TECHNICAL RESEARCH REPORT

Energy-Efficient Routing for Connection-Oriented Traffic in Wireless Ad-hoc Networks

by Anastassios Michail and Anthony Ephremides

CSHCN TR 2001-8
(ISR TR 2001-15)

Energy-efficient Routing for Connection-Oriented Traffic in Wireless Ad-Hoc Networks [*]

Anastassios Michail and Anthony Ephremides
Department of Electrical & Computer Engineering
and Institute for Systems Research
University of Maryland
College Park, MD 20742
tassos@ieee.org, tony@eng.umd.edu

Abstract

We address the problem of routing connection-oriented traffic in wireless ad-hoc networks with energy efficiency. We outline the trade-offs that arise by the flexibility of wireless nodes to transmit at different power levels and define a framework for formulating the problem of session routing from the perspective of energy expenditure. A set of heuristics are developed for determining end-to-end unicast paths with sufficient bandwidth and transceiver resources, in which nodes use local information in order to select their transmission power and bandwidth allocation. We propose a set of metrics that associate each link transmission with a cost and consider both the cases of plentiful and limited bandwidth resources, the latter jointly with a set of channel allocation algorithms. Performance is captured by call blocking probability and average consumed energy. A detailed simulation model has been developed and used to evaluate the algorithms for a variety of networks.

1 Introduction

Energy efficiency is important in the design of battery-operated wireless devices that are used in wireless networks. While users' demand for improved and more sophisticated functionalities of wireless devices increases rapidly, improvements in battery technology come at a slower pace. Therefore the possibility to provide network layer solutions (such as routing algorithms) in a way that takes into consideration energy expenditures presents a novel opportunity.

In this paper we study the problem of energy-efficient routing of connection-oriented traffic in wireless ad-hoc networks. A wireless ad-hoc network is an autonomous system of wireless routers, fixed or mobile, that communicate with each other via radio links without the support of any fixed wired infrastructure. In such a topology, the set of links and their capacities are not determined a priori, but depend on multiple factors, such as distance between nodes, signal power, interference, noise and propagation characteristics. Due to the limited transmission range of wireless radios, transmitted information cannot always be received by all nodes and communication is achieved by use of multiple hop paths between the sender and the receiver. Wireless ad-hoc networks constitute a typical paradigm of networks whose performance and operability depends crucially on battery power. The fact that most nodes are likely to play the role of a relay node, having to draw on their energy resources even when they do not need to engage in communication themselves, illustrates the importance of energy efficiency.

A crucial choice in wireless transmission is that of the RF power level. Due to the nonlinear attenuation of the received signal power with distance, a transmission over multiple short hops may require less total power than a single transmission over one long hop. On the other hand, multiple short transmissions could result in significant overhead and routing complexity along with utilization of a larger amount of network resources, thereby potentially increasing the overall energy consumption. Note also that nodes consume energy not only during transmission, but also when they receive, store and process information. The use of sophisticated algorithms that deal with congestion, or of more efficient coding schemes that perform better in bandwidth constrained links, results in needs for additional signal processing by the wireless transceivers and hence in demand for more energy. Nodes that have to dedicate

part of their resources (bandwidth, transceivers and energy) for relaying messages may be overused for routing functionalities, while others will remain idle for longer intervals, due to the topology characteristics. Such an "unfair" utilization causes certain nodes to exhaust their energy reserves faster and be forced to turn their power "OFF" which leads to performance degradation or even network partitioning.

Another crucial issue associated with the choice of the transmission power level is the interference caused to non-intended recipient nodes located in the vicinity of the transmitter. Transmitting at higher power reduces the efficiency of bandwidth re-use and causes increased interference for a fixed allocation of bandwidth resources. On the other hand, if a path consisting of multiple short hops is used, the total power required for transmission may be lower, but there is need for efficient scheduling mechanisms to avoid conflicts among the consecutive links of a path.

In this paper our focus is on source initiated unicast (single-source, single-destination) connection-oriented traffic and on developing routing algorithms that are capable of identifying paths connecting the source to the destination that provide the required resources from end to end, and subsequently keep them reserved until the completion of the session. Such resources are represented by node transceivers, energy and bandwidth availability (frequency channels, time slots or CDMA orthogonal codes, depending on the access scheme). At the same time, and in order to support connection-oriented traffic, nodes must be capable of receiving and transmitting simultaneously, necessitating therefore the use of sufficiently separated frequency channels. We choose to use frequency division multiple access (FDMA), which introduces the difficult problem of frequency scheduling, i.e. assigning non-interfering frequencies to transmitting nodes. CDMA schemes do not allow nodes to handle simultaneous transmission and reception in the same frequency band and are not considered here. Time division multiple access (TDMA) is of course another option but is conceptually similar to FDMA while at the same time introduces an even more difficult slot scheduling problem.

Related work on multihop networks that support connection-oriented traffic is for multicast routing. In [1], the authors study the effects of wireless network characteristics and of energy constraints on multicast protocol operation and propose an algorithm that exploits the node-based nature of wireless

3

communications for multicasting. In [2], a set of algorithms is proposed for the construction of minimum energy broadcast and multicast trees, which is extended in [3] to capture the effects of limited bandwidth resources. In [4], multicast routing algorithms that use capacity results for multiuser detectors are developed. A variety of approaches for energy efficiency in packet-switched networks have been presented in [5], [6] and [7]. In [5], an algorithm is proposed that determines a graph containing minimum power paths in a random topology. In [6], the authors propose a suite of algorithms that based on network flow theory try to balance the minimum lifetime of each flow path, by redirecting or augmenting the flow of certain paths and by identifying traffic splits that optimize energy consumption. However, these principles cannot be applied in the case of sessions where a path must be reserved end-to-end for the whole duration of a session. A different approach is taken in [7] where a model is presented that overcomes the complication that arises with the interference caused by increasing the traffic on a link. This model allows extension of optimal routing methodology for wire-based networks to do minimum-energy-and-delay routing for packet radio networks.

We propose a set of algorithms that jointly address the issues of transmission power levels (a physical layer function), route discovery (a network layer function) and bandwidth reservation (a MAC layer function). In particular, each node determines its transmission power, transmission frequency and next-hop neighbor, based on local information of network parameters (ie., transmission power, energy reserves, availability of transceivers and frequency channels), with the objective of identifying unicast routes that optimize performance as captured by the overall blocking probability and the average energy expenditures. Our approach is characterized by three innovative features. First, we address the unicast problem which is not characterized by the combinatorial complexity of multicasting; in fact under simplifying assumptions regarding interference and node resources minimum-energy solutions <u>can</u> be found. However, the reduced amount of complexity allows us to extend our approach to study also the effects of local interference and limited node resources (e.g. transceivers) without the additional requirements that multicasting would impose. Moreover, even though some objectives may be parallel, the actual algorithms, metrics and trade-offs are quite different as we will see in the sequel. Secondly,

we convert session routing to link metric based, even though algorithms based on minimum-distance paths are normally applicable to packet-switched networks (where the cost of using a link is typically the estimated packet delay). In particular, in telephony networks it is hard to define such metrics since energy is not a concern and delay is not an appropriate metric. Unlike telephone networks, we are able to map the overall objectives (blocking probability and energy consumption) to individual link metrics. Finally we evaluate the effects of receive and processing power in addition to transmission power. Even though processing power typically depends on a set of network parameters, we consider constant energy depletion rate per node (for receiving and signal processing) and observe its effects on the performance of our algorithms.

The effects of mobility are not addressed in this paper; in order to assess the already complex trade-offs one at a time, we focus on developing algorithms for wireless fixed topologies. Such an approach is also preferred because as we have seen in our prior work ([8, 9]), mobility effects can be addressed through the use of soft-failure mechanisms. Moreover, the possibility to use the transmission power (or the residual energy) as a metric to select a path adds a new degree of flexibility. Thereby, in the case of a link failure we may either adjust the power to maintain connectivity or choose to reroute along an alternative path, depending on the circumstances. A similar approach has been presented in [8, 9] and has been shown to yield satisfactory results in the case of relatively low mobility. Nonetheless, there are wireless ad-hoc networks (such as sensor networks) that are inherently static and involve no mobility.

Finally we adopt a simplified propagation model that ignores the fading characteristics of the wireless channel. Fading needs to be considered in conjuction with mobility and is deferred for future research.

The rest of the paper is organized as follows. In the next section we describe our wireless network model along with the interference assumptions. We continue with an overview of our approach and a discussion on the difficulties in obtaining exact optimal solutions followed by a detailed description of the routing algorithms first for the case of limited transceivers and then for limited bandwidth resources. The simulation model is presented next and the performance evaluation and comparison of our algorithms set the base for an in-depth analysis of their main characteristics and properties. We

conclude with a summary of the most important remarks.

2   Wireless network model

We consider a network consisting of N randomly deployed nodes. Connectivity of the network depends on the distance between nodes, the maximum transmission power level and the minimum required received power at a node. We assume that all nodes may transmit at any power level P which may not exceed a pre-determined maximum value P_{max}. Received signal power varies as $d^{-\alpha}$, where d is the Euclidean distance between transmitting and receiving node and $\alpha \geq 2$ is the path-loss exponent. [1] Note also that our algorithms will be independent of the value of α, so that they are applicable in various propagation environments. We consider constant α throughout the region of interest, there are no obstacles and the antennas are omnidirectional so that all nodes within communication range of the transmitting node can successfully receive the transmission.

