Scalability of Mobile Ad Hoc Networks: Theory vs Practice

R. Ramanathan, R. Allan, P. Basu, J. Feinberg, G. Jakllari,
V. Kawadia, S. Loos, J. Redi, C. Santivanez and J. Freebersyser

Raytheon BBN Technologies
Cambridge, MA
{ramanath}@bbn.com

***Abstract*—Over the past decade, the theoretical or *asymptotic* scalability of Mobile Ad Hoc Networks (MANETs) has been extensively studied. However, the implication of these asymptotic results on finite, brigade-sized networks with real-life assumptions is not well-understood.**

We present a two-pronged study on the scalability of military networks with assumptions and goals pertinent to such networks: 1) we investigate the traffic distribution characteristics in a typical military network and show that it follows a *power law* which exhibits very good scaling properties; 2) we introduce the notion of "in practice" scalability and derive an expression for the in-practice scalability of a simple example network. Our study indicates that MANETs may well be adequately scalable in practice even if they are asymptotically unscalable, and that *military* MANETs may also even be asymptotically scalable by virtue of their traffic characteristics.

I. Introduction

Mobile Ad Hoc Networks (MANETs) have been of significant interest to the military for several decades [1], [6], [9]. More recently, significant investment has been made in several waveforms for the Joint Tactical Radio System (JTRS), most of which are based on the MANET architecture. Further, with decreasing hardware cost and increasing communication capacity, the military envisions large MANETs consisting of several thousands or more nodes (an example is the Defense Advanced Research Projects Agency (DARPA) Wireless Network after Next (WNAN) [8] system). The scientific and engineering community around these emerging MANETs clearly believes in their viability.

At the same time, information-theoretic results within the last decade appear to show that MANETs are inherently unscalable. In particular, a seminal paper by Gupta and Kumar [4] – hereinafter referred to as the *GK result/paper* – shows that the per-node transport capacity of an arbitrary MANET is $\Theta(1/\sqrt{n})$, where n is the number of nodes in the MANET. In other words, the information carrying capacity becomes vanishingly small with increasing n. This result appears to be in tension with the efforts to build large-scale MANETs – *are MANETs scalable or not?*

In this paper, we present the first thorough examination of this dichotomy on MANET scalability. A key observation is that there are two interpretations of "scalability" – we call them *asymptotic scalability* and *in-practice* scalability. Asymptotic scalability, under which most of the work along the lines of [4] lies, refers to the order of growth in the limit of some metric (usually capacity) as a function of size. In contrast, we define in-practice scalability as the number of nodes (or other parameter) beyond which a network will not work "adequately" (later in this paper, we shall develop a formal definition of "adequately"). Unlike asymptotic scalability, which is typically unqualified (e.g. "Network X does not scale"), in-practice scalability is qualified ("Network X with parameter set P scales to 1000 nodes").

The well-known GK asymptotic scalability result is based on certain assumptions about the underlying network scenario. We examine these assumptions, survey and discuss work that shows the impact of relaxing some of these assumptions. In particular, we look in detail at the traffic distribution of a military MANET. Using data from the Future Combat System (FCS) network documentation[1], we analyze the approximate distribution of traffic with respect to the number of hops. We show that the traffic distribution is fairly localized, and appears to follow a *power law* distribution with an exponent between 2 and 3. This is significant because it has been shown [10] that a MANET with power-law distributed traffic with exponent greater than 2 scales as $\Theta(1)$, that is, it is *asymptotically scalable*.

We then consider general MANETs, with assumptions (including traffic) similar to the GK paper. Although by the GK result such MANETs are asymptotically unscalable, the number of nodes to which a particular instantiation can scale "in practice" is not clear. We present a definition of in-practice scalability based on the concept of residual node capacity and derive an expression for the in-practice scalability of a line network (e.g. a convoy). A realistic instantiation of the radio and network parameters along with an all-informed multicast voice traffic model - which is the most aggressive traffic distribution possible – shows that it can scale to over 5000 nodes even though asymptotically it is regarded "unscalable". While a convoy is only one kind of topology, this illustrates how a network can at once be asymptotically unscalable and be adequately scalable in

This paper was supported under U.S Navy SPAWAR Atlantic contract number N65236-08-D-5803/0137 with funding provided by the Joint Program Executive Office (JPEO) Joint Tactical Radio System (JTRS)

[1]While FCS is no longer a program of record, the studies done for FCS are still representative of military doctrine and relevant to our study.

practice. The in-practice scalability of a network depends on data rates, antennas and other factors that are being enhanced continually by new technology, promising a bright future for MANET scalability.

