

The Capacity of Multi-Hop Wireless Networks with TCP Regulated Traffic

Sorav Bansal, Rajeev Shorey
Shobhit Chugh, Anurag Goel, Kapil Kumar^f
IBM India Research Laboratory,
Block I, Indian Institute of Technology,
Hauz Khas, New Delhi 110016, India
Email: {soravban, srjeev}@in.ibm.com

Archan Misra
IBM T J Watson Research Center,
19 Skyline Drive,
Hawthorne, NY 10532, USA
Email: archan@us.ibm.com

Abstract—We study the capacity of multi-hop wireless networks with TCP regulated traffic. We study the dependence of the capacity on the transmission range of nodes in the network. Specifically, we examine the sensitivity of the capacity to the speed of the nodes and the number of TCP connections in an ad hoc network. By incorporating the notion of a minimal acceptable QoS metric (loss) for an individual session, we argue that the *QoS-aware capacity* is a more accurate model of the TCP-centric capacity of an ad-hoc network. We study the dependence of capacity on the source application (Telnet or FTP) and on the choice of the ad-hoc routing protocol (AODV, DSR or DSDV). We conclude that persistent and non-persistent traffic behave quite differently in an ad-hoc network.

1. INTRODUCTION

It is widely recognized that network capacity is a major constraint in the effective deployment of multi-hop wireless networks. In networks where nodes use the same physical channel, the transmission range of individual nodes is a key determinant of capacity, since it effectively determines the extent of spatial reuse possible. Capacity studies of multi-hop, ad-hoc wireless networks typically concentrate on the MAC layer, and investigate the effect of various parameters, such as the radio transmission range, the node density or the average distance between session endpoints, on the maximal achievable throughput. For example, [1] shows how an increase in n , the number of nodes, causes the average throughput of an individual node to degrade as $O\left(\frac{1}{\sqrt{n \log n}}\right)$ when the nodes are randomly distributed. Similarly, [2] studies the behavior of the IEEE 802.11 MAC layer [3] and shows how the end-to-end throughput available to each node degrades as $O(3)$ for random traffic patterns, and remains constant if the sessions exhibit appropriate *localization* properties.

We have recently [4] studied the capacity problem from the standpoint of the transport layer and showed how the capacity metrics at the TCP layer (*assuming persistent TCP sources*) behave quite differently from the corresponding idealized link-

layer metrics. All these studies have, however, focused on the transport-layer throughput in *static*, multi-hop networks and do not consider the impact of mobility on the overall network capacity. In this paper, we extend the earlier studies on TCP-centric capacity in ad-hoc networks in the following ways:

(i) We examine how variations in the mobility rates impact the goodput achieved at the TCP layer.

(ii) We study how the capacity is affected by the traffic load (number of TCP connections) in the network, taking care to ensure that the offered loads are feasible in the sense that they do not cause violation of the associated Quality of Service (QoS) constraints.

(iii) We study the sensitivity of capacity to two different traffic sources, representing two extremes of TCP-based applications. We consider both persistent or greedy (e.g., FIT) traffic, as well as non-persistent or intermittent (e.g., Telnet, HTTP) traffic.

(iv) We study how our choice of the ad-hoc routing protocol (AODV, DSR or DSDV) affects the capacity achieved by TCP sources, and how the optimal value of the transmission range R varies with the choice of the routing protocol.

Our studies assume that all nodes are identical in the sense that they all use the same transmission range R , the same buffer lengths B and follow the same mobility pattern; we study the properties of TCP traffic as these configurable parameters are varied. Our focus is thus on evaluating the right choice of R under different operating conditions. The common physical channel is assumed to have a bandwidth C , for our studies with IEEE 802.11 LANs, we have used $C = 2$ Mbps. Since the capacity definition for TCP traffic is not immediately apparent, we define the network's TCP-centric capacity as the total (cumulative) goodput achieved by all TCP sessions.