Given P_{max}, the distances between nodes, and the minimum required received power for error-free communication, we can determine the communication range of all nodes and the connectivity map of the network. For notation purposes we define the set $\mathcal{R}^{(i)}$ of node i to be the set of nodes within transmission range of i. All nodes $j \in \mathcal{R}^{(i)}$ are reachable by i and the power required to support a link between i and j is denoted by $P_{ij} \leq P_{max}$. Complete knowledge of the set of neighbors located within transmission range indicates the potential recipients of a transmission but is not sufficient for determining whether a connection can be established, since the required resources must also be available. These resources are modeled as follows:

(a) Transceivers: node i has C_i communication transceivers and can therefore support up to C_i sessions simultaneously. The number of "reserved" or "occupied" transceivers (B_i) varies with time according to the network state. We denote by R_i the residual capacity (ie, number of not-in-use transceivers) of node i; thus $R_i = C_i - B_i$.

---

[1]Assume here that the path loss depends only on the distance between transmitter and receiver ignoring for simplicity any possible antenna height difference which would make the dependence three-dimensional.

(b) Frequency channels: a total of m frequency channels $(f_1, f_2, \cdots, f_m)$ are available for use. Each node i maintains a separate *channel status* vector for every $j \in \mathcal{R}(i)$. The channel-status vector of node i for transmission to node j is given by $f^{(i,j)} = [f^{(i,j)}(1), \cdots, f^{(i,j)}(m)]$, where for all k we define:

$$
f^{(i,j)}(k) = \begin{cases} 1 & \text{if } k^{th} \text{ channel is available} \\ 0 & \text{otherwise} \end{cases} \tag{1}
$$

(c) Energy: at time t, node i has a residual amount of energy $E_i^R(t)$ and the initial amount of energy of node i is denoted by E_i^o. We assume that all nodes keep track of their residual energy at all times and only nodes with nonzero residual energy can participate in the network. [2]

Sessions are source-initiated and all nodes may generate connection requests. In order to admit a new session, a path p must exist from the source to the destination that meets the following requirements:

(i) all nodes $i \in p$ must have one transceiver available for use at the time of the request, to be reserved throughout the duration of the session, ie $\forall i \in p, R_i \neq 0$.

(ii) all nodes $i \in p$ must have at least one frequency channel available for transmission; moreover, a conflict-free channel allocation must exist that satisfies the following conditions:

   − a node cannot transmit and receive in the same frequency

   − a node cannot simultaneously receive more than one signals in the same frequency

   − a node cannot transmit simultaneously to more than one neighboring nodes since unicast is strictly considered here; in a broadcast senario this would an acceptable and in fact desirable situation.

Each node maintains up-to-date information about the identities of its one-hop neighbors, the corresponding required transmission power levels, as well as their residual capacity, energy and the status of their frequency channels. All nodes may periodically broadcast updates of the above information to the nodes that are located within transmission range, so that they are used by the routing protocol. This

---

[2]modern battery-monitoring technology permits accurate knowledge of battery energy reserves

can be done through a narrow-band dedicated control channel that is not part of the communication resources for payload transmission. In addition, the energy required for these message exchanges is considered negligible.

2.1 Interference model

In order to determine conflict-free channel assignments, nodes must be aware of the frequencies they are allowed to use at any time (along with the maximum allowed power levels). The restrictions presented in (ii) above can be expressed via a simplified interference model in which we assume that a receiver ignores interference from simultaneous neighboring transmissions if the distance from the location of the interfering source exceeds d_{max}. In particular, we make a binary decision of whether we can allow or not a transmission, without a detailed calculation of SINR at every node. This is somewhat of a simplification that is consistent however, with commonly made assumptions on connectivity in ad-hoc networks. Nonetheless, a more accurate model for interference can be incorporated in our model. We do not consider it in this paper so as not to prematurely complicate an already difficult problem.

The use of frequency f_k for transmission over a link (A, B), results in blocking of the following neighboring links:

Primary conflicts:

P1: Any link (v, A), $v \in \mathcal{R}(A)$, because A cannot receive and transmit at the same frequency.

P2: Any link (A, v), $v \in \mathcal{R}(A)$, because A cannot transmit to more than one nodes simultaneously.

P3: Any link (v, B), $v \in \mathcal{R}(B)$, because B cannot receive from more than one nodes at a time.

P4: Any link (B, v), $v \in \mathcal{R}(B)$, because B cannot transmit and receive at the same frequency.

Secondary conflicts:

S1: Any link (v, u) with $u \in \mathcal{R}(A)$ and $v \in \mathcal{R}(u)$, if and only if $d(A, u) \geq d(A, B)$. Transmission over (v, u) is allowed only in the case that transmission over (A, B) is not received by u.

S2: Any link (u, v) with $u \in \mathcal{R}(B)$ and $v \in \mathcal{R}(u)$, if and only if $d(u, v) \geq d(u, B)$. Transmission over (u, v) is allowed only in the case that it is not received by B.

Consider for example the senario of figure 1. If we assume that an ongoing session from node A to node B is using link (A, B) at frequency f_k, the status of the affected neighboring links during this transmission is shown in table 1.

3 The energy-efficient routing problem

Our objective is to develop algorithms that determine an appropriate unicast path for each newly arriving session request. If we were to assume that ample bandwidth and unlimited number of transceivers were available so that all sessions could be admitted, the problem would reduce to that of determining the minimum energy paths. Let us initially address only the transmission energy (we evaluate the effect of receive and processing energy later in our simulations), so that the total energy of a path is the aggregate energy consumed by its nodes. Since we are considering session traffic, all transmitting nodes along a particular path transmit throughout the duration of a session and the total transmission energy is proportional to the aggregate power needed to maintain the path. Due to the limited number of transceivers and/or frequencies, the minimum aggregate power path cannot always be used and sessions may sometimes be admitted along suboptimal paths (in terms of energy requirements) in order to minimize the number of blocked calls, or not be admitted at all if no paths exist with the required resources. As a consequence, the global performance measures that incorporate the characteristics of our problem are the *blocking probability* P_b and the *average energy per session* E_s. We also consider a third performance metric that captures simultaneously the blocking probability as well as the expended energy. In particular, we introduce what we call global performance yardstick Y as the average call acceptance ratio per energy unit, i.e.

$$Y = \frac{1 - P_b}{E_s} \tag{2}$$

Of crucial importance in developing an algorithm that attempts to minimize P_b under constraints in E_s (or equivalently maximize Y) are the assumptions we will make on the available information about the network state. An ideal situation occurs when "complete" information on the network state is available. Such information includes the full connectivity map with detailed transmission power levels,

the frequency channels being used or blocked by every node, the number of available transceivers and the residual energy at each node. At the establishment or termination of a call this information must be updated and be made available either to all nodes or to some central coordinator. Even if we neglect the cost of collecting and maintaining complete information, an exact optimal solution cannot be obtained unless we know a priori the full traffic pattern. In practice, we can only develop greedy algorithms that work on a per call basis and attempt to maximize the expected reward associated with each incoming call. Even in that case, in order to obtain the optimal solution an exhaustive search of the complete solution space is required which becomes impractical as the network size grows.

In this paper, we concentrate on developing underline{distributed} heuristic algorithms that rely only on local information. Each link (i, j) is associated with a distance metric that indicates the cost of using that link and may incorporate energy and capacity information as well as interference effects. If the cost of using link (i, j) is represented by $D_{i,j}$, the cost of using a path p will be given by $\mathcal{C}_p = \sum_{(i,j) \in p} D_{ij}$. Given the selection of the link metric, we apply distributed Bellman-Ford [10] algorithms to determine the minimum cost path. Link metrics are defined in such a way, that only links connecting nodes with available transceivers and frequency channels have finite cost. However, if frequency channels are limited, these schemes provide no guarantee that a feasible frequency allocation will exist along the minimum cost path. Such an allocation cannot be directly related to the number of available frequency channels but, rather, to the identity of those channels. Therefore, as we will show in section 5, algorithms that can determine appropriate frequency allocations are also needed and must be developed.

4 Unlimited bandwidth resources; limited transceivers case

A sufficiently large number of frequency channels is considered, so that there always exists a conflict-free channel allocation scheme; however, the number of transceivers per node is finite and equal to C. A new call request may be rejected only when no path exists between origin and destination with at least one transceiver per node. Since the minimum power path may not always be available, we can search for the lowest total power path in the subgraph defined by the nodes with nonzero residual capacity

and energy and their corresponding links. We propose the following three link metrics:

(a) Metric M1

Metric M1 is a direct measure of the power needed to transmit over a link, provided both nodes have nonzero residual capacity and energy. The cost of using link (i, j) is defined by,

$$D_{ij}^{(1)} = \begin{cases} P_{ij} & \text{if } R_i, R_j \neq 0 \text{ and } E_i^R, E_j^R \neq 0 \\ \infty & \text{otherwise} \end{cases} \tag{3}$$

where P_{ij} is the power needed for successful transmission from i to j. Metric M1 will always provide the minimum power path available, which may not be advantageous in terms of overall network performance. A minimum power path will usually consist of multiple short hops and therefore will occupy more network resources, which could result in blocking of more new calls. It is also possible, depending on the traffic patterns, that some paths get heavily utilized and act as bottlenecks (in a static topology the minimum power path will be the same until any node is blocked), while others consist of lightly used nodes. Finally, if processing power is not negligible compared to transmitter power, multi-hop paths could sometimes result in larger energy expenditures.

(b) Metric M2

To address the problem of congested nodes, we define metric M2:

$$D_{ij}^{(2)} = \begin{cases} \frac{P_{ij}}{\min\{R_i, R_j\}} & \text{if } R_i, R_j \neq 0 \text{ and } E_i^R, E_j^R \neq 0 \\ \infty & \text{otherwise} \end{cases} \tag{4}$$

Metric M2 favors links that are not heavily utilized and tries to spread the traffic load evenly over all paths by increasing the cost of links that connect nodes with smaller residual capacity.