II. Asymptotic Scalability: A brief overview

In this section, we survey and interpret information-theoretic results on asymptotic scalability and provide the background for our contributions in later sections.

Asymptotic analyses deal with behavior "in the limit", or the "order" of growth of a quantity. The notation of $O(n)$, $\Omega(n)$ and $\Theta(n)$ indicate the upper, lower and "upper and lower" bounds on order of growth respectively. For example, if $f(n) = a \cdot n^2 + b \cdot n + c$, then $f(n)$ is $O(n^2)$ and $\Theta(n^2)$ – note that lower order terms and constants are ignored in asymptotic analysis.

We begin with the seminal Gupta Kumar result [4]. They considered the *transport capacity*, that is, the sum of products of bits and the distance they are carried over per unit time, of a MANET in a given area as the number of nodes is increased. They consider two kinds of networks – *random*, in which nodes are placed with a uniform probability in the plane; and *arbitrary*, in which one is allowed to place nodes anywhere. They consider two interference models – the *protocol model* in which interference depends on the distance, and the *physical model* which is based on pathloss-based Signal-to-Noise Ratio (SNR) at a receiver. Their results for the four combinations are shown in the table below.

Transport capacity results from [4]		
	Protocol Model	Physical Model
Arbitrary Nets	$\Theta(\frac{1}{\sqrt{n}})$	$\Theta(\frac{1}{n^{1/\alpha}})$
Random Nets	$\Theta(\frac{1}{\sqrt{n} \cdot log(n)})$	$\Theta(\frac{1}{\sqrt{n}})$

The most well-known of these results, namely the one for arbitrary networks using a protocol model, states that the per-node capacity falls off as $1/\sqrt{n}$ as n increases. To see this informally and intuitively, consider an equivalent and likely more realistic "expanding area" model (we assume a circle), where the density remains constant as the nodes increase[2]. Since nodes at some threshold distance apart can transmit simultaneously without causing interference, the one-hop capacity increases as $\Theta(n)$. However, the average number of hops is roughly proportional to the diameter of the circle, which increases as $\sqrt{n}$ (since $\pi \cdot (\frac{D}{2})^2 \sim= n$). Thus, the total network capacity is $\Theta(n/\sqrt{n}) = \Theta(\sqrt{n})$, which implies the per-node capacity is $\Theta(1/\sqrt{n})$.

We note that these results are agnostic to routing, Media Access Control (MAC) or any other control protocols. That is, *no matter what protocols are used*, these results hold. Further, increasing the bandwidth or adding more frequencies does not change the scalability [4].

[2]Gupta and Kumar show that the capacity is maximized when nodes transmit at the smallest power that keep them connected, which makes this model equivalent

Relaxation of GK Assumption	Asymptotic Scalability	Relevance/Comments
Stationary => Mobile	O(1)	Good news, relevant in the context of DTNs, but delay is unbounded
Omni => Directional	$O(1/\sqrt{n})$	Bad news, but constant is increased by *2π/ beamwidth*, increases capacity in practice
No coop => Coop. MIMO	O(1)	Good news, but cooperative MIMO is research area
No MPR => MPR	O(1)	Good news, but MPR requires sophisticated hardware
Uniform Traffic dist => Power-law distributed traffic	O(1) to $O(1/\sqrt{n})$ depending on exponent	Perhaps great news since military traffic is *not uniform?*

Fig. 1. Summary of scalability results some of the GK assumptions are relaxed. Refer to the text for citations containing the results.

The GK results are based on a number of assumptions: omni-directional antennas, stationary networks, no cooperation amongst the nodes, and uniformly random traffic patterns, single packet reception, to name a few. Post-GK research has looked at the scalability of MANETs under a relaxation of those assumptions. The results are summarized in Figure 1.

A randomly mobile network, interestingly, is asymptotically scalable [3]. Intuitively, this is because, given sufficient time, the destination node will come within a small number of hops of the source. However, this result does not bound the delay which may be excessively large in practice. Directional antennas only increase the constant [14], and as mentioned earlier, asymptotic results do not change with constants. However, constants do matter in practice – we revisit this in section IV. Use of sophisticated physical layer techniques – distributed MIMO [11], Multi-Packet Reception (MPR) [2], etc. change the effect of interference and have been shown to asymptotically scale. Finally, it has been shown that if the traffic is non-uniform, in particular if it follows a "power law", then the network may scale. We investigate this in the next section.

