Throughput of Ideally Routed Wireless Ad Hoc Networks

Gábor Németh*, Zoltán Richárd Turányi, and András Valkó
Traffic Lab, Ericsson Research, Hungary

We investigate the throughput of ad hoc networks using ideal shortest path routing between randomly selected source-destination pairs. Our results not only confirm the $n^{-1/2}$ decay of the throughput that was published earlier in the literature, but also provides an approach to acquire the relation of various network parameters. This way the effects of changes in the routing, traffic generation algorithms or other network characteristics can be predicted using a first order approximation.

I. Introduction

At the physical layer the capacity of ad hoc wireless networks is constrained by the mutual interference of concurrent transmissions between nodes. We study an ad hoc network model where n nodes communicate in random source–destination pairs. *Gupta* and *Kumar* [3] showed that for static random ad hoc networks using a general routing algorithm the capacity available for each node decays as $\frac{1}{\sqrt{n}}$. Other works [1, 2, 5, 6] delve into the problem of optimizing various parameters of the transmission (e.g., power consumption or medium access control), and try to devise routing protocols that fit for particular user profiles or scenarios on the same network.

Our present study focuses on the general properties of the per node throughput available for a fixed ad hoc wireless networks (capacity) using an ideal routing process. The nodes in our model do not move, which modifies the customary definition of an ad hoc network to a backboneless network of wireless nodes occupying a flat topology. Our network model includes ideal collision avoidance and transmission through shortest paths as explained in the following Section. We introduce an alternative description of network throughput approximation that verifies the claims of [3] and extends the results by providing the relations of the various network parameters that can change with topology or traffic generation algorithm. We have investigated the throughput of various non-planar network topologies, and the results generalize the $\sqrt{n}$ dependence of the average call length parameter. We also check the validity of our model by simulation and give examples of application.

The rest of this paper is structured as follows. Section II gives a detailed description of the traffic and radio models we used in our work. Section III introduces our network capacity description model and emphasizes the consistency with the results of [3]. Section IV presents simulation results for changes in various network parameters and how their effects can be interpret using the model of Section III. Finally, Section V summarizes our findings.

*E-mail: gabor.nemethi@eth.ericsson.se

II. Model

Our simulations use a stationary ad hoc network where the randomly located nodes are distributed independently and uniformly over a rectangular area of unit size. To determine the neighborship relations we introduce the following definitions.

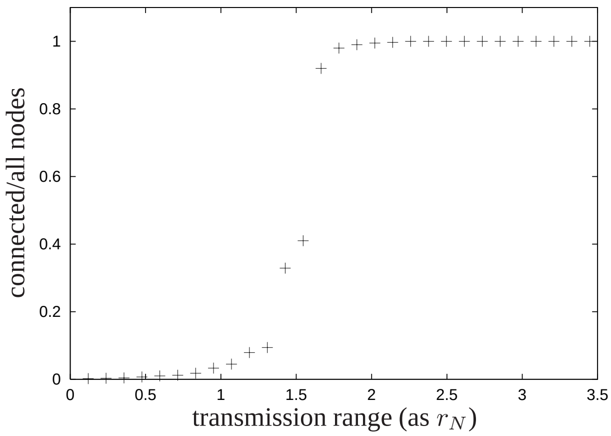

Figure 1: The relative size of the largest connected component is plotted against the transmission range. *(quoted from [4])*

The *transmission range* (r) is the maximum distance of radio communication. Two nodes are said to be neighbors if they reside within transmission range. Simple connectivity tests show that in the case of small transmission range the communication graph is not connected in general. This means that many nodes are not able to communicate with all other nodes. The transmission range that causes the network to be connected at an acceptable level – e.g., 99.9% of the nodes are in one connected component – is a function of the number of nodes. Let us now introduce a parameter called the *normalized transmission range* r_N:

$$r_N = \frac{r}{r_0} \text{ and } r_0 = \sqrt{\frac{1}{n\pi}} \tag{1}$$

where n is the number of nodes and r_0 denotes the radius of the planar disk, whose area is equal to one node's share of the unit area. The normalized transmission range determines the connectivity of the graph irrespectively of the number of nodes. Our experience shows that for $r_N > 2.7$ the graph is connected with high probability (see Figure 1 for the relationship of the transmission range and the connectedness of the network).