In [4], we have studied the achievable throughput in a static ad-hoc network, when the individual flows perform TCP-based congestion control. We specifically explore the interaction of TCP's flow control with the scheduling operations at the MAC layer. The throughput of a single TCP flow is a function of both its end-to-end loss probability p and the round-trip delay RTT ; for moderately low values of p , the throughput of a persistent TCP flow varies [6] as $p = \frac{1}{I} \frac{C}{RTT}$ where I is a proportionality constant. Since changes to the transmission radius R affect both p and RTT , the throughput of an individual

^f Shobhit Chugh, Anurag Goel and Kapil Kumar completed the B. Tech degree from the Department of Computer Science and Engineering, Indian Institute of Technology, New Delhi in May 2002. Shobhit Chugh and Anurag Goel are now with Logistics.com at Burlington, MA, USA. Kapil Kumar is with STMicroelectronics, India.

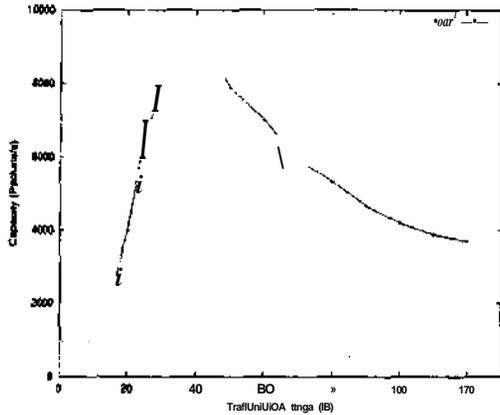


Fig. 1. Capacity versus Transmission Radius for a static network

TCP session turns out to be heavily dependent on R .

If R is very small, the average degree of connectivity of the graph is low, leading to a corresponding increase in the number of hops (N) in an individual data path. As an increase in N increases both the RTT and p , the individual session throughput p decreases with decreasing R , leading to an overall drop in the system throughput. On the other hand, if R is larger than a certain value, then the resultant MAC-layer channel interference limits the number of concurrent transmissions and effectively throttles the capacity of individual hops. In such a situation, the TCP sessions are prevented from better exploiting the network by the larger delays caused due to collisions and backoffs at the MAC layer. Accordingly (as shown in Figure 1), when the area of the network and the number of active sessions is fixed, the capacity is a bell-shaped function of the R . For values of R larger than the optimal value R' , the network is *MAC-layer constrained*, with the channel interference dominating the throughput. For values of R smaller than R' , the network is *TCP-layer constrained*, with the TCP sessions being unable to pump enough packets into the network.

In this paper, we extend our analysis to study both the effect of introducing mobility in such ad-hoc networks and of varying the number of TCP flows. In addition, we shall also investigate the sensitivity of our results to the choice of a specific ad-hoc routing protocol. The performance of different ad-hoc routing protocols has been extensively studied in literature (e.g., see [7], [8], [9]); however, these studies have not explicitly considered the differential impact of varying R .

The rest of the paper is organized as follows. We enumerate the various simulation parameters that we have used in our studies in Section II. In Section III, we vary the node mobility in the network and study how it affects the capacity, for both persistent and intermittent TCP sources. In Section IV, we first study the effect of varying the number of sources on the system capacity, and subsequently (Section V) introduce the notion of minimal QoS constraints to define a more accurate measure of

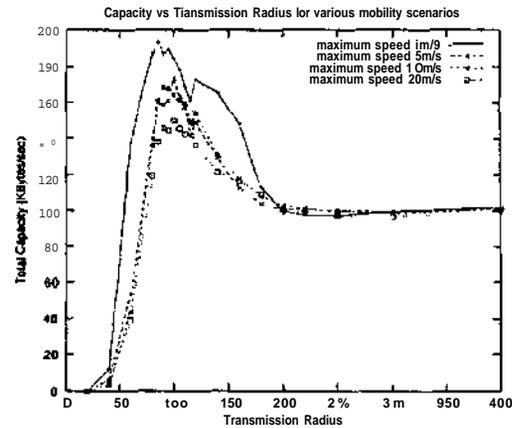


Fig. 2. Capacity versus Transmission Radius with Varying Speed (RP Traffic)

the TCP-centric capacity. In Section VI we study three ad hoc routing protocols (AODV, DSR and DSDV) from the point of view of network capacity. Finally, Section VII has the conclusion.