III. Scalability with Military Traffic

In this section, we examine the nature of military traffic, and its distribution as a function of the number of hops. In particular, we have taken the Future Combat System (FCS) Brigade Combat Team organization and doctrine, and traffic studies developed by the Capability Development and Integration Directorate-Gordon (CDID-G) and, using some reasonable assumptions, estimated the fraction of traffic as a function of the hop-distance of the traffic. As most FCS documentation is FOUO (For Official Use Only), we shall only present abstracted information in this paper.

A. Methodology

The CDID-G traffic model groups traffic by *echelons*, namely, fire team, squad, platoon, company, battalion and brigade. Thus, we are given what a team transmits to a squad, what a squad transmits to a platoon, what a platoon transmits to a company, etc. The number and type of messages, frequency of transmission, and specific "To" and "From" addresses are aggregated into a bandwidth requirement for traffic between each echelon and every other echelon in the brigade. The advantage of this model over other typical models such as force-on-force and scripted models is that it is not tied to any specific situation or, in the case of force-on-force simulations, the skill level of players. It is the combination of the results of many such simulations and therefore generic enough for our study.

Mapping the traffic distribution within and between echelons to the number of hops depends on three factors: the nominal spacing of elements within an echelon, nominal spacing of elements between echelons, and the radio range to connect elements.

To determine a realistic hop count, the differences in the ranges of the radios found in each echelon must be taken into account. For example, a radio found in a rifle squad does not have the same range as a radio found on a battalion command and control vehicle. We chose two nominal ranges to represent the 12 types of terrestrial radios currently found in a BCT – 1 kilometer for platoon and below radios, and 10 kilometers for company and above radios.

We assumed that traffic generated at the lower echelons (i.e., team) traveling to the higher echelons (i.e., brigade) would use the lower range initially and then progress to the higher radio range as the message passed through the company to the battalion. We also assumed that airborne or satellite links were not used.

Depending on how the brigade elements were deployed, we identified two equally realistic scenarios, which result in somewhat different distributions and are interesting to contrast in this paper.

- *Consolidated*. In this it was assumed that the brigade command elements are concentrated at one place.
- *Evenly distributed*. In this it was assumed that the brigade fields three tactical operations centers (Main, Forward, Rear) and brigade traffic is dispersed between them.

B. Traffic Distribution and Analysis

The Consolidated (CD) and Evenly Distributed (ED) traffic percentages for each hop are shown in Figure 2. It is clearly evident that there is a concentration around the lower numbers of hops. While it is obviously not uniform, it does not appear to be steep enough to be exponential. We are therefore led to consider a *power law* distribution which is not as steep as exponential, and may have theoretical connections with hierarchies [7]. In what follows, we first describe the power law, and then discuss the data in Figure 2 vis-a-vis the power law.

Informally, an event is said to be power-law distributed if the probablity of the event having a value x is proportional to $x^{-\alpha}$, where α is the *power law exponent* that captures how quickly the probability falls off – the higher the power law exponent, the quicker the probability falls off with increasing x [13]. Many natural pheonomena such as earthquakes, wealth distribution, population of cities, etc. have been shown to follow a power law.

Consolidated Scenario

Hops	% Total Traffic
1	83.6
2	0.3
3	10
4	4.5
5	0.3
7	1.3

Evenly Distributed Scenario

Hops	% Total Traffic
1	28.7
2	27.8
3	37.4
4	4.5
5	0.3
7	1.3

Fig. 2. Aggregated traffic distribution as a function of number of hops, for FCS BCT. No traffic went 6 hops.

Let h denote the number of hops (first column in the figures) and $P(h)$ the probability that traffic goes h hops (second column divided by 100). By definition,

$$P(h) = C \cdot h^{-\alpha} \quad (1)$$

where C is some constant.

