TECHNICAL RESEARCH REPORT

Routing Session Traffic in Fixed All-Wireless Networks under Energy and Bandwidth Limitations

by Anastassios Michail and Anthony Ephremides

CSHCN TR 2001-9
(ISR TR 2001-16)

ROUTING SESSION TRAFFIC IN FIXED ALL-WIRELESS NETWORKS UNDER ENERGY AND BANDWIDTH LIMITATIONS [*]

Anastassios Michail
Anthony Ephremides
Department of Electrical and Computer Engineering
and Institute for Systems Research
University of Maryland, College Park, MD 20742

ABSTRACT

In this paper we study the effects of limited bandwidth resources in the development of energy-efficient routing algorithms for connection-oriented traffic in fixed wireless ad-hoc networks. A frequency division multiple access scheme is considered, in which nodes must schedule their transmissions by selecting frequency channels from a limited set in an interference-free fashion. In our earlier work, we had developed a set of algorithms for determining end-to-end unicast paths based on link metrics. We argue that in order to address the effects of limited frequency resources, such algorithms must be coupled with channel allocation mechanisms for providing conflict free frequency assignments over selected routing paths. To these ends, we propose a set of link metrics for selecting candidate routing paths and a set of heuristics for frequency allocation and evaluate their performance using our detailed simulation model.

INTRODUCTION

Next generation wireless tactical networks are expected to provide survivable multimedia communications among fixed or mobile users in the digital battlefield. An important limitation of wireless terminals and radios that will be used in such ad-hoc environments is their constraints in battery power due to the increasing demand for small, light-weight, portable devices. In prior studies ([1],[2]) we addressed the problem of routing connection-oriented unicast traffic with energy efficiency. We studied the trade-offs that arise by the flexibility of wireless nodes to transmit at different power levels and defined a framework for formulating the problem of session routing from the perspective of energy expenditure. Our preliminary approach was based on the simplifying assumption of plentiful bandwidth and any node could access the channel on demand, without need for contention and without interfering with neighboring nodes. In this paper, we shift our focus to the effects of limited bandwidth resources on energy-efficient unicast routing algorithms. *Wieselthier et. al.* have studied a similar problem in [3] for the case of wireless broadcasting and multicasting. We strictly consider unicast here and even though the objectives may be parallel, the actual algorithms, metrics and trade-offs are quite different.

The problems of bandwidth allocation and transmission power selection are inherently coupled. A node that increases its transmission power to reach a remote receiver, will possibly interfere with a larger set of neighboring nodes. On the other hand, if a path consisting of multiple short hops is used, the total power required for transmission may be lower, but bandwidth allocation complexity increases since it involves a larger set of nodes. Therefore the routing decision must be based on both energy and bandwidth considerations and to these ends we concentrate on developing routing algorithms that jointly achieve efficient usage of the available energy and interference-free channel allocations.

In the next section we describe our wireless network model. We continue with a description of the proposed algorithms followed by a summary of the most important simulation results. Note that throughout this paper, we assume a sufficiently large number of transceivers and consider networks that involve no mobility (e.g. sensor networks). Nonetheless, mobility effects can be addressed at a later time through the use of soft-failure mechanisms ([4]). Moreover, the possibility to use the transmission power as a metric to select a path adds a new degree of flexibility since nodes may adjust their power to maintain connectivity in the events of link failures.

---

[*] Prepared through collaborative participation in the Advanced Telecommunications/Information Distribution Research Program (ATIRP) Consortium sponsored by the U.S. Army Research Laboratory under Cooperative Agreement DAAL01-96-2-0002. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

WIRELESS NETWORK MODEL

We consider a network consisting of N nodes that may transmit at any power level $P \leq P_{max}$. Any node i can establish a direct link to any node j located $d \leq d_{max}$ distance units away, provided an available frequency channel can be found and the required transmission power is $P_{ij} \leq P_{max}$.

A total of m frequency channels $(f_1, f_2, \cdots, f_m)$ are available for use. Frequencies may be reused provided they do not cause interference. Although the use of a frequency division multiple access (FDMA) scheme introduces the difficult problem of assigning non-interfering frequencies to transmitting nodes, it is the most appropriate for our problem. Code division multiple access (CDMA) schemes do not allow nodes to handle simultaneous transmission and reception in the same frequency band and cannot be used. On the other hand, in time division multiple access (TDMA), the need to assign specific time slots results in a more difficult problem and is deferred for future research.

Each node i maintains a separate *channel status* vector for every neighbor. The channel-status vector of node i for transmission to node j is given by $f^{(i,j)} = [f^{(i,j)}(1), \cdots, f^{(i,j)}(m)]$, with:

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

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 and 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:
(i) nodes cannot transmit and receive in the same frequency
(ii) a node cannot simultaneously receive more than one signals in the same frequency
(iii) a node cannot transmit simultaneously to more than one neighboring nodes [1].

Each node maintains up-to-date information about the identities of its one-hop neighbors, the required power levels, and the status of the 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. [2]

ALGORITHMS

Our objective is to develop algorithms that determine an appropriate unicast path for each newly arriving session. A session can be admitted only if a path exists in which all nodes may simultaneously transmit and receive, using a conflict-free frequency allocation scheme. Since energy-efficiency is of paramount importance, the ideal algorithm should select among all available paths one with minimum aggregate transmission power which would result in minimum energy expenditures for the session under consideration. In the optimal case, the selected path should also have the least effect on the blocking of future calls.