(c) Metric M3

Both metrics M1 and M2 rely on the power level required to successfully transmit to a one-hop neighbor but ignore the important parameter of the residual energy at each node. Under certain circumstances, it may be preferable to route a call over a path that consists of nodes with larger amounts of residual energy, even though this may result in additional energy consumption by the session. This is so because conservative usage of the remaining energy at a low-energy node that occupies a crucial

position may prolong the connectivity and, hence, the lifetime of the network. Thus, we propose the use of link metric M3 defined by

$$D_{ij}^{(3)} = \begin{cases} W_p \frac{P_{ij}}{P_{max}} + W_e \frac{E_j^o}{E_j^R} & \text{if } R_i, R_j \neq 0 \text{ and } E_i^R, E_j^R \neq 0 \\ \infty & \text{otherwise} \end{cases} \quad (5)$$

where W_p and W_e are weights that may be adjusted to favor either of the two terms. Note that in the beginning of network operation the second term is equal to one for all nodes and therefore our metric is similar to M1. As the residual energy of every node begins to drop, the second term will increase and when the amount of residual energy is low it will dominate the value of the metric. Even though the intuition behind the selection of metric M3 is to "favor" paths consisting of nodes with higher residual energy levels, it can be easily proved that if the algorithm has to choose between a single-hop and a multi-hop path, then M3 (with $W_p = W_e$) will always select the former. Despite this limitation, M3 is very appropriate in situations when the choice is between multi-hop paths; in that case the residual energy of the intermediate nodes becomes crucial. [3]

5 Limited bandwidth resources; unlimited transceivers case

In this section we shift our focus to the case of limited frequency channels and assume a sufficiently large number of transceivers per node. Unlike the case of limited transceivers, the existence of at least one frequency channel in each link of the path is not sufficient to guarantee admission of a session. Instead, an interference-free allocation of channels must be determined. Such an allocation is not directly related to the number of channels available for each node but rather to the identity of these channels. Another limitation is that the nodes of a path cannot select which frequency to use independently; for every channel assignment made over one link, neighboring nodes that experience interference must update their blocked-frequency table before they can make their assignment. Therefore a channel allocation algorithm must generate the frequency assignments on a hop-by-hop basis. To address these issues, we develop algorithms that evolve in two stages: (a) a minimum cost path (as measured by energy

---

[3]The proof is rather straightforward and it has been omitted due to space limitations. It can be found in [11]

and blocked resources) is first determined (the *candidate path*) and (b) if an interference-free channel assignment can be determined along that path, the call request is admitted; else the call is rejected. We discuss in the following sections our proposed methods for achieving both objectives (a) and (b).

5.1 Link metrics for determining minimum cost path

In this section, we propose the *minimum power metric* (MPM) and the *power and interference based metric* (PIM) for determining the cost of the links.

(a) Minimum power metric (MPM)

MPM is a direct measure of the power needed to transmit over the specific link, provided that at least one frequency channel is free for transmission. The cost of using link (i,j) as expressed through MPM is defined by:

$$D_{ij}^{(MPM)} = \begin{cases} P_{ij} & \text{if } \sum_{k=1}^{m} f^{(i,j)}(k) > 0 \\ \infty & \text{otherwise} \end{cases} \tag{6}$$

Clearly MPM is similar to metric M1 (see equation 3) in that it accounts for transmission energy requirements and is expected to produce the minimum transmission power path available.

(b) Power and interference based metric (PIM)

To address the blocking effects a transmission may cause to neighboring nodes we define PIM by incorporating into MPM interference effects. We first introduce the following notation:

- $B^{(i,j)}(k)$ denotes the set of "links" (i.e. transmitter - receiver pairs) that are blocked whenever node i transmits to j over frequency f_k

- $|B^{(i,j)}(k)|$ is the cardinality of $B^{(i,j)}(k)$

- $|E|$ is the cardinality of all transmitter-receiver pairs.

PIM is then defined as follows:

$$D_{ij}^{(PIM)} = \begin{cases} \frac{P_{ij}}{P_{max}} + \frac{|B^{(i,j)}(k)|}{|E|} & \text{if } \sum_{k=1}^{m} f^{(i,j)}(k) > 0 \\ \infty & \text{otherwise} \end{cases} \tag{7}$$

13

i.e. it is the sum of the transmission power and the number of blocked resources; note that we normalize each of these quantities (with the maximum transmission power and the maximum number of transmitter-receiver pairs, respectively) so that both terms take comparable values in the (0,1] interval.

5.2 Frequency allocation algorithms

Once a candidate minimum cost path has been identified, an interference-free channel allocation must be determined. We propose the following frequency allocation mechanisms:

(a) Link-by-link greedy frequency allocation (LLG)

LLG is motivated by the use of similar greedy channel allocation schemes in linear cellular networks. Channel allocation is performed along the candidate path in a hop-by-hop manner, starting from the origin node and moving towards the destination, selecting an available channel for each link (and updating blocked frequencies after each allocation).

We illustrate how this scheme operates by a simple example. Consider an isolated path from the source node S to the destination D as shown in figure 2(a). Node S needs to establish a session to node D via the path that consists of the nodes A,B and C. We list in parentheses below each link the frequency channels that are available at every instant and above each link the frequency channel that gets reserved at every step. We show how frequency allocations are made link by link from node S towards D and which one- and two-hop neighbors are blocked from transmitting according to our interference model. In this example we were fortunate, in that all links had available channels when they were needed. If however, link (S, A) had chosen frequency f_3 (as shown in figure 2(b)) it would have left links (A, B) and (B, C) with only one channel available and clearly the call would have been blocked.

Next we describe the algorithm for the general case of a path p that consists of k nodes $i_1, i_2, \cdots, i_k$. Each link (i_j, i_{j+1}) is associated with a "pool" of available frequency channels denoted by $F(i_j, i_{j+1})$. Every node i_j is aware of its immediate next hop neighbor i_{j+1} in the path.

LLG algorithm:

1. $j = 1$

2. i_j randomly selects a frequency channel $f_x \in F(i_j, i_{j+1})$ for transmission over (i_j, i_{j+1})

3. Block neighboring links according to the interference model and update their $F(\cdot, \cdot)$.

4. $j = j + 1$

5. If $j = i_k$ terminate; else goto 2.

In LLG there always exists a possibility that the path may run out of resources, even though all nodes may have had initially at least one channel available. Clearly, given a path and the available channels of each link we have a finite number of permutations, some of which result in feasible assignements, whereas others don't. Nevertheless, LLG can be implemented in a fully distributed manner, without the complexity of an exhaustive search.

(b) Most congested link first (MCLF) frequency allocation

In order to increase the possibility of producing a feasible allocation, we propose a second heuristic in which we assume full knowledge of the path and the available frequencies at each node along the path. Such information allows us to give priority to nodes with smaller numbers of available channels so that they can make their reservations first (since those are the nodes more likely to run out of resources). We describe the algorithm for a path p; let $E_p = \{\text{all} \ (i, j) \in p\}$ and assume that $\forall (i, j) \in E_p$, $F(i, j)$ and its cardinality $|F(., .)|$ are known

MCLF algorithm:

1. Sort elements of E_p in increasing order starting with minimum value of $|F(., .)|$

2. Remove first element (i, j) of E_p

3. If $|F(i, j)| > 0$, randomly select $f_x \in F(i, j)$ for transmission over (i, j); else go to 6.

4. Block neighboring links according to interference model and update their $F(\cdot, \cdot)$.

5. If $E_p \neq \emptyset$ go to 1.

6. Terminate

Note that there still exists some arbitrariness and unpredictability, in the way the frequency channels

are selected, as was the case with LLG, but we believe that the probability of success is higher, since we expect to avoid situations where nodes with many available channels may block neighboring nodes in the path with a single channel, just because of an unfortunate choice of transmission frequency.

6   Simulation model

We evaluate the performance of the proposed algorithms for a variety of networks. Each node's location is randomly generated within a square region of 100×100 units and we model both sparse and dense topologies by examining networks of sizes between $N = 10$ and 50 nodes. As was noted in section 2, existence of a link depends only on the Euclidean distance between two nodes and the propagation loss exponent which is considered equal to $\alpha = 2$. All links are full-duplex and error-free and, without loss of generality, we assume that a node transmitting at a power level of $P_o = 0.1$ is received by all nodes located within distance $d \leq d_o = 10$ units. Using this as a reference, we may compute the transmission range corresponding to our choice of P_{max} by,

$$\frac{P_{max}}{P_o} = (\frac{d_{max}}{d_o})^\alpha \tag{8}$$

For each network size examined, we generate 100 connected random topologies. Note that we are only interested in connected topologies (performance of a partitioned network in which not all nodes may communicate with each other was not considered). All nodes have equal amounts of initial energy E^o and unless otherwise specified this energy level is sufficient for the duration of the simulation. Initially, we assume that energy is only consumed during transmission and for the whole duration of a session. The effects of processing power are considered in a separate section.

In all the experiments, calls arrive independently at each node according to a Poisson process with average rate λ. We consider values of λ between 0 and 1 and for every new call arrival, the destination is uniformly selected among the remaining nodes. Average session durations are exponentially distributed with average duration $\mu = 1$.

Performance is measured by the blocking probability P_b, the average energy per session E_s and

the performance yardstick Y that we defined in equation 2. In some of the experiments we were also interested in additional performance metrics, such as the the average lifetime of the nodes or the average path length. Lifetime of a node is defined as the time from the beginning of the simulation until the node has exhausted its energy reserves.

7 Performance results: unlimited bandwidth resources

7.1 Blocking probability (P_b)

Figures 3 and 4 illustrate graphically the blocking probability as a function of the offered load per node, for $N = 20$ and the cases of $d_{max} = 30$ and $d_{max} = 50$. Each node is equipped with 5 transceivers and we plot the average value of P_b for metrics M1, M2 and two different cases of M3, one that accounts only for the residual energy of the receiver ($W_p = 0, W_e = 1$) and one that equally considers transmission power and residual energy ($W_p = 1, W_e = 1$). We observe that in all cases P_b increases as the offered load increases. In both figures 3 and 4, metric M1 exhibits the worst performance; the reason is that it always searches for the lowest total power path available which is usually a path with a larger number of hops, and therefore it results in utilization of larger number of transceivers. Metric M1 results also in heavier utilization of some paths while at the same time others may consist of idle nodes. Metric M2 partially solves this problem and the improvement in P_b is greater as d_{max} increases since denser connectivity provides additional paths and incoming traffic can be spread over the network more effectively. This is shown in figure 4 for $d_{max} = 50$.