Since the sum of probabilities must equal 1, we have

$$\sum_{h=1}^{7} (C_\alpha \cdot h^{-\alpha}) = 1 \quad (2)$$

Clearly, the constant will depend on α. To proceed further, we need to determine approximately what α we are interested in. As we shall see later, α = 2 appears to be a thresholding point in terms of scalability [10]. Therefore, we consider α = 2 and α = 3. Expanding and solving for C_α, for α = 2 and 3, we get C_2 = 0.66, and C_3 = 0.83. Plugging these back in equation 2 we get "true" power law distributions (that is, the distribution if a 7 hop network followed power law precisely) for exponent 2 and 3 respectively as $P_2(h) = 0.66 \cdot h^{-2}$ and $P_3(h) = 0.83 \cdot h^{-3}$.

We plot the actual probabilities for each hop and compare with the probabilities based on "true" power law distributions with exponent 2 and exponent 3. The way to do this is to take the cumulative distribution function and plot it on a log-log (doubly log) scale [13]. That is, we plot $P^C(h) = log(P(H > h))$ versus $log(h)$. Figure 3 shows the Consolidated scenario along with cumulative counterparts of P_2^C and P_3^C. Figure 4 shows the Evenly Distributed Scenario along with the cumulative P_2^C and P_3^C.

C. Implications on asymptotic scalability

The GK results assume that traffic is uniformly distributed in terms of hop distance. The analysis for the FCS BCT done above shows clearly that this is far from true, and that there is a heavy localization. The level of localization depends upon the scenario, but even in the ED scenario, there is only

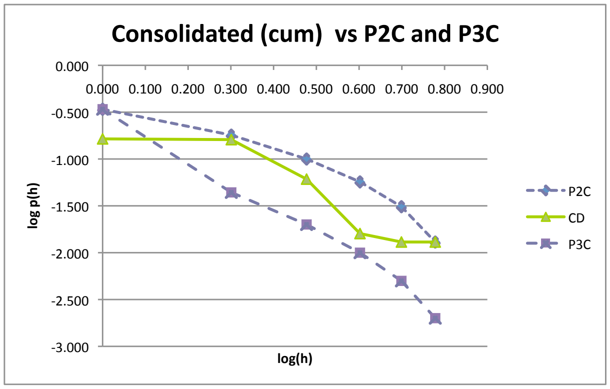

Fig. 3. Comparison of CD scenario with power law distributions of exponent 2 (upper dashed) and 3 (lower dashed). Cumulative probabilities plotted on doubly log scale.

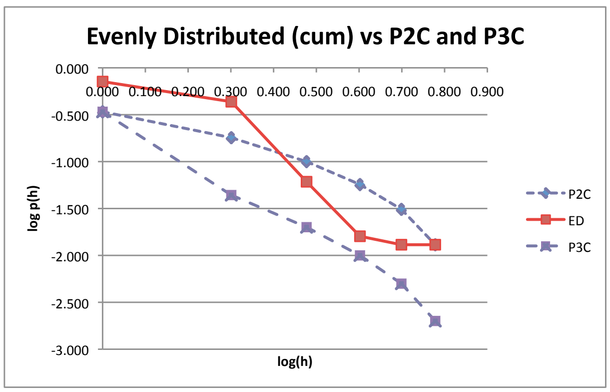

Fig. 4. Comparison of ED scenario with power law distributions of exponent 2 (upper dashed) and 3 (lower dashed). Cumulative probabilities plotted on doubly log scale.

a small fraction of traffic at larger hops. Thus, GK results do not seem to apply to military traffic scenarios.

Our analysis of the distribution vis-a-vis the power law distribution, shows up some clear relationships. The CD scenario seems to fit fairly well with a power law distribution with an exponent somewhere between 2 and 3. The ED scenario is less of a fit, but we note that *after* 3 hops, it is sandwiched between P_2^C and P_3^C. As per a more precise definition of power law, it is sufficient that there exist a h' above which the distribution is a power law [13].

As mentioned earlier, theoretical results on scaling law with power law distributed traffic is different from that for uniform traffic. In [10] it is shown that the scaling has four distinct regimes depending on the power law exponent. Specifically, if $\alpha < 1$, then the per node capacity scales as $O(1/\sqrt{n})$, which is identical to GK result for uniform traffic; if $\alpha = 1$, then it is $O(ln(n)/\sqrt{n})$ and if $1 < \alpha < 2$, it is $O(n^{\frac{(\alpha-2)}{2}})$, and when $\alpha = 2$ it is $O(1/ln(n))$. These are better scaling properties than with uniform, but still *unscalable* in the asymptotic sense; but if $\alpha > 2$ then [10] shows that the capacity scales as $O(1)$, that is, it is asymptotically scalable.