II. SIMULATION PARAMETERS

Our performance studies are carried out using simulations performed on the *ns-2* simulator [10]. While we have experimented with a variety of node densities and layouts, we report all results using a representative 50 node ad hoc network. The nodes are distributed randomly and move about in an area of 500 m x 500 m. For our studies, we set the interference range to be twice the transmission range. A fixed number of TCP connections are run for a duration of 500 seconds and the capacity is calculated by summing the TCP goodputs over all the connections. Results are averaged over a minimum of 10 separate runs. While TCP Reno is used as the transport layer, the data sources ("the application") are chosen to be either persistent (FTP) or intermittent (Telnet). Unless otherwise specified, results are reported using DSR as the ad-hoc routing protocol. Node mobility is modeled using the Random Waypoint model [7], with the *pause time* of all nodes set to 0 in all simulations.

III. CAPACITY WITH VARYING NODE MOBILITY

We first study the effects of mobility on two different classes of application - persistent and non-persistent. While the persistent traffic (FTP) is greedy and attempts to inject packets whenever permitted by TCP's congestion window, the non-persistent traffic (Telnet) produces only sporadic bursts of packets. Hence, as will be seen later, while the effects of interference are clearly visible in the case of FTP, the MAC-layer interference is not so critical in applications such as Telnet.

The capacity of a network with 40 FTP connections with different mobilities has been plotted in Figure 2. In Figure 3, we

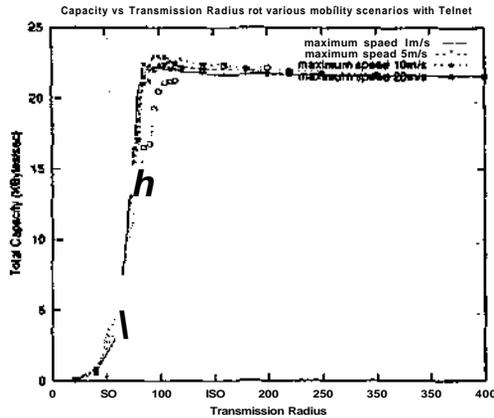


Fig. 3. Capacity versus Transmission Radius with Varying Speed (Telnet)

plot capacity versus transmission range with varying mobility for Telnet traffic. The speed of a node is uniformly distributed between 0 4s and a maximum value (shown in the figures).

In Figure 2, we see that the capacity of the network decreases with increasing node speed. Clearly, the overhead of route re-establishment, and the fraction of packets dropped due to routing failures, increases with increasing node mobility. Furthermore, the optimal transmission radius R' (corresponding to maximum capacity) shifts to the right (i.e., R' is higher) with an increase in the node speed. In other words, we need a higher transmission range to counteract the high mobility in the network. Note that, as in the case of a static topology used in Figure 1, the shape of the capacity versus transmission range plot is bell-shaped for mobile networks as well. Moreover, we observe that the increase in R' with larger mobility rates is not very dramatic; accordingly, it appears that a single well-chosen value of R' will ensure reasonably good (although not necessarily optimal) network performance, even if the node speeds cannot be predicted precisely.

For Telnet traffic (Figure 3) and fixed node speeds, the capacity increases with increasing transmission radius till a value of R' , after which the capacity does not change appreciably with R . Since telnet traffic is sporadic in nature, we do not observe the interference effect visible with FTP. In other words, due to the lower average packet arrival rate, the network is never MAC-layer constrained even at large values of R , there are very few requests for concurrent access to the 802.11 channel. While the number of non-interfering concurrent transmissions possible in the network does dip with an increase in R , the telnet goodput remains unaltered. This is also the reason that the network capacity with Telnet application is significantly lower than that with the FTP (persistent) source. Increasing R beyond R' does not result in any further increase in the throughput; the number of packets transferred in a single burst is usually too small to allow TCP to take advantage of

