Increased Connectivity at Lower Cost: The Case for Multi-radio Nodes in Multi-hop Wireless Networks

Stefano Basagni
ECE Department
Northeastern University
Boston, MA 02115
basagni@ece.neu.edu

András Faragó
Dept. of Computer Science
Univ. of Texas at Dallas
Richardson, TX 75080
farago@udallas.edu

Maurizio A. Nanni
ECE Department
Northeastern University
Boston, MA 02115
mnanni@ece.neu.edu

Dung T. Tran
Dept. of Computer Science
Univ. of Texas at Dallas
Richardson, TX 75080
dtt063000@utdallas.edu

***Abstract*—We address multi-radio networks, i.e., wireless networks where the nodes are equipped with multiple air interfaces. We analyze, both analytically and via simulation, various gains that the multi-radio environment can provide. First we investigate the gain in network connectivity by modeling the topology of a multi-radio network by a multigraph. The gain is captured by introducing the novel graph theoretic concept of the multigraph advantage. It is the surplus of connectivity over the sum of the individual connectivities, as we put together several graphs to form a multigraph sum. We first prove that in the traditional random graph model it results in a strict super-additive behavior. We validate the theoretical results via simulations and show that similar phenomena occur in geometric random graph models. We then investigate, via ns2-based simulations, the nodal energy consumption as well as the end-to-end packet latency needed to route packets in a multi-radio network.**

I. Introduction

Wireless networking solutions, while becoming more and more ubiquitous, also have growing diversity. Existing and emerging systems, such as variants of IEEE 802.11 wireless LANs (802.11a, 802.11b, 802.11g), IEEE 802.15 personal area networks, ZigBee, Bluetooth, broadband wireless (IEEE 802.16), as well as a wide variety of sensor networking solutions and many proprietary systems are likely to co-exist in the foreseeable future. In parallel with the growing diversity, the radio interfaces are rapidly getting fairly inexpensive and physically small. As a result, it has become a reasonable option to equip wireless network nodes with several radio transmitters/receivers, each of which may follow a different mode of operation in a different frequency band. This creates an environment where the network effectively has *multiple physical layers*. For example, many laptops or PDAs have both an IEEE 802.11 card and a Bluetooth interface. Given the tendency of decreasing price and shrinking physical size, it is quite likely that such multi-radio nodes will soon become very common.

Many of the ubiquitous wireless network nodes will likely be capable of operating effectively with multiple physical layers. The current deployment of multiple radio interfaces on a single node serves very different purposes. For instance, an IEEE 802.11 card is usually used to connect a device to the Internet, or to provide links for multi-hop routing. Bluetooth, on the other hand, is used mainly for cable replacement, i.e., for connecting external devices (e.g., mouse and keyboard) to a device wirelessly, or to synchronize data between, say, a computer and a cellular phone. An important question, however, arises about whether deploying multiple (wireless) interfaces in a cooperative way (i.e., for implementing common networking functions such as routing, broadcast, etc.) will actually lead to significant improvement in the overall network performance. In other words, the critical question is: Will multi-radio based wireless networking solutions provide enough performance gain to justify the investment?

In this paper we contribute to the research on multi-hop, multi-radio networks by providing quantitative evidence of the advantage of deploying nodes with multiple radios. The contribution of this paper is twofold. First we analyze the *gain in network connectivity* by modeling the topology of a multi-radio network by a multigraph. The gain is characterized by introducing the novel graph theoretic concept of the *multigraph advantage*, which is the surplus of connectivity over the sum of the individual connectivities, as we put together several network topology graphs to form a "multigraph sum." As a first step, we formally prove that, in the traditional random graph model, connectivity exhibits a strict super-additive behavior. This strictly super-additive nature of connectivity means that we get a surplus over the sum of the individual connectivity parameters, which points towards increased network reliability and fault tolerance. Our theoretical results are validated via simulations, through which we also show that similar super-additive properties can be observed for various geometric random graph models, more often used to model multi-hop networks.