Metric M3 achieves much lower P_b, especially when the transmission range increases, mainly because of its property to favor direct links from source to destination (such links are more likely to exist for higher d_{max}), but also because of the tendency to spread the traffic more evenly among paths in order to balance energy expenditures among all nodes. However, this occurs at the cost of increased energy expenditures as we will see in subsequent results.

Similar results have also been observed for 10- and 50-node topologies [11, 12] but are not included in this paper due to space limitations.

To verify our intuition that metrics M1 and M2 provide, on average, paths with larger number of hops, we computed the average number of hops per accepted session for the above simulations. Results for the cases of $\lambda = 0.1$ and 0.5 are summarized in table 2. Each cell in the table consists of the average number of hops per session and the corresponding standard deviation. These results clearly indicate that metrics M1 and M2 select, on average, paths with larger number of hops.

7.2   Average energy per session (E_s)

Figures 5 and 6 depict the average energy per accepted session (E_s) versus the arrival rate λ, again for $N = 20$ and for the cases of $d_{max} = 30$ and $d_{max} = 50$. As was expected, algorithms based on the first two metrics (M1 and M2) result in lower E_s, since they admit sessions on the lowest transmission power path available. However, this causes higher P_b as we observed in the previous section. Note also that for lower values of d_{max} (figure 5), E_s starts decreasing slowly for higher values of λ. The reason for such behavior is that when traffic load (and therefore blocking) is high, calls that require fewer hops from source to destination (ie fewer node resources) have a higher chance of being admitted. By increasing the transmission range ($d_{max} = 50$ in figure 6), we increase the network connectivity and hence the possibility to select between a multi-hop path and a direct link from source to destination. Metrics M1 and M2 primarily search for the multi-hop path (which would lead to lower E_s) and if not available the typical alternative is a direct link (with higher E_s). When P_b is higher the direct link is more likely to exist and this is why E_s exhibits the observed increase. On the other hand, metric M3 always looks for the direct link first and hence for increased range its behavior is not significantly affected by blocking.

7.3   Effect of node density

In this section we explore the effects of network size and node density on performance. We focus our attention on metrics M1 and M3 (with $W_p = W_e = 1$ and $d_{max} = 50$). In figures 7 and 8 we plot P_b as a function of the offered load, for network sizes of 10,20 and 50 nodes. Similarly, in figures 9 and 10, E_s is plotted for the same set of network parameters. We observe that while for metric M1 P_b increases with network size, for metric M3 it decreases and in particular it does not vary if we increase the size from

20 to 50 nodes. The two metrics react differently when the node density changes; in particular, if we increase the number of nodes and assume a fixed transmission range, then we increase the connectivity of the network (which implies reduced average distance between nodes and hence more short distance links). In that case, metric M1 attempts to route the calls over multiple short hop links and blocks more resources so that future calls experience higher P_b. Instead, metric M3 favors direct links (for $d_{max} = 50$ there are a lot of direct links) and hence future calls that cannot use the shortest paths still have good chance to get through along multi-hop paths. On the other hand, E_s decreases with network size because a denser network provides additional short hops. Note that in the case of M1 this decrease is more dramatic whereas for M3 it is not as significant.

We also conclude that metric M3 is less sensitive to network parameters. It is less sensitive to network size, since the difference both in P_b and E_s is relatively small from 20 to 50 nodes whereas for metric M1 it is more substantial. It is also less sensitive, in terms of P_b, to the average node degree. As shown in figure 11, in all 100 random topologies P_b did not vary as much (for the case of M3) with the average node degree, unlike M1 and M2.

7.4 Performance yardstick (Y)

In equation 2 we defined the global reward yardstick Y in order to address the interdependencies between blocking probability and energy usage on network performance. Figure 12 shows how Y varies with the arrival rate for $N = 20$ and $d_{max} = 30$. Note that even though metrics M1 and M2 seem to achieve higher values of Y, if λ increases (higher blocking) then all metrics exhibit similar performance. Similar conclusions can be drawn also from figure 13 where we have increased the number of nodes to 50. Here M3 outperforms the rest of the metrics for high values of arrival rates.

7.5 Effect of receive and processing power on energy consumption

Throughout the results presented so far, we assumed that the node processing power was negligible compared to transmission power. In actual systems however, wireless transceivers consume a significant amount of energy for signal processing and other functionalities that are necessary for relaying sessions.

A key feature of our simulation model enables us to quantify the effects of non-neglibible processing power. Although the actual processing energy depletion pattern depends on the processor and the operations it performs, we have assumed that each wireless transceiver consumes a constant amount of processing power (denoted P_{proc}) whenever it is "serving" an active session. This is not unreasonable, since receive and processing power do not depend on distance from the transmitter.

For simulation purposes we consider 100 random topologies with $N = 20$ nodes. We have pre-calculated the average transmission power for all sample networks in the case of $P_{max} = 2.5$ and have found it equal to $\bar{P}_{tr} = 1.1$. Given a fixed arrival rate ($\lambda = 0.5$), we ran simulations for all metrics for different values of processing power P_{proc} equal to 0, 1%, 10%, 30% and 50% of the average transmission power. The results are depicted in figure 14. We observe that E_s has a higher slope for metrics M1 and M2 compared to M3. Even though for $P_{proc} \leq 0.2 \times \bar{P}_{tr}$ the metrics M1 and M2 perform better, after P_{proc} exceeds 25% of $\bar{P}_{tr}$ the metric M3 overtakes them. This is because M1 and M2 favor the use of multiple small hops which implies that multiple nodes are engaged in processing.

7.6   Network "lifetime"

In all previous experiments we assumed that nodes always have sufficient energy to continue operating until the completion of the simulation. In this paragraph, we study the performance of networks under the effect of node failures due to the exhaustion of their energy reserves. We show how the admitted traffic gets affected when some nodes run out of battery power and turn their transceivers "OFF". We also compare the average - per node - energy consumption for different metrics.

In figure 16(a) we plot the cumulative number of accepted calls versus the simulation time for the case of the network of figure 15(b) and for $\lambda = 0.3$. All experiments were run for an adequate number of events so that the network would get partitioned. We observe that as nodes begin to switch to OFF state, the rate of accepted calls starts decreasing. In order to maintain acceptable performance, an algorithm should try to balance the energy consumption evenly among all nodes so that the time until the first node turns OFF is maximized. Metric M3 achieves that to some significant extent, since the

interval from the beginning of the simulation until the first node turns its power OFF is longer. The results are not sensitive to the rate of offered traffic, as far as the relative performance of the metrics is concerned. Of course higher traffic load results in shorter network lifetime for all metrics.

When we increased the maximum transmission range to $d_{max} = 50$ (see figure 15(c)), we observed that use of metric M3 results in more rapid network partirioning. In figure 16(b), we plot the cumulative number of accepted calls versus the simulation time ($d_{max} = 50$). Due to the nature of M1 to "favor" lower power paths, the first node will turn OFF faster under M1 than under M3. However, due to the increased transmission range, we observe from the connectivity map that there exist many direct links between nodes which will be "favored" by M3. Even though M3 tries to maintain a balance among all nodes in terms of energy consumption, most of its nodes will turn OFF shortly after the first node does so, and the average time until the network gets partitioned is shorter under M3 than under M1.

To verify that M3 exhibits better fairness properties in terms of energy expenditures per node, we ran some experiments for all three connectivity maps of figure 15 and examined the energy reserves of all nodes before any node had run out of energy reserves. Let $E = [E_0, E_1, \cdots, E_{N-1}]$ represent a vector of energy expenditures, with E_i being the ratio of energy spent versus initial energy for node i. Each cell in table 3 consists of 4 elements, the mean and standard deviation of E (top row) and the minimum and maximum values of E (bottom row) respectively. We observe that the average energy per node for the case of M1 does not vary by increasing the transmission range, whereas this is not the case for M3. Note also that for M3, E_s exhibits smaller standard deviation, which can be translated as a fairness property of M3, in that the energy levels are more balanced among all nodes.

8 Performance results: limited bandwidth resources

8.1 Evaluation of frequency allocation algorithms

The two frequency allocation heuristics LLG and MCLF are compared in terms of P_b under the assumption that the minimum cost path is given by MPM metric. We have simulated 100 random networks with $N = 20$ for the parameters shown in table 4. Figure 18 graphically illustrates the relative perfor-

mance of LLG and MCLF for $m = 6$ and $m = 9$ (where m is the number of frequency channels). As was expected, MCLF provides consistently a slight improvement over LLG. By increasing the number of channels we get lower values of P_b and MCLF provides relatively better improvement. Of course the trade off that needs to be considered is the need for centralized operation under MCLF versus the fully distributed nature of LLG.

In order to gain additional insight into the performance of the frequency allocation schemes, we compare them against exhaustive search mechanisms. Of course such a comparison can only be applied in a small and limited number of topologies due to the high complexity of exhaustive search. Tables 5 and 6 summarize the performance results (P_b and E_s respectively) for the topologies shown in figure 17. The first element of each cell in these tables (ExSrch) corresponds to a scheme in which for each new call request a search is performed among all paths and all frequency allocations and the one with minimum total power is selected. To illustrate the amount of complexity of ExSrch, it suffices to say that for the network of figure 17(b), when $m = 5$ there exist nearly a million possible route and frequency assignments between nodes 4 and 6. The next element in each cell of the results (ESMP) corresponds to a scheme that limits the search to the minimum power path only (Exhaustive Search of Minimum-cost Path) in order to reduce the size of the solution space. In a sense such a mechanism operates as an admission control policy, in which a session gets blocked from the system if no feasible allocation can be placed along the minimum cost path. The last two elements of each cell correspond to the heuristics described earlier in the paper (MCLF and LLG).