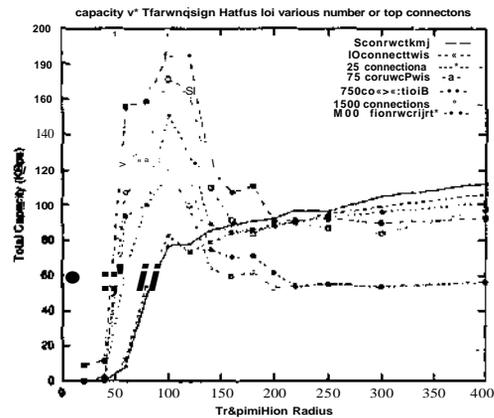


Fig. 4. Capacity with Varying FTP Connections

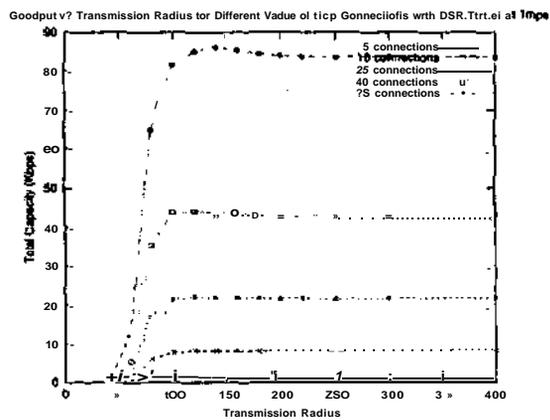


Fig. 5. Capacity with Varying Telnet Connections

the smaller loss probability and round-trip delays. Hence for light non-persistent traffic, the TCP goodput depends solely on the connectivity of the network.

IV. VARYING NUMBER OF TCP CONNECTIONS

We now study how changes to N , the number of active TCP flows, affects the overall system throughput. In general, we can clearly expect the TCP goodput for an individual session to degrade with an increase in the offered load. In essence, an increase in N leads to a potential increase in both p and RTT , since the larger load leads to more frequent buffer overflow and larger buffering delays. Accordingly, Equation 1 implies a drop in the TCP throughput. However, the effect on the overall system capacity is unclear, since this reduction may or may not be offset by an increase in the number of distinct flows.

This phenomenon is studied in Figure 4 where the number of FTP connections is varied from 5 to 200 in an ad hoc network

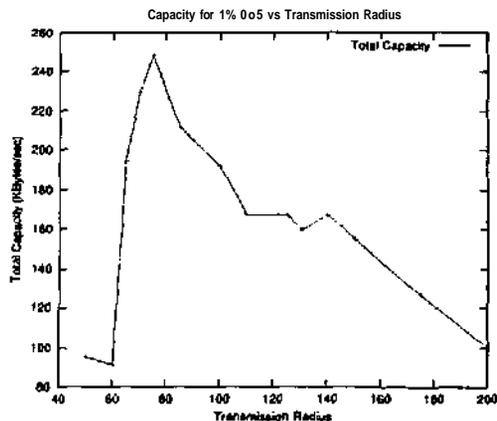


Fig. 6. Total Network Capacity versus Transmission Radius

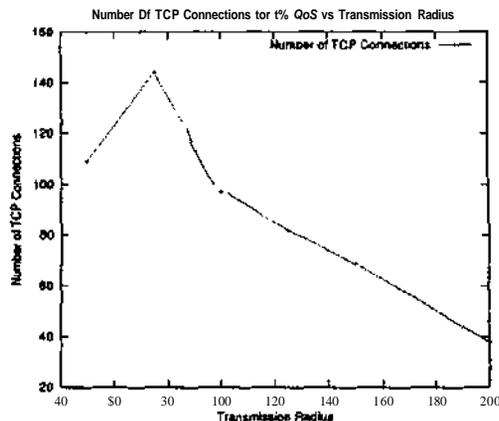


Fig. 7- Number of TCP Connections versus Transmission Radius

with 50 nodes. The distribution of node speed is uniformly distributed between 0 m/s and 1 d s . As the number of TCP connections increase, the network capacity increases initially. However, the capacity begins to degrade beyond 750 TCP connections in the ad hoc network. It is worth noting that the drop in throughput for values of R larger than R' is more acute for larger values of N . When the network becomes MAC-constrained and nodes must perform exponential backoff more frequently to access the channel, the individual nodes are unable to clear their packet buffers at a sufficiently high rate. Accordingly, the buffering losses and delays are higher for higher values of N , leading to a sharp drop in system throughput-in essence, the system is now nearer to *congestion collapse*. Figure 4 thus illustrates an important point: *if the number of persistent TCP flows cannot be accurately estimated in advance, it is better to adopt a conservative approach and set R to a smiler value. If the chosen value of R is larger than R' , the network suffers a much stiffer penalty.*

Figure 5 plots the results obtained by varying the number of Telnet sessions. Due to the rather sporadic injection of packets, the overall traffic load is always rather low for Telnet sources. Accordingly, the network is always *source-constrained*, even at large values of R . Accordingly, the capacity of the network is seen to linearly increase with an increase in the number of TCP sessions. As seen earlier in Figure 3, the system capacity saturates at a certain value R' of the transmission range.