This last result is very promising in light of our observations – the exponent need not be very high, just greater than 2 (even marginally) for asymptotic scalability. The FCS BCT traffic does seem to have this property, and is an encouraging sign. However, our analysis was based on a single set of data and the number of hops is limited. Therefore, we cannot at this point definitively conclude that all military traffic is power law with exponent greater than 2, without analysis of more data sets. Nonetheless, it is clear that at the very least GK results need to be revisited for military networks; and there is some evidence that military traffic might well result in asymptotically scalable networks.

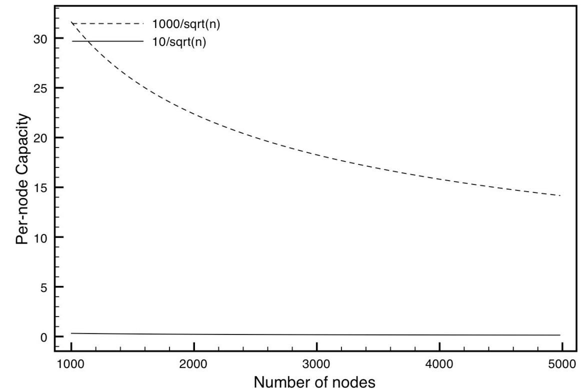

Fig. 5. $C/\sqrt{n}$ for C=10 (lower) and C=1000 (higher) for brigade sized networks. Asymptotically, they are the same, but it is clear that in practice C=1000 is significantly better.

IV. In Practice Scalability

The previous section showed that the FCS BCT and other such echelon-based networks are likely to have power-law-distributed traffic with exponent greater than 2, which by [10] makes them scale asymptotically. However, non-military MANETs and military MANETs with traffic distribution unlike the ones analyzed above may not have power-law-distributed traffic, or the exponent may be less than 2. In this section, we discuss the number of nodes to which *general* MANETs - regardless of the traffic distribution – can scale to. A network may be asymptotically unscalable, yet scale to several thousands or tens of thousands of nodes in practice, depending on the *constants* in the expressions.

To see the impact of constants, consider two networks with per-node capacity $O(1/\sqrt{n})$. In one the constant is 10, and the other is say 1000. Figure 5 shows how these curves look on brigade-sized networks. Clearly, for the military network practitioner, the network with a constant of 1000 is significantly better (e.g. at 3000 nodes, it gives a 20x improvement in capacity). Thus, even though these networks have the same asymptotic scalability (the curves meet up at some very large node number), it is obvious that for the purposes of fielding military MANETs, there is a major difference which would have been missed had we only looked at asymptotic scalability.

As another example, consider the use of directional antennas in MANETs, [12] which effectively increases the constant. It has been shown [14] that a random MANET with transmit and receive beamwidth of β increases the capacity by $4\pi^2/\beta$. Consider a random MANET that scales to 100 nodes with omni-directional antennas. A 30 degree beamwidth would make it scale to 14,400 nodes, which is entirely acceptable in practice. At the same time, per asymptotic scalability definition, this is an unscalable network.

In general, asymptotic analysis is not suitable for answering several questions that are of practical interest: how many nodes does a given network scale to, what is the effect of different topology classes, traffic and physical layer parameters, which parameters most affect my network's performance, etc. While there are specific studies [5], and Gupta-Kumar [4] do work out the actual expressions with constants enroute to the asymptotic results, there is no definition and framework for studying such *in-practice scalability* in general. In the remainder of this section, we introduce a simple definition of in-practice scalability and apply it to a line network (e.g. convoy).

We begin with a few definitions. Consider a node m in the MANET. The *available capacity* $\mathsf{A}(m)$ indicates the amount of data that can be handled by m. The *used capacity* $\mathsf{U}(m)$ denotes the amount of data load at m in transporting a required set of flows. The *blocked capacity* $\mathsf{B}(m)$ denotes the capacity that is unusable by node (for example, due to contention or energy considerations). The *residual capacity* $\mathsf{R}(m)$ is the difference $\mathsf{A}(m)$ - $\mathsf{U}(m)$ - $\mathsf{B}(m)$.

The scalability of a network depends fundamentally on the balance between available and used resources. Clearly, the network can support the offered flows if and only if the residual capacity at every node in the network is positive. We assume that all nodes are loaded equally on average, for instance, by the use of a load-balancing routing protocol. Heterogeneous loading can be considered by identifying the most congested node and computing its used capacity, but this is beyond the scope of this paper.