In our model all nodes are identical, each one has the same transmission range equal to the normalized transmission range. It may happen that there are isolated nodes, that are outside the transmission range of all other nodes.

We model traffic as constant bit rate calls. Every node attempts to initiate the same number of calls toward randomly selected destinations. Each successful call occupies 1 unit of capacity at each affected node.

Our radio model is as follows. Every node is able to transmit and receive data in a fixed amount of capacity units. We assume that the ideal arbitration allows the nodes to fully utilize the available capacity for data transfer. This way no overhead occurs due to collision avoidance. If a node transmits, all nodes within its transmission range – including itself – use up 1 unit of capacity. The calls go hop by hop and therefore every intermediate node needs to spend three units of capacity for a single call. Once when it receives the call from the previous node in the path, second time when it transmits towards the destination, and for the third time when the next node in the path transmits the data further. This roughly models a radio technology with one common channel shared by all nodes, similar to CSMA/CA.

Neither the topology nor our traffic model changes with time. We study the throughput of the network at independent moments of time. The simulation starts by distributing the nodes, and after determining the neighborship graph, we calculate the shortest path routes using Dijkstra's algorithm.

If node i intends to transmit to node k, then after determining the shortest path from *i* to *k*, admission control is performed to check if every node in the path and all their neighbors have enough free capacity for this call. If not, then the call is blocked.

It is hard to find a single quantity that can meaningfully describe the throughput of a network, even for such a simple model, as the one used in the present paper. We have found that one way to characterize the throughput is by the number of successful calls delivered by the system. At a given blocking probability b the total number of successful calls N^b_{calls} is reproducible for a given set of network parameters, and is relatively easy to interpret (in the following index b will be omitted). Thus we have chosen N_{calls}/n to represent the throughput of the system, and will call it the *per node critical load*. In order to determine N_{calls}/n, we carry out subsequent simulations with different call intensities and calculate the resulting blocking probability. The *offered load*, that is, the number of call attempts is adjusted in iterative steps until call blocking is close to the fixed blocking value. The output of the simulations are the per node critical load of the system and the average path length of the calls.

III. Analysis

We present an alternative approach to [3] for measuring the maximum traffic a stationary ad hoc network can deliver. Our heuristic model can be adapted for various surface topologies or changes in the traffic generation algorithm, and the model is able to provide the relation of the parameters used.

The idea for approximation of the capacity for given values of the network parameters is based on the averages of node utilization: u_i is defined for every node i as the number of capacity units used up by the node. The distribution of u_i is independent of the neighborship realization if the blocking probability and normalized transmission range are fixed, and thus we take the expectation value: $\bar{u}$. The total number of utilization units in the whole system is calculated as follows. The distribution of call lengths is also constant and its average is $\bar{k}$ (in hops). A single call of average length will use up $\bar{k}\, r_N^2$ units of capacity thus the total number of units needed equals the average utilization times the number of nodes:

$$\bar{k}\, r_N^2\, N_{\text{calls}} = \bar{u}\, n, \text{ or}$$

$$\frac{N_{\text{calls}}}{n} = \frac{\bar{u}}{r_N^2\, \bar{k}} \tag{2}$$

For a given simulation run the parameter $\bar{u}$ can be determined and the approximation for N_{calls}/n can be calculated.