Increased connectivity, achieved by adding radio interfaces to the same node, might suggest a higher nodal energy consumption. Therefore, we set to investigate the energy expenditure due to routing over multi-hop paths when using multiple radio interfaces. What we observe is the quite counter-intuitive fact that connectivity gain actually comes with *reduced nodal energy consumption* for transmitting and receiving packets. A closer look to the new network structure obtained by multi-radio nodes lead us to observe that a multi-radio topology has an enriched set of possible routes with respect to its

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

single-radio components, allowing routes that would not exist in any of them. This richer set of routes makes it possible to save energy. Through ns2-based simulations performed on networks with increasing traffic we observe that the energy needed to route packets through the multi-hop paths of a network with multi-radio nodes is, on average, up to 30% less than that required to route the same packets on single-radio networks. Furthermore, exploiting the richer sets of routes provided by multi-radio links also yields remarkable improvements for what concern end-to-end packet latency. We observe that packets in a multi-radio network are delivered, on average, up to 70% faster than in single-radio networks.

The remainder of the paper is organized as follows. Section II introduces the multigraph advantage, i.e., a general model for multi-radio network topology analysis. In Section III we investigate through mathematical analysis the multigraph advantage in the classic random graph model. Section IV presents our experiments for validating the multigraph advantage results, and for our investigation of nodal energy consumption and packet latency. Finally, Section V concludes the paper.

II. A Model for Network Topology Analysis

Our fundamental model of the network topology for a network with multiple physical layers is a *multigraph*, possibly with labeled edges. As it is well known, a multigraph differs from an ordinary graph in that multiple edges can connect any pair of vertices [2]. Naturally, the vertices represent network nodes, while the edges represent the links in the various physical layers, e.g., that are generated by the different radio interfaces. The labels on the edges distinguish their affiliation to the various layers. As graphs have long been used to model network topologies, one may rightfully ask the question at this point: Can a multigraph lead to any *essential* new insight? In what follows we show that this model yields interesting and nontrivial novel problems.

The multigraph advantage. Consider a wireless multi-hop network with antennas with beamforming capabilities and let its network topology be represented by the graph G_1 in Figure 1(a). Assume now that we equip the same nodes with another physical layer, such as a second radio with an omni-directional antenna. This second physical layer generates another network topology, shown by graph G_2 in Figure 1(b). If we put the obtained topologies together, we obtain the topology of the combined network, shown by the multigraph G in Figure 1(c).

Let us call the above merging operation of graphs the *multigraph sum* of G_1 and G_2 and let us denote this operation by $\uplus$. Therefore, we obtain G as $G_1 \uplus G_2$. The operation can be extended to more components in a natural way, so we can take the multigraph sum of any number of graphs or multigraphs:

$$\biguplus_{i=1}^{N} G_i = G_1 \uplus G_2 \uplus \ldots \uplus G_N.$$

Note that by assigning technology-dependent labels to the edges, the multigraph sum becomes invertible, that is, we can recover each individual G_i from $\biguplus_{i=1}^{N} G_i$.

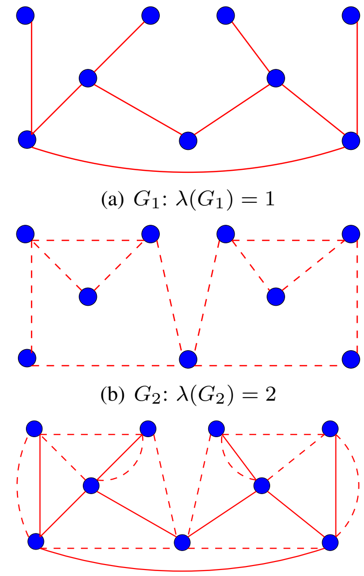

(a) G_1: $\lambda(G_1) = 1$

(b) G_2: $\lambda(G_2) = 2$

(c) $G = G_1 \uplus G_2$: $\lambda(G) = 4$

Fig. 1. The multigraph advantage: Edge connectivity λ

Having introduced the multigraph sum, let us now consider a curious property of it, which we call *multigraph advantage*. We explain it through an important graph parameter, the *edge-connectivity* (or simply *connectivity*, for short), denoted by $\lambda(G)$. This is the minimum number of edges that one needs to remove for disconnecting the graph. In other words, this is the size of a *minimum cut* in the graph.