In most of the cases, ExSrch provides better performance in terms of blocking probability. In fact this improvement is even higher for larger d_{max} since additional paths become available and ExSrch examines a larger solution space. Of course this comes at the cost of higher E_s since the call admission is not now restricted only to the minimum cost path.

It is very interesting to point out that in some isolated situations either the ESMP scheme or one of the heuristics (LLG or MCLF) resulted in better performance than the ExSrch. For example, in the case of $d_{max} = 35$, $\lambda = 0.5$ and $m = 6$, ESMP and MCLF provided lower blocking than ExSrch. Similarly,

for $d_{max} = 35$, $\lambda = 0.7$ and $m = 3$, ExSrch had the worst performance among all methods. Even though such a behavior was not expected, it is a consequence of the fact that ExSrch does not guarantee a global optimum since it works on a per call basis. In fact, a global optimum would be obtainable only if complete knowledge of the traffic pattern was available prior to the beginning of the simulation. Of course, for a given call and a given network state, no heuristic can provide a better solution than the ExSrch. Sometimes however, there may exist more than one valid frequency assignments along the minimum cost path; every valid assignment though, has the same effect on blocking of neighboring nodes (ie the number of blocked transmissions does not depend on the selected channel but it depends only on the power level and the receiving node) but may have different effect on future calls, depending on the future traffic characteristics. Hence, it is <u>possible</u> that by accepting a call via the exhaustive search (a call that would have been otherwise blocked via one of the heuristics) future calls may be adversely affected. Of course, all the heuristics work on a per call basis, so, on average, such situations are not very likely to occur and, as our results indicate, they do happen only rarely.

It is also of interest to observe the difference between the worst and best performing algorithm for each set of arrival rates. For the case of $m = 3$, when the traffic load is light, there is significant difference in P_b between the values of ExSrch and the rest, whereas for heavy traffic they are fairly close. On the other hand for $m = 6$, the difference seems to be greater for heavy traffic, which is intuitively clear, since the heuristics are expected to perform less poorly when traffic is light. A possible explanation for the case of $m = 3$ is that because of the very small number of channels, it becomes less likely to determine a feasible assignment along the minimum power path; by contrast, the ExSrch has the flexibility of using higher power paths, even the highest available (which could be a long single-hop direct transmission) as long as a channel assignment can be found. The increased values of E_s, produced by ExSrch in table 6, verify our explanation.

8.2 Performance characteristics of link metrics for use with link-by-link greedy frequency allocation scheme

In this section we evaluate the relative performance of MPM and PIM metrics for joint use with the LLG channel allocation algorithm. Figures 19 − 24 illustrate graphically the relative performance of MPM and PIM. We have simulated networks of $N = 10$ and $N = 20$ nodes with $d_{max} = 30$ and $d_{max} = 50$. For each pair of N and d_{max} we have run simulations for 100 randomly generated topologies and we have computed the average values of our performance measures, namely P_b, E_s and Y, given that the number of frequency channels remained constant ($m = 6$). From the graphs shown in figures 19 and 20 we observe that the use of PIM provides better performance in terms of blocking probability in all cases. Note that when the network becomes denser (larger N), P_b increases. It is also of interest that PIM achieves better improvement in P_b relatively to MPM for larger d_{max}. By contrast, MPM performs better when the performance metric is the average energy per session (figures 21 and 22). Clearly, this improved performance can be attributed to the fact that MPM's only criterion for selecting the candidate path is the minimum power consumption. Notice again that when the transmission range increases, for equal sized networks MPM provides much better improvement in the values of E_s. Finally, figures 23 and 24 show the performance yardstick versus the call arrival rate. Except for the case of $N = 20$ and $d_{max} = 30$, MPM seems to outperform PIM.

9 Conclusions

We developed and compared a set of heuristics that determine end-to-end unicast paths with sufficient bandwidth and transceiver resources in fixed all-wireless topologies, for routing connection-oriented traffic under energy constraints. We examined first the case of limited transceivers and concluded that a link metric combining transmission power requirements and residual energy information (M3) provides best performance in terms of blocking probability and energy consumption and achieves better fairness in terms of energy expenditures per node. In the case of limited bandwidth resources and infinite transceivers, we addressed jointly the problems of route selection and of channel allocation.

We proposed a scheme to select a candidate path (as the minimum cost path in terms of transmission energy and blocking) and then to produce an interference-free frequency allocation. The whole area of energy-efficient ad-hoc networking which is currently receiving increased attention, involves numerous and diverse trade-offs. The work reported in this paper concerns only one facet of it. [4]

References

[1] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. Algorithms for energy-efficient multicasting in static ad hoc wireless networks. *To Appear in Mobile Networks and Applications (MONET)*, 2001.

[2] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. Energy-efficient broadcast and multicast trees in wireless networks. *To Appear in Mobile Networks and Applications (MONET)*, 2001.

[3] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. Algorithms for bandwidth-limited energy-efficient wireless broadcasting and multicasting. In *Proceedings Milcom 2000, Los Angeles*, October 2000.

[4] C. Sankaran and A. Ephremides. Multicasting with multiuser detection in ad-hoc wireless networks. In *Proceedings of the 2000 International Zurich Seminar on Broadband Communications*, February 2000.

[5] V. Rodoplu and T. Meng. Minimum energy wireless networks. *IEEE JSAC*, 17, August 1999.

[6] J. H. Chang and L. Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *Proceedings Infocom 2000, Tel-Aviv, Israel*, March 2000.

[7] R. Berry, S. Finn, R. Gallager, H. Kassab, and J. Mills. Minimum energy and delay routing in packet radio networks. In *Proceedings of the 3rd Annual ARL FEDLAB Symposium*, February 1999.

[8] A. Michail and A. Ephremides. A distributed routing algorithm for supporting connection-oriented service in wireless networks with time-varying connectivity. In *Proceedings 3rd IEEE Symposium on Computer and Communications (ISCC 98), Athens, Greece*, June 1998.

---

[4]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.

[9] A. Michail, W. Chen, and A. Ephremides. Distributed routing and resource allocation for connection-oriented traffic in ad-hoc wireless networks. In *Proceedings of the Conference on Information Sciences and Systems, Princeton, New Jersey*, March 1998.

[10] D. Bertsekas and R. Gallager. *Data Networks*. Prentice-Hall, 1992.

[11] A. Michail. *Routing and Scheduling Algorithms in Resource-limited Wireless Multi-hop Networks*. PhD thesis, University of Maryland, College Park, 2000.

[12] A. Michail and A. Ephremides. Energy-efficient routing of connection-oriented traffic in ad-hoc wireless networks. In *Proceedings of the 11th IEEE International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC), London, UK*, September 2000.

Link	Status	Conflict type
(A, C)	Blocked	P2
(C, A)	Blocked	P1
(C, D)	Free	-
(D, C)	Free if $d(A, B) < d(A, C)$	S1
(B, K)	Blocked	P4
(K, B)	Blocked	P3
(K, L)	Free if $d(KL) < d(KB)$	S2
(L, K)	Free	-

Table 1: Frequency blocking status for example of figure 1

	$\lambda = 0.1$		$\lambda = 0.5$	
Metric	$d_{max} = 30$	$d_{max} = 50$	$d_{max} = 30$	$d_{max} = 50$
M1	4.06 ; 0.59	3.89 ; 0.44	3.62 ; 0.44	3.51 ; 0.36
M2	4.10 ; 0.58	3.98 ; 0.41	3.62 ; 0.42	3.46 ; 0.31
M3(0,1)	2.87 ; 0.38	2.79 ; 0.30	1.59 ; 0.14	1.59 ; 0.14
M3(1,1)	2.87 ; 0.38	2.80 ; 0.30	1.60 ; 0.14	1.60 ; 0.14

Table 2: Average number of hops and standard deviation per admitted session for $N = 20$

$N = 10, \lambda = 0.3$			
Metric	$d_{max} = 30$	$d_{max} = 40$	$d_{max} = 50$
M1	0.314 ; 0.256	0.277 ; 0.165	0.277 ; 0.165
	(0.057 - 0.844)	(0.086 ; 0.536)	(0.089 ; 0.537)
M3	0.318 ; 0.231	0.287 ; 0.079	0.328 ; 0.100
$(W_p = W_e = 1)$	(0.059 - 0.744)	(0.171 - 0.383)	(0.147 - 0.488)

Table 3: Consumed energy per node for networks of figure 15

N	20
P_{max}	2.5
d_{max}	50
λ	$\in [0.1, 0.9]$
μ	1
m	6 and 9
Link Metric	MPM

Table 4: Simulation parameters for comparing LLG with MCLF

	$m^f = 3$		$m^f = 6$		
λ	$d_{max} = 35$	$d_{max} = 40$	$d_{max} = 35$	$d_{max} = 40$	Algorithm
	0.193	0.129	0.006	0.005	ExSrch
0.1	0.220	0.190	0.013	0.009	ESMP
	0.224	0.187	0.013	0.010	MCLF
	0.226	0.197	0.015	0.013	LLG
	0.442	0.373	0.110	0.075	ExSrch
0.3	0.442	0.399	0.122	0.093	ESMP
	0.446	0.409	0.129	0.101	MCLF
	0.457	0.413	0.134	0.104	LLG
	0.556	0.506	0.258	0.202	ExSrch
0.5	0.555	0.519	0.250	0.208	ESMP
	0.559	0.521	0.253	0.218	MCLF
	0.565	0.528	0.262	0.209	LLG
	0.634	0.593	0.357	0.297	ExSrch
0.7	0.627	0.594	0.359	0.304	ESMP
	0.626	0.589	0.362	0.297	MCLF
	0.629	0.595	0.368	0.325	LLG