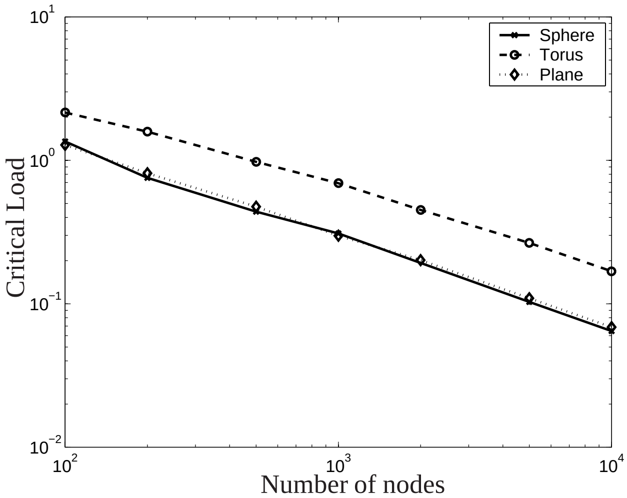

Figure 2: Per node throughput decays $\propto 1/\sqrt{n}$ independently of the arena topology.

We can see that for fixed r_N the capacity decreases as $\frac{1}{\sqrt{n}}$ with the number of nodes. This can be attributed to the fact that the average path length of the calls increases with $\sqrt{n}$ thus one call uses up more and more capacity. An approximation of the overall throughput parameter can be given based on geometrical considerations as follows.

Assuming that $\bar{u}$ is independent of n (2) yields the $\frac{1}{\sqrt{n}}$ rule as our first result. If the definition in (1) is substituted and knowing that $\bar{k} \propto \sqrt{n}$ we obtain

$$\frac{N_{\text{calls}}}{n} \propto \frac{1}{\bar{k}}$$

In Figure 2 the per node critical load is plotted against the number of nodes for a normalized transmission range of 3.0. Besides for the planar square of unit area the Figure also displays the results for nodes scattered on a unit–surface sphere and torus scenarios which are to represent different surface topologies. For the case of the sphere the capacity equals to that of the planar square – this was proven in [3] analytically. Actually the average path lengths in the $n \to \infty$ limit do not differ for the two cases. This is not so for the case when all nodes at the border of the square are in the neighborship of the nodes at the opposite border of the square: the plane is folded into a torus. In this case the per node critical load still decreases with $\frac{1}{\sqrt{n}}$ but because $\bar{k}$ is halved, the critical load is twice than in the planar case, according to (2).

IV. Simulation results

Simulation runs of the model were carried out in order to study the impact of various parameters on the capacity of the system. Unless otherwise stated all simulations are run with parameters $n = 1000$, $b = 10\%$ and $r_N = 3.0$.

IV.A. Power Control

In this scenario we analyze the effect of an ideal power control implementation. In this case the hosts are able to change their transmission range according to the demand: they transmit with power just enough to reach the next node. The neighborship relations are still calculated using the normalized transmission range, which serves as maximum transmission range in this case, and actual transmissions affect only those nodes that are closer to the transmitting node then the next node in the path.

With the use of (2) we can predict the effect of the power control: the transmission ranges reduce, which causes the right hand side to increase. Figure 3 shows the per node critical load parameter plotted against the size of the system. The results without power control are also given for comparison. It appears that the use of the power control mechanism can increase the capacity of the ad hoc network by as much as 35% for a given number of nodes but the system retains its overall worsening property with $n \to \infty$.

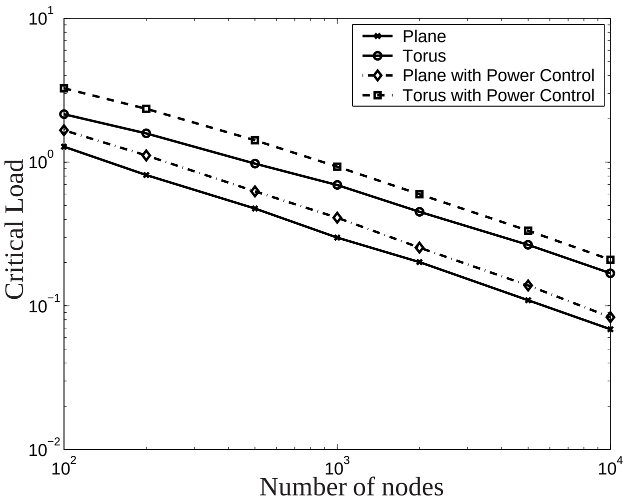

Figure 3: Using power control can result in +35% in the per node throughput, still the $\propto n^{-1/2}$ decrease remains as expected.