V. QOS-AWARE CAPACITY

We have seen that an unbounded increase in the number of persistent (FTP) connections can ultimately lead to a drop in the system capacity. Figure 4, however, does not consider the associated issue of QoS; in particular, it does not incorporate the fact that an increase in the number of sessions, typically leads to a decrease in the performance metrics of an individual session. For a more accurate description of capacity, we

have to additionally determine the maximum number of active sessions that can be simultaneously present in the network without causing an unacceptable degradation in the quality of an individual session.

We now attempt to answer the question - what is the maximum number of TCP connections that can be accommodated in an ad hoc network with K nodes, without causing a violation in the QoS metrics of each individual connection? We point out the inadequacy of the capacity graph in Figure 4 and argue that Quality of Service (QoS) constraints should be taken into account in order to answer the above question.

We now assume that the maximum packet loss rate is fixed at 1% for minimum acceptable quality of service. In Figure 6, we plot the maximal total capacity (subject to the constraint of an upper bound of 1% on the packet loss rate) versus the transmission radius R for our 50 node network. We see that capacity is maximum when R is approximately 75 meters. Figure 7 plots the maximum permissible number of TCP connections (subject to the 1% loss constraint) versus R . We see that the maximum acceptable number of TCP connections shows a peak as well, and begins to drop fairly sharply as R is increased beyond an optimal value. Further, we see that a transmission radius of - 75 meters corresponds to the maximum number of TCP connections, i.e., ~ 140 . Note that Figure 7 conveys more information than Figure 6 - it enables us to obtain an upper bound on the number of simultaneous TCP connections (for 1% QoS) that can be permitted in a K node ad hoc network. These graphs also convey the appropriate value of the transmission radius R' that yields maximum system capacity and maximum number of TCP connections. It is therefore clear that, even with QoS constraints imposed in the network, the total capacity versus R behavior exhibits the bell-shaped behavior seen earlier in Figure 1 and Figure 4. However, the optimal transmission radius R' under QoS constraints can be appreciably different from the optimal range R' obtained sim-

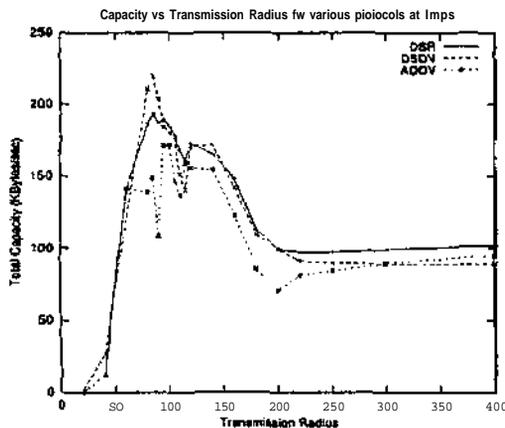


Fig. 8. Capacity with Different Ad Hoc Routing Protocols (Speed: 1 m/s)

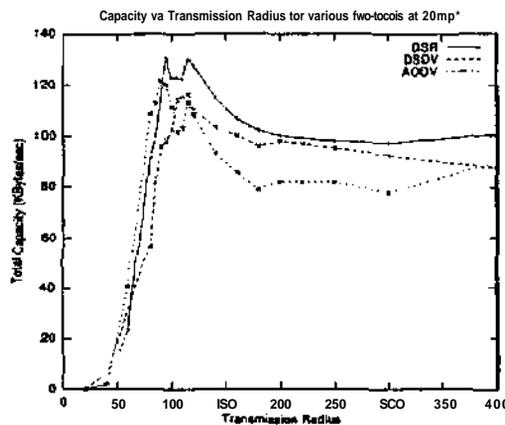


Fig. 9. Capacity with Different Ad Hoc Routing Protocols (Speed: 20 m/s)

ply from an aggregate capacity standpoint. Accordingly, future studies of ad hoc network capacity should take into account the appropriate QoS constraints (packet loss rate, latency, etc) as well.