Table 5: Blocking probabilities for topology of figure 17

λ	$m^f = 3$		$m^f = 6$		
	$d_{max} = 35$	$d_{max} = 40$	$d_{max} = 35$	$d_{max} = 40$	Algorithm
	1.409	1.472	1.416	1.393	ExSrch
0.1	1.329	1,326	1.406	1,387	ESMP
	1.336	1.306	1.389	1.376	MCLF
	1.315	1.309	1.403	1.382	LLG
	1.364	1.489	1.430	1.454	ExSrch
0.3	1.252	1.287	1.379	1.371	ESMP
	1.276	1.278	1.396	1.374	MCLF
	1.236	1.237	1.346	1.380	LLG
	1.290	1.454	1.466	1.596	ExSrch
0.5	1.194	1.274	1.361	1.391	ESMP
	1.170	1.252	1.352	1.401	MCLF
	1.177	1.235	1.324	1.370	LLG
	1.232	1.457	1.426	1.393	ExSrch
0.7	1.148	1.250	1.355	1.396	ESMP
	1.143	1.236	1.344	1.381	MCLF
	1.108	1.189	1.342	1.377	LLG

Table 6: Energy expenditures for topology of figure 17
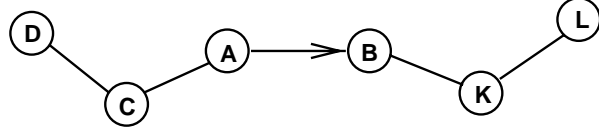
Figure 1: Example network to illustrate interference model

(a) Case 1

Call Request Blocked
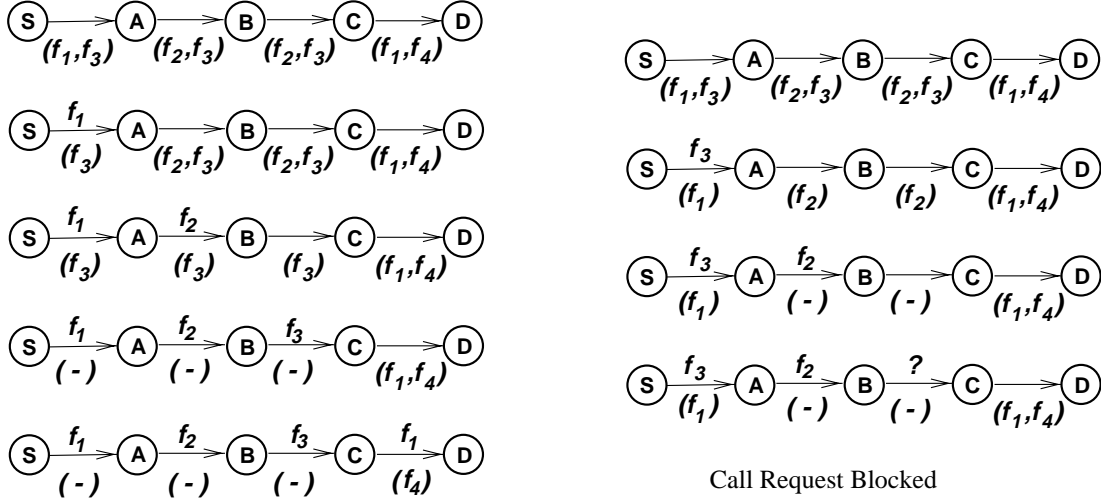
(b) Case 2

Figure 2: Example of LLG frequency allocation algorithm
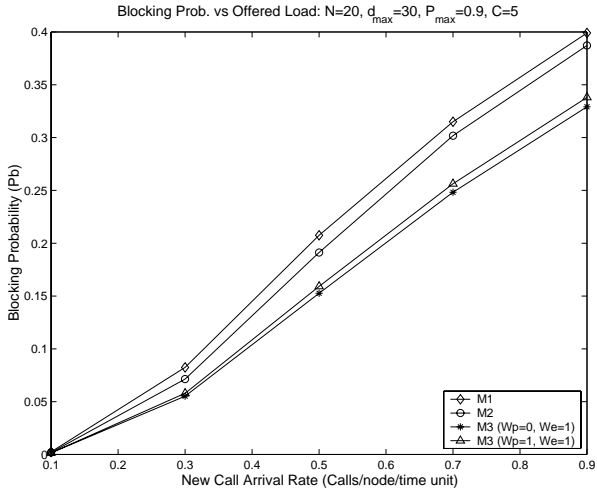
Figure 3: Blocking probability versus per node arrival rate for $N = 20$, $d_{max} = 30$
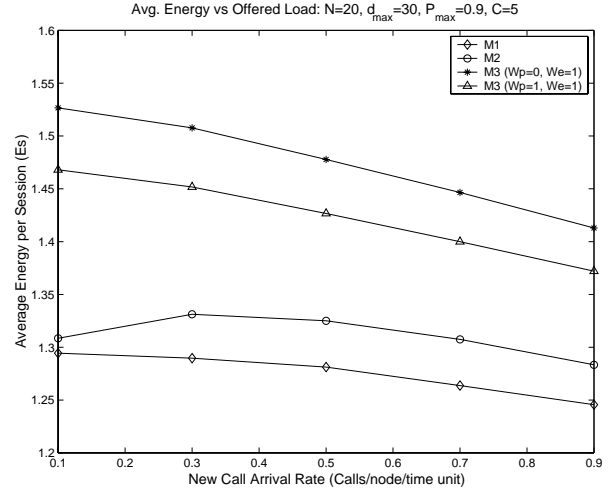
Figure 5: Average energy per session versus per node arrival rate for $N = 20, d_{max} = 30$
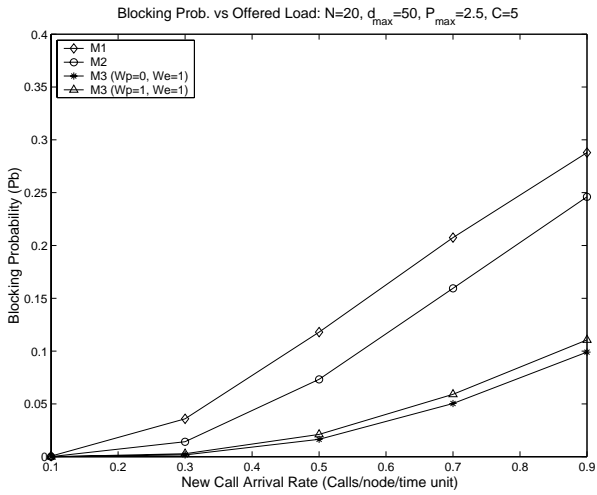
Figure 4: Blocking probability versus per node arrival rate for $N = 20, d_{max} = 50$
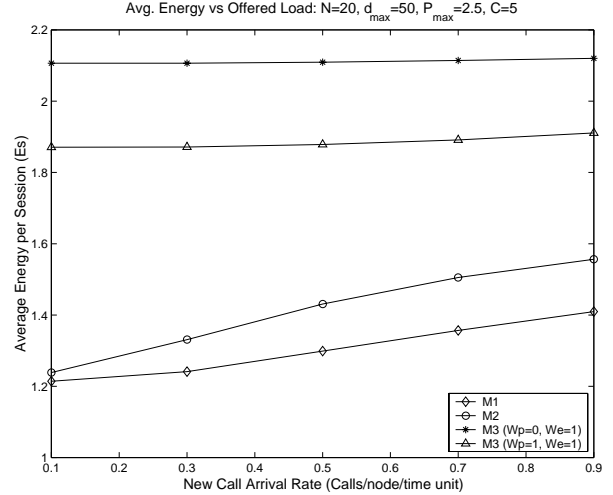
Figure 6: Average energy per session versus per node arrival rate for $N = 20, d_{max} = 50$
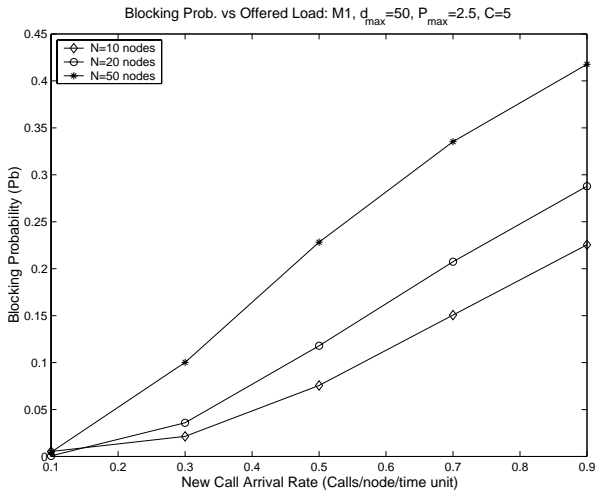
Figure 7: Blocking probability vs arrival rate for variable network sizes, M1 and $d_{max} = 50$
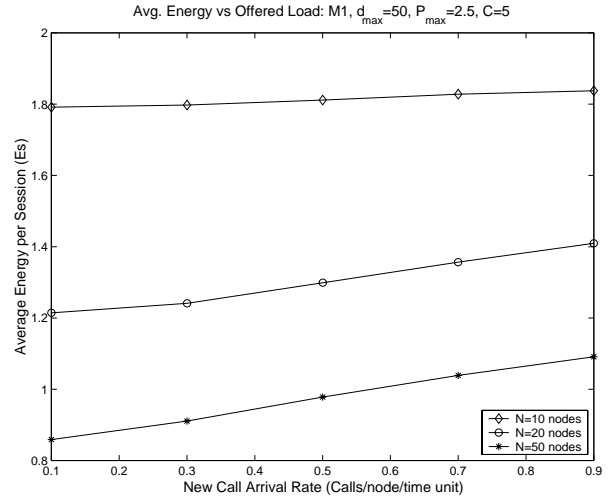
Figure 9: Energy per session vs arrival rate for variable network sizes, M1 and $d_{max} = 50$
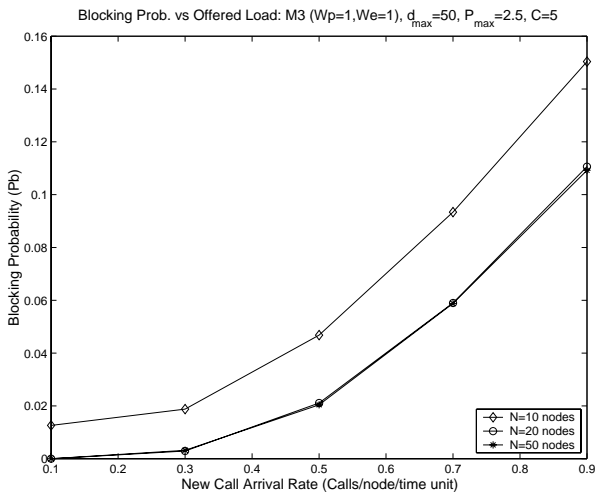
Figure 8: Blocking probability vs arrival rate for variable network sizes, M3 and $d_{max} = 50$
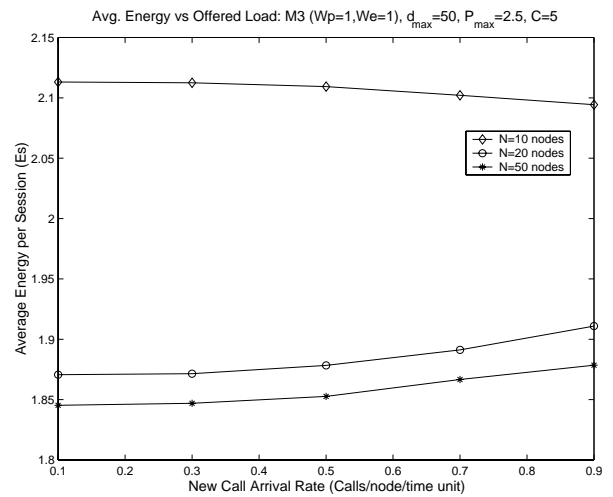
Figure 10: Energy per session vs arrival rate for variable network sizes, M3 and $d_{max} = 50$