Thus, it suffices to take any sample node and determine if its residual capacity is positive. Thus, we drop the m in the paranthesis for the remainder of this discussion.

Both U and B may have multiple components. For example, U typically consists of traffic load, routing control overhead, MAC overhead, etc. Channel contention is a key component of B . Denoting the j_{th} component with a subscript j, we have

$$\mathsf{R} = \mathsf{A} - \sum_j \mathsf{U}^\mathsf{j} - \sum_j \mathsf{B}^\mathsf{j} \tag{3}$$

Now, for each component j, assuming node homogeneity, the contention-based blocking is a result of a certain number of nodes wanting to do the same thing as the considered node. Thus,

$$\mathsf{B}^\mathsf{j} = \gamma_j \cdot \mathsf{U}^\mathsf{j} \tag{4}$$

Then,

$$\mathsf{R} = \mathsf{A} - \sum_j (1+\gamma_j) \cdot \mathsf{U}^\mathsf{j} \tag{5}$$

We call γ_j the *contention factor*. In wireless networks, the contention factor depends upon how many nodes are in the neighborhood of a node, and on the medium access control protocol.

We consider a system to be able to work adequately if and only if every single node in the system is able to withstand the load on average, that is residual capacity is positive. This brings us to the definition of in-practice scalability.

Definition IV.1. *The* in-practice *scalability of a network is the number of nodes X such that for all $n \leq X$, $\mathsf{R} > 0$, and for all $n > X$, $\mathsf{R} \leq 0$*

As an application of the above, consider a line network, representing, for instance, a convoy of vehicles. Suppose the network has N nodes, is topologically stationary (the convoy may move but the links between nodes aren't broken), each node has X transceivers each with a rate r, and packet error probability is e. Suppose each node originates a flow of L bps at duty cycle d that is *multicast*[3] to an "all-informed" group, that is, to all nodes in the network. We assume that the network uses a link state routing protocol with overhead of Q bps orignated per node. We ignore the medium access control overhead (typically, this does not increase with size as long as the density is about the same).

By the GK result, the above network is asymptotically unscalable since it is an "arbitrary" network and meets the assumptions made in the analysis, including the uniform traffic assumption. What is the in-practice scalability of this network? We give an approximate analysis of this below.

Using equation 5, we have, for a typical node,

- $\mathsf{A} = r \cdot X$. We assume that the X transceivers are assigned X orthogonal frequencies and hence the capacity is essentially multiplied by X (one can do better than this if more than X frequencies are available, but we take a conservative view).
- For the traffic flow (F) component, $\mathsf{U}^F = L \cdot (N-1) + L = L \cdot N$ since every node needs to forward L bps from every other node and from itself. Further, $\gamma_F = 2$ since the receiver and its neighboring node must defer.
- For the control overhead (O) component, since link state floods, $\mathsf{U}^O = N \cdot Q$, and since Link State Updates (LSUs) are broadcast, $\gamma_O = 4$ since for each of the sender's two neighbors, they as well as *their* neighbors must defer so that both neighbors get the LSU.

Each packet constituting L is subject to an error probability of e and needs to be retried. The expected number of transmissions is therefore $1/(1-e)$.

[3]We use multicast rather than unicast for two reasons: first, military traffic is overwhelmingly multicast; second, it represents a more aggressive demand than unicast and represents a worse case.

Further, the effective load is L multiplied by the duty cycle. Substituting in equation 5.

$$\mathsf{R} = r \cdot X - (1+2) \cdot (\frac{L \cdot d}{1-e} \cdot N) - (1+4) \cdot N \cdot Q \quad (6)$$

Using definition IV.1 and equation 6, the in-practice scalability of the line network is

$$N = \frac{r \cdot X}{\frac{3 \cdot L \cdot d}{1-e} + 5 \cdot Q} \quad (7)$$

Let us consider a reasonable realistic instantiation of the above. We assume r = 10 Mbps, number of transceivers X = 4, L = 260B voice packets at 6.25 pps which, including about 80B worth of IP, subnet and PHY headers amounts to approximately 17 kbps, d = 10%, e = 0.05, and Q = 206B LSUs every 5 seconds = 329.6 bps.