IV.B. Alternative Routing

According to Section II the routing algorithm is rather naive: if a node becomes blocking, then all subsequent calls that need to hit this particular node will be blocked as well. The above effect may cause a significant amount of network capacity never to be used, which behaviour can be corrected by applying alternative routes.

We extend our model using a heuristic alternative routing algorithm as follows. At the beginning of every simulation run the shortest paths for all $(src, dest)$ pairs are calculated. During the admission control step if a node on the path does not have enough free capacity for the call, then instead of counting the call as blocked we simply recalculate the shortest path using non-blocking nodes only. This way a call is blocked only if $dest$ is unreachable from src through non-blocking nodes.

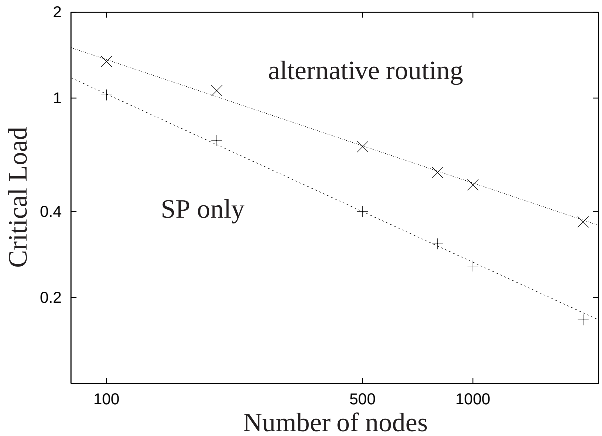

Figure 4: Switching from "shortest paths only" to "alternative" routing makes significant difference in the exponent of throughput decrease with increasing number of nodes.

As the first application of the extended routing algorithm we re-visit the critical load vs. system size relationship. In the case of alternative routes we presume that the nodes will be able to utilize more of their capacity:

$$\bar{u}_{\text{AR}} > \bar{u}_{\text{SP}}$$

Of course using the alternative routes will introduce the increase of the average path length $\bar{k}_{\text{AR}}$, yet the longer routes will have an exponentially dropping weight (see Figure 5 and explanation in Section IV.C) and thus we can safely neglect the increase of $\bar{k}$ compared to that of $\bar{u}$. This way it is a natural conclusion that the load carried by the system will be higher in amplitude, and (2) becomes

$$\frac{N_{\text{calls,AR}}}{n} = \frac{\bar{u}_{\text{AR}}}{r_N^2 \, \bar{k}} > \frac{N_{\text{calls,SP}}}{n} \quad (3)$$

Additionally with the increase of n there will be an exponentially growing number of alternative paths, thus the advancement of N_{calls}/n will increase as well.

Figure 4 summarizes the results of the measurements: a significant increase in both the value of the per node critical load and its derivative can be observed. The average critical load is $\propto n^{-0.59}$ for "shortest paths only" and $\propto n^{-0.43}$ in the case of alternative routes.

IV.C. Path Length Distributions

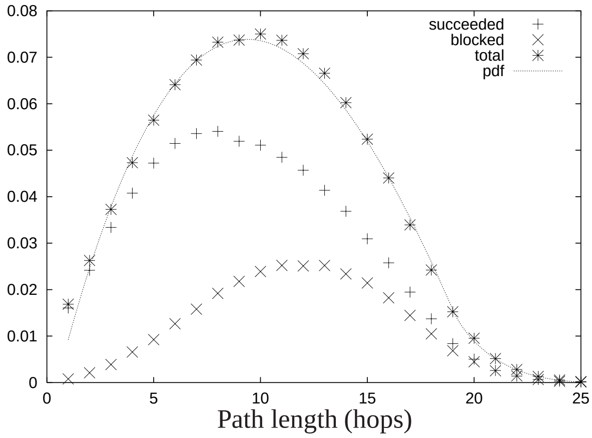

Figure 5: Typical distribution of the length of succeeded, blocked and all calls. Plot is generated using a network of $n = 500$, $r_N = 3.0$ and $b = 30\%$.

