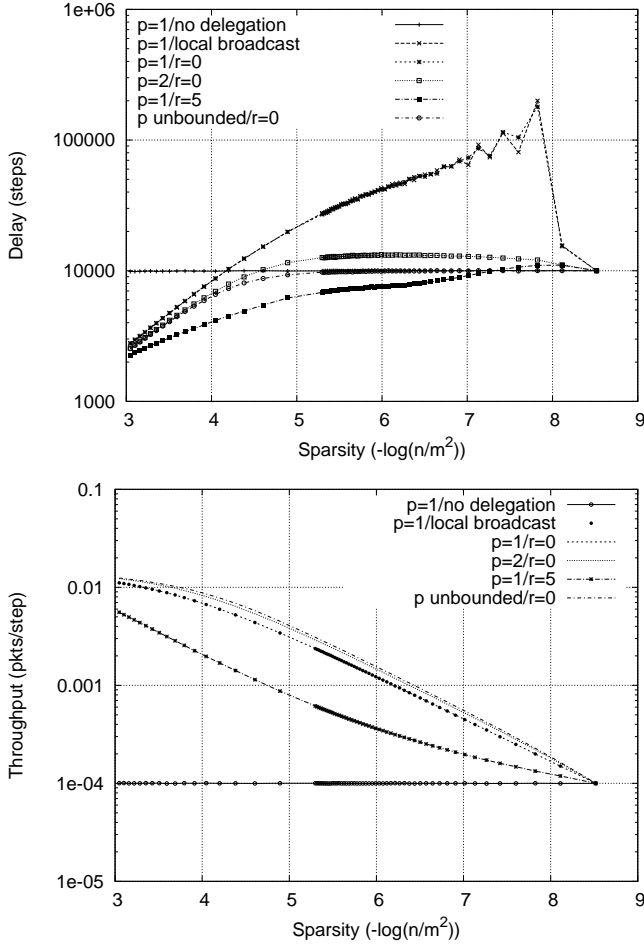


Fig. 2. MANET on a complete graph when n is varied with $m = 100$.

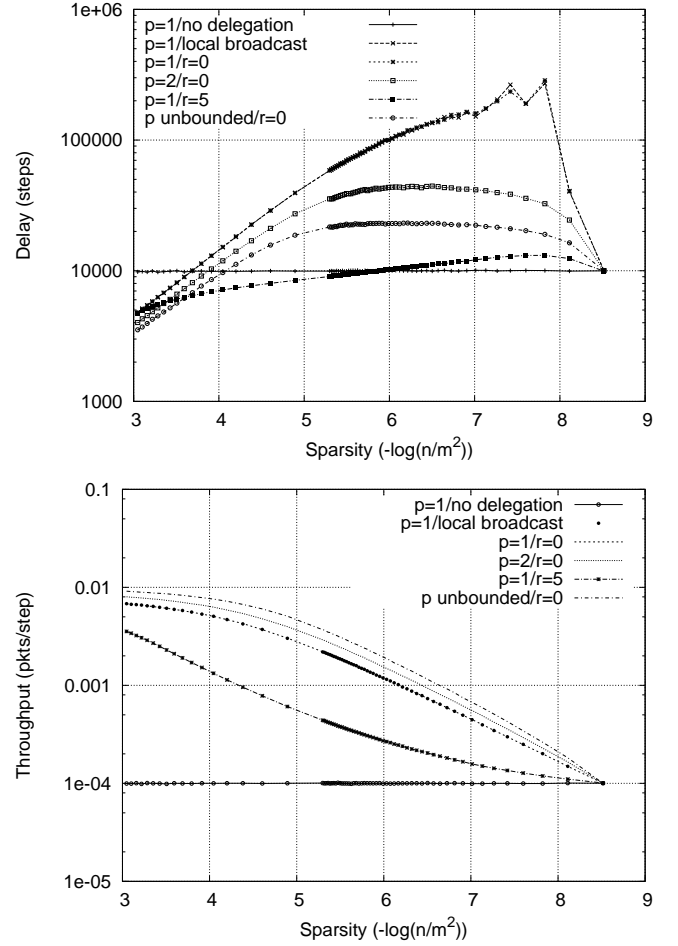


Fig. 3. MANET on a two-dimensional torus when n is varied with $m = 100$.

[8]). In the case of the torus, the increase in delay for small $2 < n \leq 150$ has a component in addition to the one for the complete graph. The greedy relaying strategy is not as effective on the torus as it is for the complete graph. This is essentially because of the slower mixing of a walk on a torus. A mobile node on a torus has fewer movement choices at a vertex than on a complete graph, which results in a more “correlated” movement of meeting nodes. Because of this increased correlation, delegated packets are not spread out uniformly across all nodes. Hence, the fate of the delay of many packets may end up relying on the meeting time of a few relay nodes. Packet repetition mitigates this artifact and reduces delay because packets now have a higher chance of being spread uniformly across different relay nodes.

IV. CONCLUSIONS AND FUTURE WORK

In this short paper, we reported our results on the throughput-delay tradeoff for sparse MANETs which show that the scaling behavior of small and sparse networks is different from dense networks studied in the past. What kind of motion and traffic model and on what graph captures practical AUV missions? How does the increased delay region scale with m ? What is the analytical explanation for the observed throughput and delay? We intend to explore such questions

about sparse networks in our future work. This research was supported in part by grant N00014-05-1-0666 from the U.S. Office of Naval Research.

REFERENCES

- [1] G. Sharma, R. Mazumdar, and N. Shroff, “Delay and Capacity Trade-offs in Mobile Ad Hoc Networks: A Global Perspective,” in *Proc. of IEEE INFOCOM*, Barcelona, Spain, Apr. 2006.
- [2] T. Spyropoulos, K. Psounis, and C. Raghavendra, “Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks,” in *Proc. of ACM SIGCOMM WDTN-05*, Philadelphia, PA, Aug. 2005.
- [3] A. E. Gamal, J. Mammen, B. Prabhakar, and D. Shah, “Throughput-Delay Trade-off in Wireless Networks,” in *Proc. of IEEE INFOCOM*, Hong Kong, Mar. 2004.
- [4] R. M. de Moraes, H. R. Sadjadpour, and J. Garcia-Luna-Aceves, “On Mobility-Capacity-Delay Trade-off in Wireless Ad Hoc Networks,” in *Proc. of IEEE/ACM MASCOTS*, Volendam, The Netherlands, Oct. 2004.
- [5] R. C. Shah, S. Roy, S. Jain, and W. Brunette, “Data MULEs: Modeling a Three-Tier Architecture for Sparse Sensor Networks,” Intel Research, Seattle, WA, Tech. Rep., Jan. 2003.
- [6] N. Bansal and Z. Liu, “Capacity, Delay and Mobility in Wireless Ad Hoc Networks,” in *Proc. of IEEE INFOCOM*, San Francisco, CA, Mar. 2003.
- [7] M. Grossglauser and D. Tse, “Mobility Increases the Capacity of Ad Hoc Wireless Networks,” *IEEE/ACM Trans. on Networking*, vol. 10, no. 4, Aug. 2002.
- [8] D. Aldous and J. A. Fill, Reversible Markov Chains and Random Walks on Graphs. [Accessed 15-June-2007]. [Online]. Available: <http://www.stat.berkeley.edu/~aldous/RWG/book.html>