33
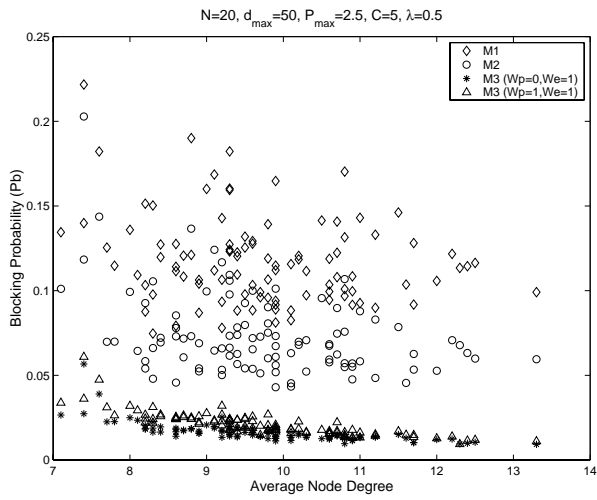
Figure 11: Blocking prob. vs average node degree, $N = 20, d_{max} = 50, \lambda = 0.5$
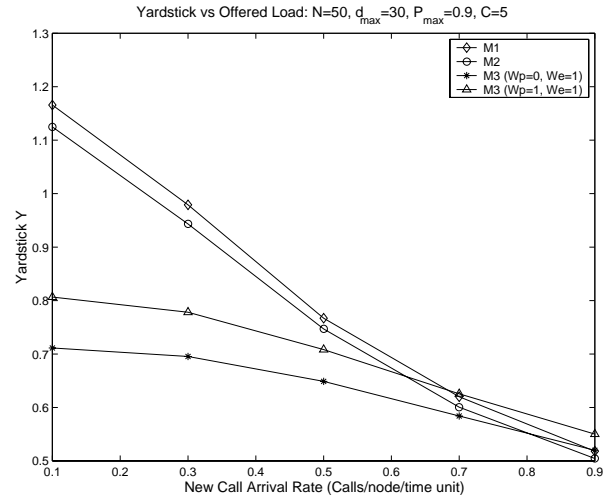
Figure 13: Yardstick Y vs. per node arrival rate for $N = 50, d_{max} = 30$.
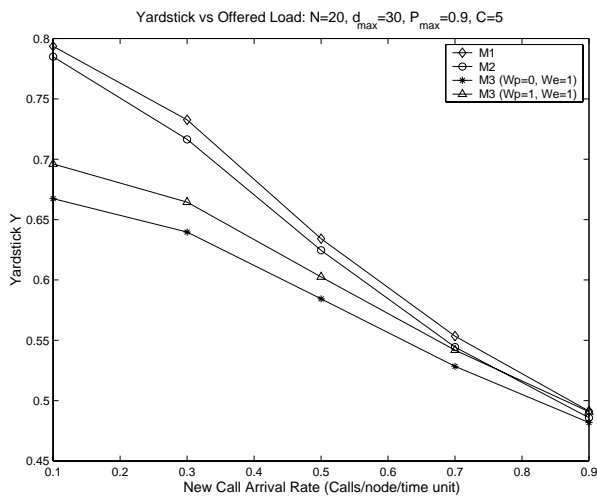
Figure 12: Yardstick Y vs. per node arrival rate for $N = 20, d_{max} = 30$.
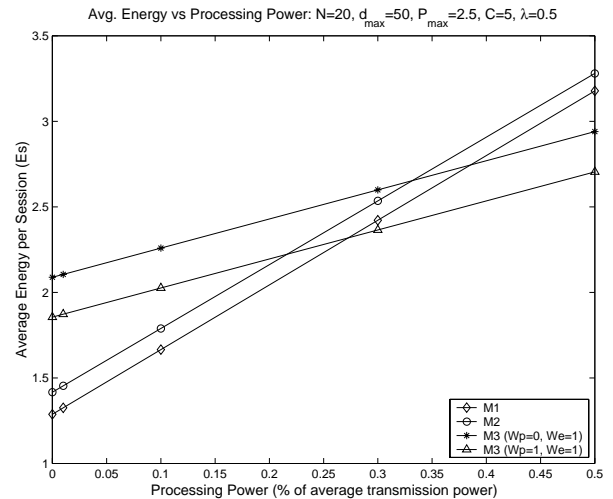
Figure 14: Effect of processing power on E_s for $N = 20, d_{max} = 50$

(a) $d_{max} = 30$ (b) $d_{max} = 40$ (c) $d_{max} = 50$
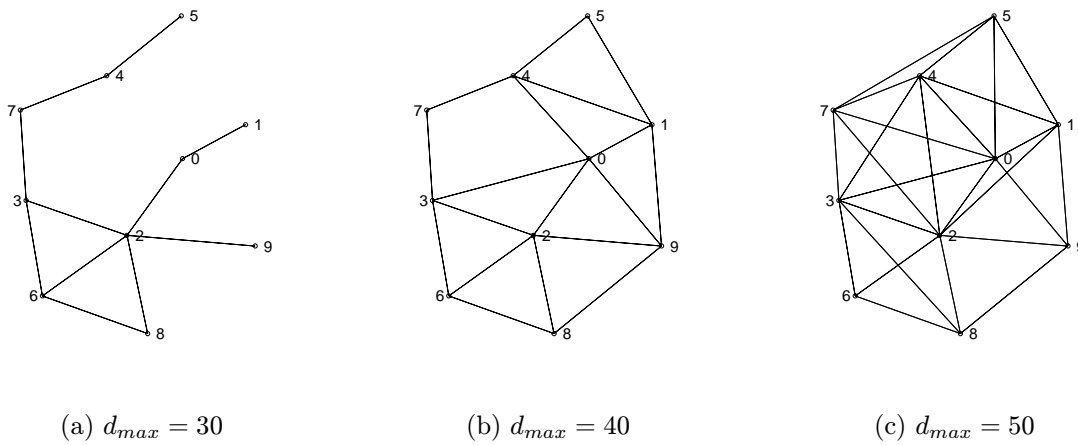
Figure 15: Sample network connectivities with N=10 for studying effects of energy exhaustion

(a) $N = 10$, $d_{max} = 40$ and $\lambda = 0.3$ (b) $N = 10$, $d_{max} = 50$ and $\lambda = 0.3$
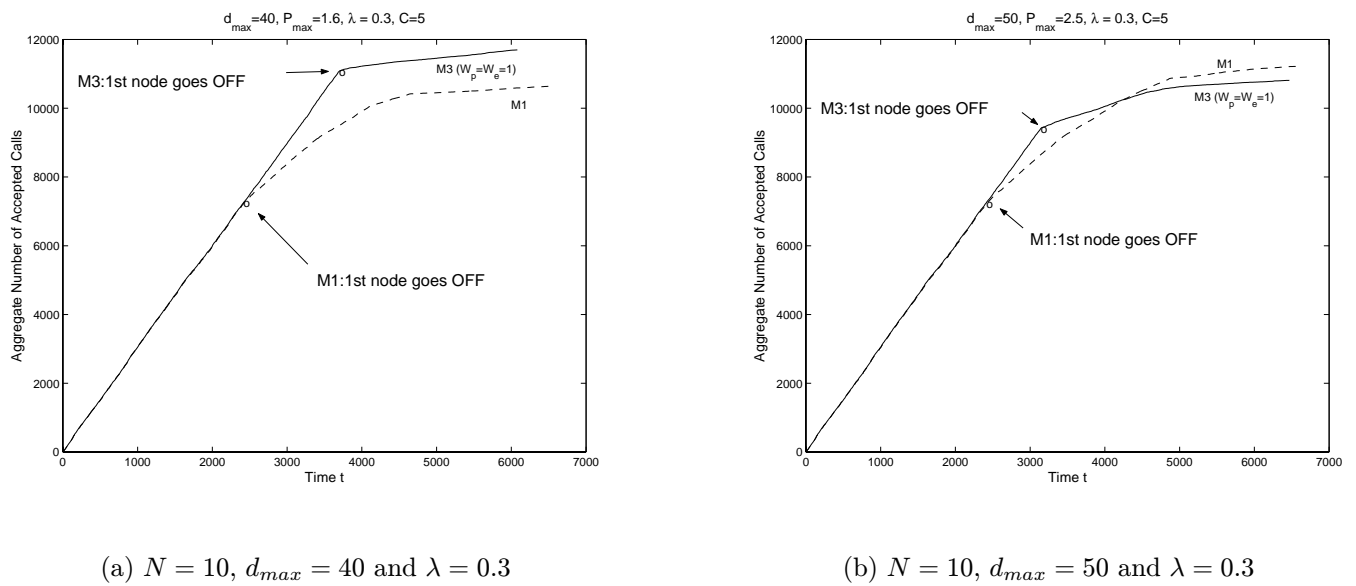
Figure 16: Cumulative number of accepted calls vs time