When a call has been built up between the source and destination nodes, the length of the call path is measured by the number of hops the call needs to travel. In the present Section the observed probability distributions of successful and blocked calls are to be discussed.

The path length distribution of succeeded and blocked calls can be compared to the same distribution for all calls. In a previous work [4] we have shown that the hop count distribution of all calls is perfectly fitted by the probability distribution of the distance of two randomly chosen points in $[0;1] \times [0;1]$. This property does not change with the above change of routing, as it is shown on Figure 5.

Analyzing the Figure one can notice that the expectation value of blocked calls is shifted to higher values, whereas that of the successful ones is less than the average call length. This observation confirms that in principle both *pdf*s are generated by skewing the "original" *pdf* of all calls (labeled *total* on the Figure). The probability that a call of length k is non-blocking (or blocking) is roughly

$$\begin{aligned} P_{\text{nonbl}}(k) &\simeq (1-p)^k \\ P_{\text{bl}}(k) &\simeq 1-(1-p)^k \end{aligned} \qquad (4)$$

respectively, where $p \in [0;1]$ depends on the network parameter settings.

In order to obtain the plot data on Figure 5, the *pdf*s in (4) shall be multiplied by *total*, which defines the frequencies of calls of a given length. As an example Figure 6 displays the *succeeded* and *blocked* data sets borrowed from Figure 5 and the *pdf*s $P_{\text{nonbl}}(k)$ and $P_{\text{bl}}(k)$ using $p = 0.032$.

IV.D. Node Utilization Distributions

Next some details of the node utilization parameter u are discussed. In (2) we used an average for this parameter over all nodes and assumed it to be independent of the system size. Figures 7 and 8 demonstrate the differences in the u-histograms made by the changes in n and r_N respectively. The former of the two figures demonstrates how $\bar{u}$

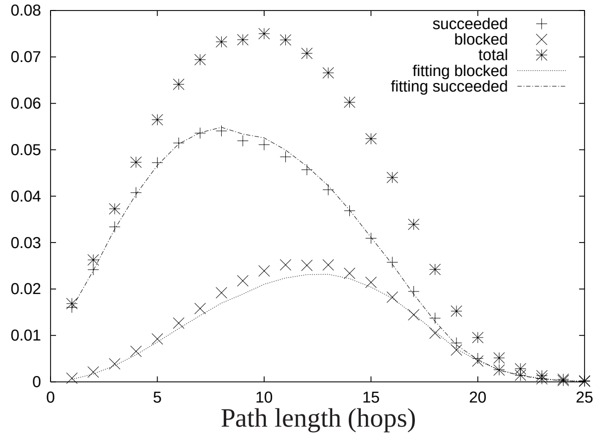

Figure 6: Deriving distribution of succeeded and blocked calls, see (4). $n = 500$, $r_N = 3.0$, $b = 30\%$.

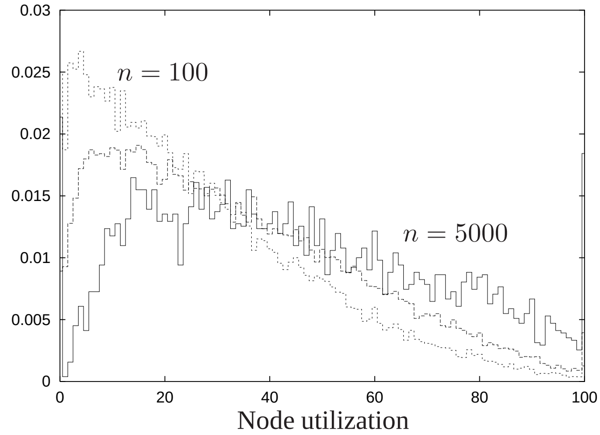

Figure 7: Variation of the node utilization distribution with the system size. The basic shape of the curve remains the same while the average utilization grows with increasing Ns.