In case G is a multigraph, the parallel edges are all counted when we consider the size of a cut. The connectivity provides an important characterization of the network topology. For example, it tells how vulnerable is the network to link failures. Connectivity also shows how rich is the network in link-disjoint routes. This is because it guarantees that between any pair of nodes there are at least $\lambda(G)$ disjoint routes. (This fact is a consequence of Menger's theorem in graph theory [2].)

Now one can easily see that

$$\lambda(G_1 \uplus G_2) \geq \lambda(G_1) + \lambda(G_2) \tag{1}$$

always holds. More generally, we have

$$\lambda\left(\biguplus_{i=1}^{N} G_i\right) \geq \sum_{i=1}^{N} \lambda(G_i). \tag{2}$$

The reason is that the size of any given cut in the multigraph sum is just the sum of the sizes of the corresponding cuts in the components. Thus, the connectivity of the multigraph sum cannot be smaller than the sum of the connectivities of the components. In fact, from the above reasoning one might first expect that inequalities (1) and (2) always hold as *equalities*. This is, however, not the case. Let us look at the graphs in Figure 1. One can directly check that $\lambda(G_1) = 1$ (Figure 1(a)) and $\lambda(G_2) = 2$ (Figure 1(b)). On the other hand, as shown in Figure 1(c), $\lambda(G) = 4$, which is strictly greater than the sum.

The above small example shows what we call *multigraph advantage*. Regarding connectivity, it means that the connectivity of the multigraph sum can be *strictly larger* than the sum of the connectivities of the components. Thus, the network topology has a *quantifiable* extra benefit from taking the multigraph sum, representing multiple physical layers.

One may reasonably ask at this point: Is this behavior just the consequence of specially chosen examples or is it somehow typical? In the following sections we show that it is quite typical.

III. The Advantage in a Random Network Model

A. Choosing Random Graphs

The type of random graph that is most frequently used to model mobile wireless networks is the *random geometric graph* [9]. In the simplest case it is generated by randomly placing points in a planar domain, typically a regular one, such as a square. The points represent the randomly positioned network nodes. Any two points are connected with an edge (representing a network link), whenever their distance is at most a given value r, that stands for the transmission radius.

There are various generalizations of the above simple model. For example, the domain can be irregular and it may be part of a higher dimensional space, possibly with a different metric. To our knowledge, the most general version that still has a geometric flavor is based on pre-metric spaces [4] subsuming all previously studied geometric random graph variants.

On the other hand, the oldest and most studied random graph model is the one in which all edges are chosen *independently* with the same probability [3]. This model is referred to as independent-edge random graph, or Erdős-Rényi random graph, or sometimes Bernoulli random graph. While this model lacks the edge correlations induced by geometry, nevertheless it has a number of important advantages, beyond the fact that it is much more amenable to mathematical analysis. For instance, in situations where the transmission radius is comparable with the domain diameter, the main reason for a missing link is not the distance, it is the presence of random obstacles or random variations in the radio propagation. In such a case an independent edge model can be adequate, or even more accurate than a geometric random graph. Furthermore, some studies have shown that if the radio propagation model is more realistic, e.g., it takes into account statistical variations around the mean power, then it tends to decrease link correlations and the actual topology graph becomes similar to the independent-edge model [6], [7]. Finally, if the network applies power control, it tends to counterbalance the effect of distance, especially if the distances are not too large. This again points to the applicability of the independent edge model.