(a) $d_{max} = 35$ (b) $d_{max} = 40$
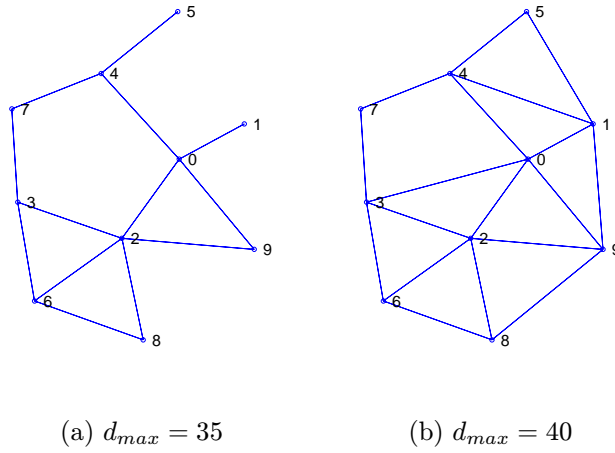
Figure 17: Example network topology for comparison of heuristics against exhaustive search methods
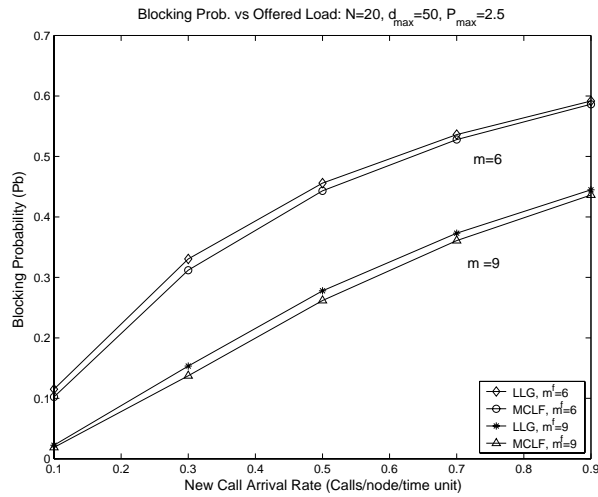
Figure 18: Comparison of LLG, MCLF through blocking probability and number of frequency channels, $N = 20, d_{max} = 50$
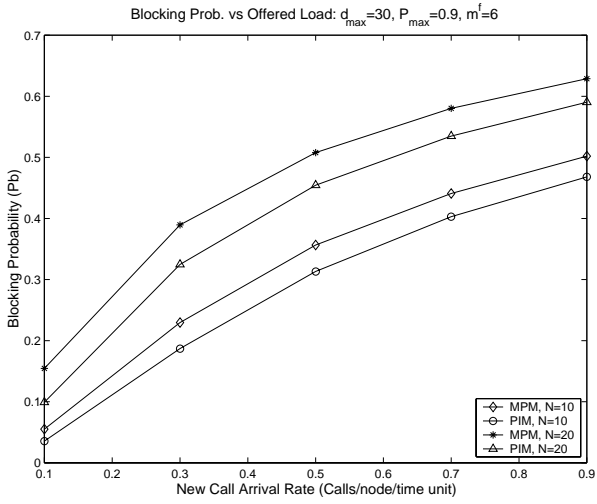
Figure 19: Comparison of MPM and PIM through blocking probability, $d_{max} = 30$
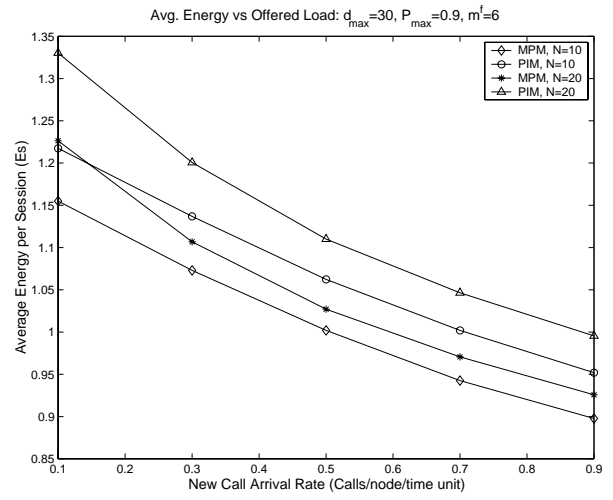
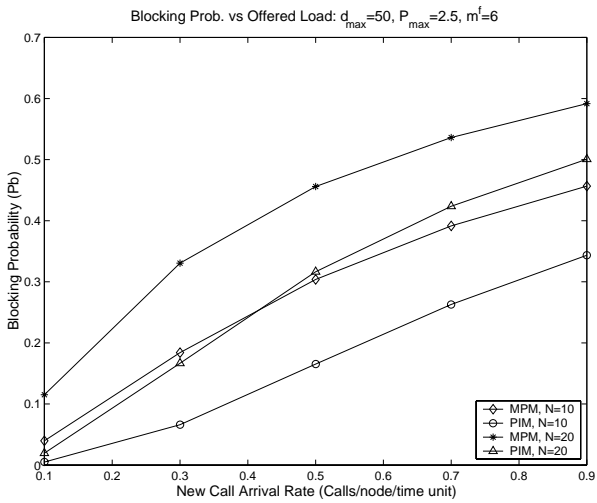Figure 21: Comparison of MPM and PIM through energy per session, $d_{max} = 30$

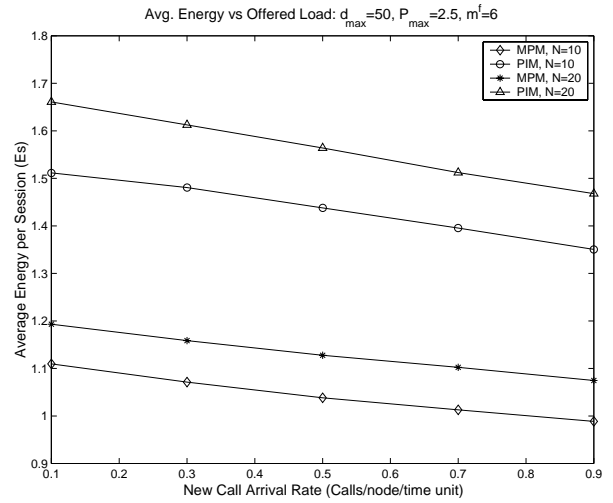Figure 20: Comparison of MPM and PIM through blocking probability, $d_{max} = 50$

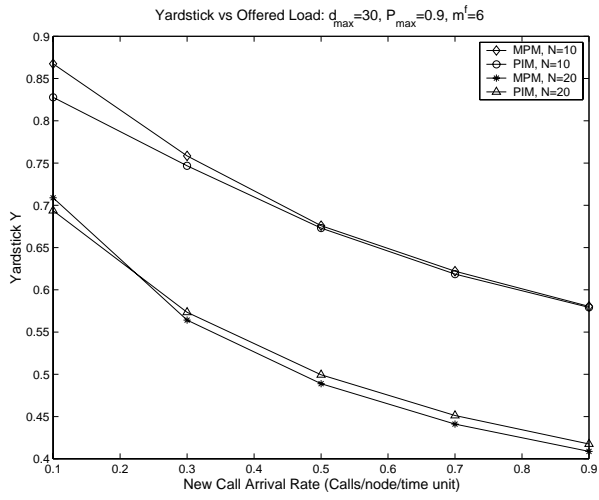Figure 22: Comparison of MPM and PIM through energy per session, $d_{max} = 50$

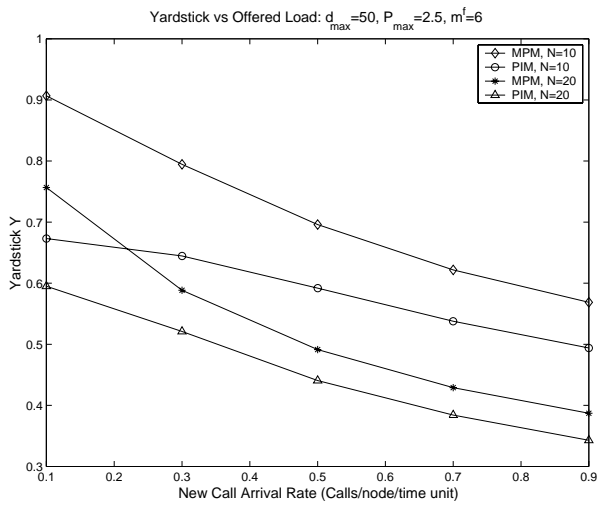Figure 23: Comparison of MPM and PIM through yardstick Y for $d_{max} = 30$

Figure 24: Comparison of MPM and PIM through yardstick Y for $d_{max} = 50$