depends on the system size: the expectation value seems to slowly shift as the network size increases. This behaviour is explained by the border effect of the arena: nodes residing near the border of the unit square are causing a systematic decrease in N_{calls} because they are affected by less calls than an average node near the center of the square (the border nodes have under-average number of neighbors and thus is less the number of calls hitting them). This is why $N_{\text{calls}}/n \propto n^{-0.59}$ instead of $\propto \sqrt{n}$.

Figure 8 demonstrates the effect caused when the transmission range is increased. According to Section III all transmissions affect $\propto r_N^2$ nodes, thus causing the average utilization to grow likewise. The weight of heavy loaded nodes increases with r_N as it is clearly visible on the Figure.

A structural difference can be noticed while comparing the u-histograms measured for the two routing algorithms. The average for alternative routing increases compared to the basic shortest paths algorithm, but the main difference – as it is seen on Figure 9 – is that instead of having a peak at low utilization values ($u \simeq 10$) and a heavy tail, the

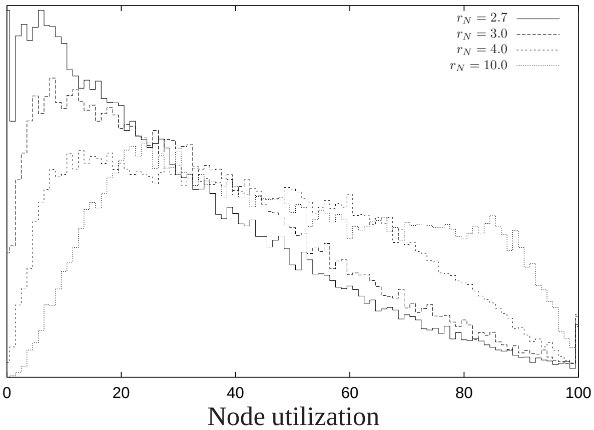

Figure 8: Histograms of node utilization plotted for different radio ranges. For high radius values the ratio of more heavily loaded nodes increases.

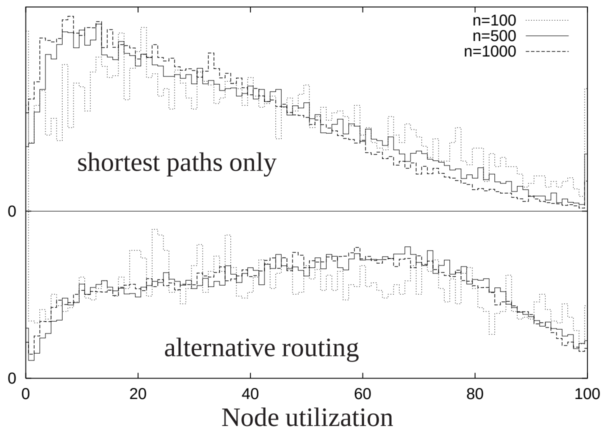

Figure 9: Utilization histograms for both routing methods. The intensities are measured in the same units on both graphs.

increased load gets more distributed: a plateau of values is formed in the range $15 \leqslant u \leqslant 75$. The present description lacks the proper modeling of the observation, but we are to address this issue in the future.

IV.E. Localization of Traffic

The decrease in throughput seems to persist for all investigated cases of n identical agents that are calling all other nodes with equal probabilities. It is possible, however, for a different communication behaviour to "break" the throughput barrier. Note that (2) yields that if $\bar{k}$ is constant, then $N_{\text{calls}} = \text{const}$ also holds, because $r_N^2 \propto n$. This way we have only to find a way that makes $\bar{k} = \text{const}$.