B. Showing Advantage

We consider the following situation. Let $G_{n,p}$ denote a random graph on n nodes, with edge probability p. Note that typically p is a function of n, i.e., $p = p(n)$, and we are interested in asymptotic properties, when $n \to \infty$.

A well known property (see [3]) is that $G_{n,p}$ is connected asymptotically with probability 1 if and only if

$$p(n) = \frac{\log n + \omega(n)}{n},$$

where $\omega(n)$ is any function that tends to infinity with n. Informally, this means that the edge probability of $\frac{\log n}{n}$ is the minimum requirement for connectivity, as $\omega(n)$ can tend to infinity arbitrarily slowly.

We investigate the multigraph advantage in a situation when the network is moderately dense in the sense that the average degree is constant times larger than the minimum needed for connectivity.

Definition 1: The random graph $G_{n,p}$ is said to be in the *moderately dense regime* if there exists a constant $c > 1$, such that the following holds:

$$p(n) = \frac{c \log n}{n}.$$

Now we consider the following situation. Let G_{n,p_1} and G_{n,p_2} be two independently drawn random graphs on the same set of nodes, with edge probabilities p_1 and p_2, respectively. Assume that they are both in the moderately dense regime. Let us compare the connectivities $\lambda(G_{n,p_1})$, $\lambda(G_{n,p_2})$ and the multigraph connectivity $\lambda(G_{n,p_1} \uplus G_{n,p_2})$. We can prove the following result about the multigraph advantage regarding connectivity. (The proof can be found in [5].)

Theorem 1: Let G_{n,p_1} and G_{n,p_2} be independently drawn random graphs in the moderately dense regime, on the same set of nodes. Let their edge probabilities be

$$p_1(n) = \frac{a \log n}{n} \quad \text{and} \quad p_2(n) = \frac{b \log n}{n}$$

with constants $a, b > 1$. Then there exists a constant $c = c(a, b) > 0$, such that the asymptotic multigraph advantage regarding connectivity is at least $c \log n$. That is,

$$\lim_{n\to\infty} \Pr\left(\frac{\lambda(G_{n,p_1} \uplus G_{n,p_2})}{\lambda(G_{n,p_1}) + \lambda(G_{n,p_2}) + c \log n} \geq 1\right) = 1$$

As shown by Theorem 1, the multigraph advantage is quite significant in the moderately dense regime, for two reasons. First, it tends to infinity as the graphs grow. Second, since it is known that in the moderately dense regime the connectivity of the random graph is $O(\log n)$ [3], the gain in connectivity is of the same order of magnitude as the connectivity of the component graphs. In other words, there is a guaranteed constant percentage of *relative gain*, and this percentage does not vanish as $n \to \infty$.

IV. Experimental Results

This section starts by validating our theoretical results on random graphs with a realistic number of nodes. We also show that our result holds for geometric random graphs, which are often used to model ad hoc and sensor network

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

topologies. These results refer to the gain in connectivity (called *multigraph gain*) obtained by combining (two) single-radio topologies into one. In the second part of this section we consider realistic multi-radio scenarios and we show the advantage of deploying multi-radio nodes in terms of nodal energy consumption and packet latency. All the results have been obtained by averaging over a number of experiments large enough to achieve 95% statistical confidence.

A. Multigraph Gain for a Finite Number of Nodes

To measure the connectivity gain obtained with multi-radio nodes we developed a simulator in C++. For computing the (multi)graph connectivity λ we implemented the min-cut algorithm of Stoer and Wagner [10].

1) Validating Theorem 1: Theorem 1 provides a theoretical *asymptotic* lower bound on the multigraph gain in connectivity. It does not say, however, anything about two things: 1) Whether the bound holds for a limited number of nodes, and 2) how tight is the lower bound, especially for a finite number of nodes. We have investigated these questions via simulation.