VI. AD HOC ROUTING PROTOCOLS

In this section, we study the effect of a choice of different routing protocols on the TCP-centric capacity. While earlier studies have indeed performed a comparative study of these routing protocols, they do not consider the impact of varying the transmission range R . We performed our studies with three of the more widely discussed ad-hoc routing protocols, namely AODV, DSR and DSDV [7]. We compare DSDV, DSR and AODV in a SO-node network with (i) low node mobility (Figure 8) and (ii) high node mobility (Figure 9). Both these graphs correspond to FTP traffic.

It can be seen that for almost all values of the transmission

radii, AODV protocol yields the lowest capacity as compared to the DSDV and the DSR protocols. The difference in capacity of DSR, AODV and DSDV protocols is more pronounced at higher mobility and at higher transmission ranges. However, for our purposes, it is more important to note that *the value of the optimum transmission range R' is fairly similar for all three ad-hoc routing protocols*. In other words, the results of our capacity analysis seem to hold (at least qualitatively) are fairly independent of the precise choice of the ad-hoc routing protocol. Since routing protocols will continue to evolve with time, verifying this protocol-independence is essential to make our results and observations meaningful for future ad-hoc networks.

VII. CONCLUSION

In this paper, we focus on the TCP capacity over a multi-hop, wireless network where all links share the same physical channel. We study the sensitivity of TCP capacity to different network parameters (node mobility, number of TCP connections), different ad hoc routing protocols and different applications (Telnet, ITP).

Several avenues remain for future work. Our studies need to be extended to cover UDP traffic and UDP-based applications, which often have stringent constraints on additional QoS metrics such as delay. The results in this paper assume the Random Waypoint model. In a recent paper [51], the authors explore the tradeoff between the number of hops from source to destination and the overall bandwidth available to individual nodes as the transmission power is varied. The results assume UDP traffic and a *modified random direction* mobility model [5]. It will be interesting to study the sensitivity of TCP traffic to different mobility models

REFERENCES

- [1] P. Gupta and P. R. Kumar. "The Capacity of Wireless Networks". *IEEE Transactions on Information Theory*, March 2000.
- [2] J. Li, et al. "Capacity of Ad Hoc Wireless Networks", *Proceedings of ACM MOBICOM '01*, July 2001.
- [3] IEEE Computer Society LAN MAN Standards Committee. "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications". *IEEE Std. 802.11-1997*.
- [4] S. Bansal, A. Misra, R. Shorey, et al. "Energy Efficiency and TCP Throughput of Multi-Hop Wireless Networks". in the *Proceedings of IEEE INFOCOM'02*, June 2002, New York, USA.
- [5] E. M. Rayer, P. M. Melliar-Smith and L. E. Moser. "An Analysis of the Optimum Node Density For Ad hoc Mobile Networks". in the *Proceedings of IEEE ICC01*, June 2001, Helsinki, Finland.
- [6] M. Muthuswamy, I. Sasmke, I. Mahdavi and T. Ott. "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm", *Computer Communication Review*, July 1997.
- [7] J. B. M. D. Malin, D. Johnson, Y. Hu and I. Jelchsva. "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols", *Proceedings of ACM MOBICom'01*, July 2001.
- [8] S. Das, C. Perkins and E. Royer. "Performance Comparison of Two On-Demand Routing Protocols for Ad-Hoc Networks", *Proceedings of IEEE INFOCOM 2000*, March 2000.
- [9] P. Johansson, T. Larsson, N. Hedman and B. Mielcencck. "Routine Performance Analysis for Mobile Ad-Hoc Networks: A Comparative Performance Analysis". *Proceedings of ACM MOBICom99*, August 1999.
- [10] "The ns-2 network simulator", <http://www.mash.cba.hawaii.edu/~ns2/>.