Let us change the traffic generation algorithm in such a way that prefers short length calls. In this scenario calls are going to be generated with a certain probability, which depends on the distance spanned by the shortest path connecting the source and destination nodes. Let h denote this path length in hops and let the probability function be the

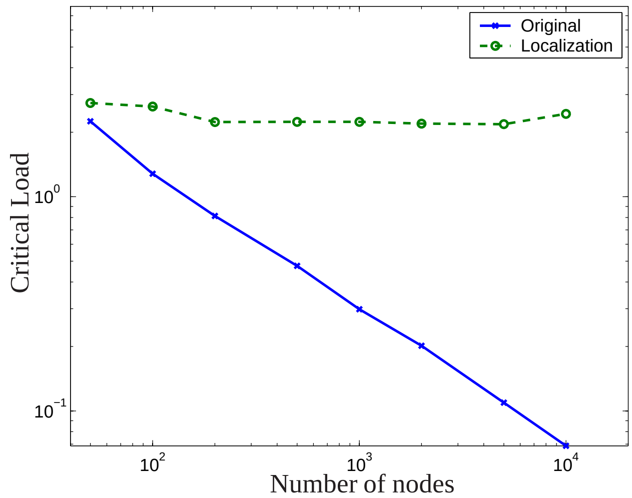

Figure 10: The throughput of the localized network is not effected by the changes of the system size.

standard one:

$$f(x) = e^{-\frac{h^2}{2\,\sigma^2}}$$

where σ plays the role of a cutoff parameter: the majority ($\simeq 67\%$) of the attempted calls are a maximum of σ hops long, thus the average path length values for different system sizes will be similar. Figure 10 shows that in such simulations the per node critical load does not fall with increasing system sizes. Actually the per node throughput will constantly hold the same value as the one for the basic non-localized case of n_0 nodes, where n_0 depends solely on σ. We explain this behavior as follows. The traffic initiated by a single node will mostly affect those that are her $1^{\text{st}}, 2^{\text{nd}}, \ldots, \sigma^{\text{th}}$ neighbors. It means that a virtual *subsystem* of the size n_0 is formed around each node: the traffic generated by a call now is limited to a much smaller environment. The traffic otherwise evolves in every such subsystem as in the whole system for the non-localized case, however, the limiting size now is n_0, not n. For $\sigma = \text{const}$ n_0 will not change with n, thus the throughput is constant for various system sizes.

Figure 10 demonstrates a numerical example for comparison of the localized and non-localized cases. We have taken $\sigma = 3.0$, which yields that up to the 3^{rd} neighbors of every node are in the subsystem. The expectation value for the count of the maximum 3^{rd} neighbors of a node is in the range $40 - 60$, which result agrees to the results of Figure 10, where at $n_0 \simeq 50$ the per node throughput values for the two traffic generation cases are nearly equal.

V. Conclusion

In our paper we devise an alternative model of the maximum throughput for stationary ad hoc networks. Our model is consistent with the $O(1/\sqrt{n})$ decay of the throughput, and provides an intuitive way to acquire first order approximation concerning the effect of qualitative changes in the topology or traffic model parameters of the network. Examples are given to verify and support the heuristics, and in order to show the applicability of our model.

References

[1] S. Das, C. Perkins, and E. Royer. Performance comparison of two on-demand routing protocols for ad hoc networks. In *Proc. IEEE Infocom*, 2000.

[2] J. Gomez et al. A power-aware routing optimization scheme for mobile ad hoc networks. Work in Progress, `draft-gomez-paro-manet-00`, 2001.

[3] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 1999.

[4] E. Kail, G. Németh, and Z. R. Turányi. The effect of the transmission range on the capacity of ideal ad hoc networks. In *Proc. WPMC'01*, 2001.

[5] P. R. Kumar et al. Protocols for media access control and power control in wireless networks. Work in Progress, 2001.

[6] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris. Capacity of ad hoc wireless networks. In *Proc. ACM Mobicom*, 2001.