In our simulation setting we consider topologies with 100, 150 and 200 nodes. The edge probabilities are chosen such that $p_1 = 0.05$ is fixed, while p_2 varies from 0.05 through 0.25. We have measured the ratio

$$\rho = \frac{\lambda(G_{n,p_1}) + \lambda(G_{n,p_2}) + c \log n}{\lambda(G_{n,p_1} \uplus G_{n,p_2})} \qquad (3)$$

which, according to Theorem 1, should tend to 1, but only asymptotically. Figure 2 demonstrates that already for a relatively small number of nodes and for a wide range of edge densities the ratio is quite close to 1. That is, our asymptotic lower bound is a good approximation of the *actual multigraph gain*, even for a relatively small number of nodes.

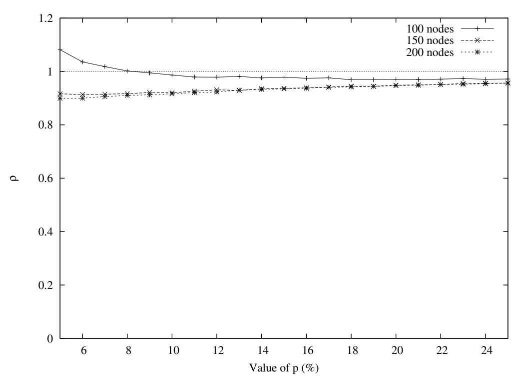

Fig. 2. Theorem 1 in experiment

2) The dependence of multigraph gain on graph densities: The result shown by Figure 2 confirm that the multigraph gain is well predicted by our formula (3). Here we investigate how the relative gain changes for varying densities of the component graphs. We have observed the following interesting behavior that is described below.

We have fixed the density (edge probability) of one graph at 0.05, while the other is varied in the range $p = 0.03 \ldots 0.25$. The number of nodes is 100, 150 and 200. The (relative) multigraph gain is shown in Figure 3(a), as a function of p. The interesting phenomenon is that it has a bell shaped curve, suggesting that there is an *optimal density difference* that maximizes the relative multigraph gain. The curves are essentially the same for different numbers of nodes, i.e., it appears that there is no significant dependence on the number of nodes.

3) Multigraph gain in other graph models: A similar qualitative behavior observed in the random graph model has been found for other graph models. Figure 3(b) shows the results of running the same experiment on random geometric graphs, which are often used to model ad hoc and sensor network topologies. The nodes in our experiment are scattered randomly and uniformly in a square area of 1000×1000 square meters. Two nodes are connected if they are within a given transmission radius r. The value of r for one graph is 180m, while for the other graph r is varied from 190 through 550.

The three curves in Figure 3(b) correspond to three experiments with 100, 150 and 200 nodes. They all show a similar bell shaped curve, suggesting that there is a density difference that maximizes the relative gain also in this case. Interestingly, while the value of the gain depends on the number of nodes, the place of the optimum remains approximately constant.

Finally, we have run the same set of connectivity related experiments on a third set of topology graphs, the so called *nearest-neighbor graphs* [8]. Nodes are generated the same way as in the previous experiments. However, the links are created such that each node is connected with its k nearest neighbors. In our experiments we have set the $k = 3$ for the first graph (sparse topology), while we have varied the density of the second graph (i.e., we have let k vary from 4 to 18). We have again observed the characteristic bell shaped curve, independently of the number of nodes (Figure 3(c)).

B. Experiments in Realistic Multi-hop Settings

We have shown, theoretically and through experiments, that deploying multiple (radio) interfaces on the network nodes of a multi-hop network increases its connectivity, which points toward greater robustness and fault-tolerance. We might object, however, that supporting multiple radio interfaces imposes a greater toll on nodal energy consumption for packet communications, and that it might also have an impact on packet transmission overhead (collisions) and end-to-end packet latency. In this section we show via simulations that this is indeed not the case, as we perform experiments that show that while adding radios positively affects packet latency, protocol overhead and nodal energy consumption for routing packets along multi-hop routes are actually less for a multi-radio than for its single-radio components.