Substituting in equation 7, we have N = 5701. That is, a line network with the above parameters can scale to 5701 nodes, even though, per the Gupta-Kumar result [4] it is asymptotically unscalable. Even with only one transceiver, it scales to more than 1000 nodes. We made pessimistic assumptions (only four frequencies, only 10 Mbps radios, and all-informed multicast), so in general this number will be even higher. This simple exercise demonstrates how a network may be asymptotically unscalable and at the same time adequately scalable in practice. Our ongoing work, beyond the scope of this paper, extends such results to other network types such as grid, random, etc. as well.

V. Concluding Remarks

We have shown that the well-known *asymptotic scalability* results from Gupta-Kumar[4] have limited relevance to the scalability of real military systems, for two broad classes of reasons. First, these results make certain assumptions which are not likely to be true for military networks. Specifically, they assume that the traffic distribution over number of hops is uniform. In contrast, our analysis of the traffic distribution of the FCS BCT under two scenarios shows that the traffic is more likely to be power law distributed with an exponent between 2 and 3. Based on another asymptotic result [10], such a distribution would result in MANET scalability even in the asymptotic sense.

Second, asymptotic analysis does not answer the question of whether MANETs – regardless of the traffic assumption – can scale to sizes of most interest to the military, namely brigade-sized networks (1000 - 5000 nodes). We introduce the notion of *in-practice scalability* as a formal way of studying finite-sized network behavior in the non-asymptotic sense. While asymptotic scalability considers generic questions (“Does network X scale?”), in-practice scalability considers the specific (“How many nodes does does network X scale to?”). An analysis of the in-practice scalability of a line network (representing a convoy for example) with realistic parameters and an aggressive, non-power-law traffic model (all-node multicast) still gives a conservative estimate of scalability to 5700 nodes in practice.

Thus, MANETs in general may well be adequately scalable in practice even if they are asymptotically unscalable, and military MANETs may also even be asymptotically scalable by virtue of their traffic characteristics.

Much work needs to be done to make MANETs scale to brigade-sized networks, particularly in the domain of control overhead, and good engineering to make a system work as designed. We also need to understand in-practice scalability and military traffic better, and experiment with large-scale MANETs. However, as we have shown in this paper, the results of Gupta and Kumar [4], “negative” as they may be, do not by themselves weaken our chances of building MANETs that are sufficiently scalable to military needs.[4]

References

[1] J.A. Freebersyser and B.M. Leiner. *A DoD perspective on mobile Ad hoc networks*, pages 29–51. Mobile Ad Hoc Networks, Edited by C. Perkins, Addison-Wesley, Boston, MA, USA, 2001.

[2] J. J. Garcia-Luna-Aceves, H. R. Sadjadpour, and Z. Wang. Challenges: Towards truly scalable ad hoc networks, 2007.

[3] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. *IEEE/ACM Transactions on Networking*, 10:477–486, 2002.

[4] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46:388–404, 2000.

[5] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *Proc. MobiCom '03*, pages 66–80, 2003.

[6] J. Jubin and J. Tornow. The darpa packet radio network protocols. *Proceedings of the IEEE*, 75:21–32, 1987.

[7] D. Krioukov, F. Papadopoulos, M. Boguna, and A. Vahdat. Greedy forwarding in scale-free networks embedded in hyperbolic metric spaces, 2009.

[8] M. Lawlor. Wireless to the nth degree. *AFCEA SIGNAL online*, July 2006.

[9] B.M. Leiner, R.J. Ruth, and A.R. Sastry. Goals and challenges of the darpa glomo program. *IEEE Personal Communications*, 3:34–43, 1996.

[10] J. Li, C. Blake, D. S. J. De Couto, C. Hu, H. I. Lee, and R. Morris. Capacity of ad hoc wireless networks. In *In ACM Mobicom*, pages 61–69, 2001.

[11] A. Ozgur, O. Leveque, and D. N. C. Tse. Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks. *IEEE Trans. Inf. Theory*, 53:3549–3572, 2007.

[12] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit. Ad hoc networking with directional antennas: a complete system solution. *IEEE Journal on Selected Areas in Communications*, 23:496–506, 2005.

[13] Wikipedia. Power law — wikipedia, the free encyclopedia, 2010. [Online; accessed 21-April-2010].

[14] S. Yi, Y. Pei, and S. Kalyanaraman. On the capacity improvement of ad hoc wireless networks using directional antennas. In *In 4th ACM MobiHoc*, pages 108–116, 2003.

[4]Research in section IV was partially sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.