1) Protocols and simulation environment: In order to assess how the multigraph advantage translate into actual multi-radio performance gain we implement a simple shortest path routing protocol by using the VINT Project *ns2* simulator and its extension ns2-MIRACLE [1]. The considered scenarios comprise 100, 200 and 300 nodes randomly and uniformly

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

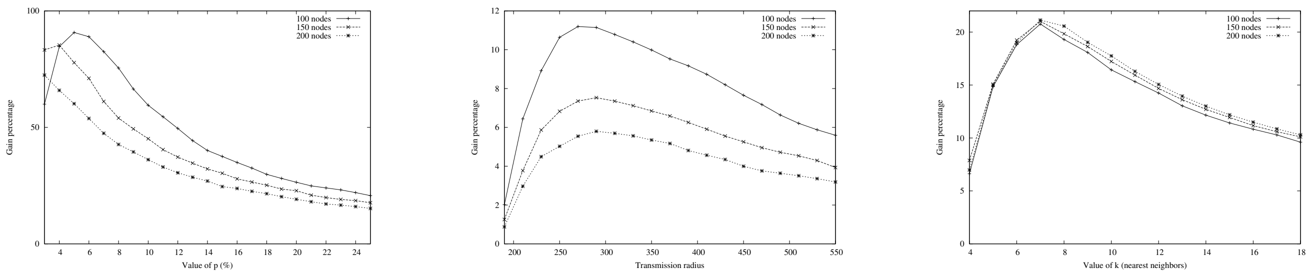

(a) In random graphs (b) In random geometric graphs (c) In k nearest-neighbor graphs

Fig. 3. Multigraph gain

deployed on a 1000m $\times$ 1000m area. Each node is equipped with two 802.11g radios, with a transmission rate of 6Mb/s. The first radio operates on channel 1, while the second on channel 11. (There is no co-channel interference between the two interfaces.) Each radio keeps transmitting on the same channel for the whole duration of the simulation. In other words, there is no need for channel switching, which implies that no scheduling algorithm is required for ensuring that transmitters and receivers are tuned on the same channel. Furthermore, since the two interfaces act on independent channels concurrent communications on the two interfaces may occur. The transmission radius of the two radios are set to $r_S = 130$m and $r_L = 180$m, respectively (the radio propagation model is free-space). The short-range radio (radius r_S) has a transmission energy consumption of 0.28W . The long-range radio (radius r_L) has a transmission energy consumption of 0.53W. For both interfaces, the receiving energy consumption has been set to 0.02W. (In order to assess the multi-radio advantage independently of any specific possible awake/asleep schedule used for energy conservation we did not consider the energy consumed by a node while being idle. This choice corresponds to equipping each radio with some form of very low cost wake-up radio technology.) On top of the link layer of each of the two interfaces, we implemented a (single) routing module. Upon receiving a packet, the routing module can choose to send it down to either radio. Our routing module implements a simple shortest path protocol where shortest paths are computed once and for all for each pair of source-destination nodes based on edge weights that represent energy consumption due to communication. In other words, we run a multi-graph all-pair shortest-path algorithm in which each edge is weighted according to the transmission and reception energy consumption of the two radios. Thus, every packet will be routed along the cheapest path, in terms of energy consumption. Nodal routing tables are set at network set-up. Each simulation run was set to last 3600 seconds. During the first 600 seconds metrics are not collected (transitory period). In the following we will refer to the remaining 3000 seconds as the simulation time. Traffic is generated according to a Poisson process of parameter λ, where λ indicates the average number of arrivals (i.e., of packets generated) per second. In our experiments we set λ to five different values, namely, $\lambda = 50, 100, 150, 200, 250$. These values impose a network traffic from low ($\lambda = 50$) to relatively high ($\lambda = 250$). For each packet generated a source node and a destination node are randomly and uniformly chosen among all nodes. Each generated packet is 500B long.

2) Investigated metrics: The metrics selected for our study concern: 1) The number of MAC layer collisions experienced by a node during the simulation time. This metric helps in understanding whether transmitting packets on links of different radios decreases the MAC overhead. This is also a metric the shed light on results for packet latency and nodal energy consumption. 2) The nodal energy consumption (in Joules), i.e., the energy drained by a node for packet transmission and reception throughout the simulation time. 3) The end-to-end packet latency (defined as the time from when a packet is generated to when it is buffered at its intended destination) for packets generated throughout the simulation time.

Every figure shows three curves: One for the multi-radio topology ($r_S = 130$m and $r_L = 180$m) and the other two for the single-radio component topologies generated by r_S and r_L.

3) MAC layer collisions: Figure 4(a) depicts the average number of collisions experienced by a node throughout the simulation time. We can observe that in single-radio networks with larger radius the number of collisions is higher than in all the other cases. However, we also observe that in the short range scenario the number of collisions is only at most 6% lower than those counted for networks with larger radius (this happens at high traffic, i.e., for $\lambda = 250$). This is because, despite it is clearly true that a smaller transmission radius reduces the number of collisions among nodes whose transmissions overlap, the topologies generated by smaller radii are sparser, and this increases the number of bottleneck nodes. These nodes have to deal with a lot of traffic, and hence they are points in the network where an elevated number of collisions occur. Finally, smaller radii correspond to longer routes, which in turn, when the traffic is high, keeps the number of collisions relatively high.

The multi-radio scenario performs remarkably better than the other two. In particular, we observe an average of 58% less collisions than those counted for r_S networks, and an average of just above 60% less collisions than those of r_L networks (high traffic scenarios, i.e., with $\lambda = 250$). This is

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

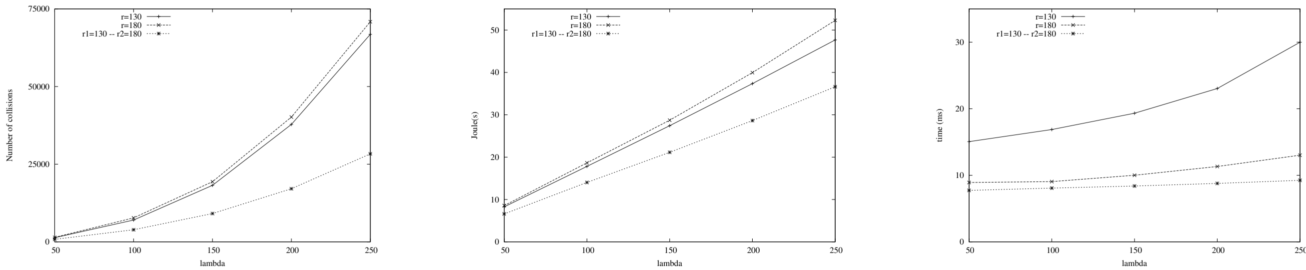

(a) Number of collisions (b) Energy consumption (c) Packet latency

Fig. 4. Average number of nodal MAC collisions, nodal energy consumption and end-to-end packet latency

because the multi-radio setting enjoys the best of two worlds: Use longer hops (i.e., avoid bottlenecks) and send packets on the shorter ones (reducing the collision probability).

4) Nodal energy consumption: Figure 4(b) depicts the average energy consumption needed by a node during the simulation time for transmitting and receiving packets. Multi-radio nodes save considerable energy with respect to their single-radio counterparts. In particular for heavy traffic load ($\lambda = 250$) a multi-radio node spend 23% less energy than that consumed by a node in r_S single-radio networks. These savings increase to 30% with respect to networks with radius r_L. Energy consumptions in the two single-radio scenarios are fairly similar: Even if a shorter transmission radius requires a packet to travel through longer routes (i.e., through more nodes), and therefore we have a higher number of transmissions, with a lower transmission radius we have a lower number of collisions. We have observed already that this number of collisions is not overwhelming. Therefore, eventually, energy consumption balances out among the two single-radio scenarios, especially at low traffic. When the traffic increases, as expected, lower radii, i.e., lower transmission powers, allow nodes to save more energy. When $\lambda = 250$ savings in networks with short range radio (r_S) average around 9% with respect to the case of long range networks (radius r_L).

5) End-to-end packet latency: Combining two topologies into a multi-radio one, we are able to deliver packets to their destination faster, to the extent depicted in Figure 4(c). For the single-radio scenario with radius r_S packet delays are, as expected, much higher than those obtained in the single-radio scenario with radius r_L. Especially at high traffic ($\lambda = 250$) in networks with radius r_S we measured average packet delays up to 2.3 times bigger than those measured in networks with larger radius. This is because the higher average route lengths corresponding to shorter hops. We also observed that for the r_S single radio scenario the delays grow remarkably faster than those in the r_L single-radio scenario. This happens because the network in the first case is getting congested with increasing values of the parameter λ. The congestion is due to the fact that with a smaller transmission radius there are many bottleneck nodes in the network that have to face a large amount of traffic (the topology of r_S-generated networks is quite sparse). In topologies generated by r_L the radius is big enough to avoid bottleneck nodes. This justifies the "flatter" behavior of the corresponding curve in Figure 4(c).

In the multi-radio scenario end-to-end delays are even lower. In particular, for $\lambda = 250$, we observe delays that are about 70% lower than those obtained in the single-radio scenario with radius r_S, and about 29% lower than those in networks with radius r_L. This is due to the fact that, as observed, collisions in the multi-radio setting are heavily reduced, so that packets can proceed more rapidly to their destinations.

V. Conclusions

This paper explored the advantage of deploying nodes with multiple radio interfaces in multi-hop wireless networks. In particular, we have shown, both theoretically and through simulation-based experiments, that multi-radio network connectivity increases following a super-additive law with respect to the single-radio case. We have also shown that the advantage of multi-radio nodes does not stop at increased connectivity, but extends to reduced energy consumption for routing packets from sources to destinations, and to delivering packets faster than single-radio nodes.

VI. Acknowledgments

This research was supported in part by NSF Grants CCF 0634847 and 0634848.

References

[1] N. Baldo, F. Maguolo, M. Miozzo, M. Rossi, and M. Zorzi. ns2-MIRACLE: A modular framework for multi-technology and cross-layer support in network simulator 2. In *Proceedings of ValueTools 2007*.

[2] C. Berge. *Graphs and Hypergraphs*. North–Holland, 1985.

[3] B. Bollobás. *Random Graphs*. Cambridge University Press, Cambridge, UK, second edition, 2001.

[4] A. Faragó. On the fundamental limits of topology control in ad hoc networks. *Algorithmica*, 49(4):337–356, December 2007.

[5] A. Farago, D. T. Tran, and S. Basagni. Improving fault tolerance via multi-radio nodes: A graph theoretic analysis. Technical Report UTDCS-27-08, Dept. of Computer Science, Univ. of Texas at Dallas, Sept. 2008.

[6] R. Gowaikar, B. Hochwald, and B. Hassibi. Communication over a wireless network with random connections. *IEEE Transactions on Information Theory*, 52(7):2857–2871, July 2006.

[7] R. Hekmat and P. Van Mieghem. Connectivity in wireless ad-hoc networks with a log-normal radio model. *ACM/Kluwer Mobile Networks and Applications (MONET)*, 11(3):351–360, June 2006.

[8] M. S. Paterson and P. P. Yao. On nearest-neighbor graphs. In *Proceedings of ICALP 1992*, LNCS 623, pages 416–426, Vienna, Austria, July 13–17 1992.

[9] M. D. Penrose. *Random Geometric Graphs*. Oxford Studies in Probability, 5. Oxford University Press, New York, NY, June 1 2003.

[10] M. Stoer and F. Wagner. A simple min-cut algorithm. *Journal of the ACM*, 44(4):585–591, July 1997.

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.