If we were to assume "complete" information on the network state (that would get updated at the establishment and termination of each call), we could use a greedy algorithm that attempts to maximize a "reward" function on a per-call basis. Such a reward function can be possibly defined either as the total transmission power required, or as a linear combination of total power and number of blocked resources. Such a method however, requires the exhaustive search of a large state space which grows exponentially with the network size and therefore would be impractical except for trivially small networks.

An alternative approach would look for suboptimal solutions that rely on distributed local information acquired through the periodic exchange of control messages between neighboring nodes. Similarly to the mechanisms developed in [1, 2], we examine heuristic algorithms that utilize link metrics to capture local parameters and apply Bellman-Ford algorithm to determine the minimum cost path. Note however, that for the case of plentiful bandwidth, existence of one available transceiver on every node in the path would result in a finite cost route which was sufficient to guarantee admission of the session. By contrast, existence of at least one available frequency channel on every node is not sufficient to guarantee admission; 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 their identities. Moreover, the nodes of a path cannot select which frequency to use (among the set of available channels) independently. For every channel assignment made over one link, neighboring nodes that experience interference must update their blocked frequency table before they make their assignment.

To address the above issues we develop algorithms that that evolve in the following two stages: (i) a minimum cost path (as measured by energy and blocked

---

[1] we strictly consider unicast here; in a broadcast scenario this would an acceptable and in fact encouraged situation.

[2] This can be done via an underlying link-level mechanism that is not the purpose of this study.

resources) is first determined (the *candidate path*) and (ii) if an interference-free channel assignment can be determined along that path, the call request is admitted. We discuss in the following sections our proposed methods for achieving both items (i) and (ii).

Link Metrics for Minimum Cost Path

(i) Minimum power metric (MPM)

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

$$D_{ij}^{(MPM)} = \begin{cases} P_{ij} & \text{if } \sum_{k=1}^{m} f^{(i,j)}(k) > 0 \\ \infty & \text{otherwise} \end{cases} \tag{2}$$

MPM accounts for energy requirements only and provides the minimum transmission power path available.

(ii) Power and interference based metric (PIM)

To address the blocking effects a transmission may cause to neighboring nodes we define PIM, a metric based on transmission power and resulting interference. We introduce the following notation:

- $B^{(i,j)}(k)$: set of transmitter-receiver pairs that are blocked whenever node i transmits to j over frequency f_k
- $|B^{(i,j)}(k)|$: cardinality of $B^{(i,j)}(k)$
- $|E|$: cardinality of the set of all transmitter-receiver pairs.

PIM is then defined as follows:

$$D_{ij}^{(PIM)} = \begin{cases} \frac{P_{ij}}{P_{max}} + \frac{|B^{(i,j)}(k)|}{|E|} & \text{if } \sum_{k=1}^{m} f^{(i,j)}(k) > 0 \\ \infty & \text{otherwise} \end{cases} \tag{3}$$

Note that we normalize the transmission power with P_{max} and the number of blocked resources with $|E|$ so that both terms (which are measures of different quantities) take values in the (0,1] interval.

Frequency Allocation Algorithms

Once the *candidate* path has been identified, an interference-free channel allocation must be determined for the nodes that will be transmitting. We propose the following two frequency allocation mechanisms:

(i) Link-by-link greedy (LLG) allocation

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 and moving towards the destination, selecting an available channel for each link (and updating blocked frequencies after each allocation). We describe the algorithm for the 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 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 assignments, whereas others don't. Since we proceed with the allocations on a hop by hop manner with random selections we do not fully search the solution space (this would be extremely complex in non-trivial networks). By contrast, LLG can be implemented in a fully distributed manner and without the complexity of an exhaustive search.

(ii) Most congested link first (MCLF) 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 to 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. The algorithm proceeds as follows:

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 randomness 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 would block neighboring nodes in the path with a single channel, just because of an unfortunate choice of transmission frequency.

PERFORMANCE RESULTS

In this section we present a set of selected results obtained through our simulation model. A more detailed discussion accompanied with additional results can be found in [5]. We assume that calls arrive independently at each node following a Poisson distribution with average rate $\lambda \in [0, 1]$. Average session durations are exponentially distributed with $\mu = 1$. Performance is measured by the blocking probability P_b and the average energy per session E_s and the number of frequency channels is denoted by m.

In order to gain better insight on the performance of the frequency allocation schemes, we first compare them against exhaustive search mechanisms. Of course such a comparison can only be applied in small topologies but it still leads to significant remarks. Table 1 summarizes performance results for a 10 node topology with $d_{max} = 35$ (13 links) and $d_{max} = 40$ (17 links). The first element of each cell (ExSrch) corresponds to a scheme that upon a new call request performs a search among all paths and all frequency allocations and selects one that minimizes the total power (MPM metric). The next element (ESMP) corresponds to a scheme that limits the search to the minimum power path (Exhaustive Search of Minimum-cost Path). In a sense such a mechanism functions 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 correspond to the frequency allocation heuristics described earlier.

Observe that the improvement we get in P_b by using ExSrch is more significant for larger d_{max} since additional paths become available and ExSrch examines more possibilities. Of course this comes at the cost of higher E_s. It is also of interest to note that in some situations ExSrch results in higher P_b (e.g. when $d_{max} = 35, \lambda = 0.7$). This is due to the fact that ExSrch does not guarantee a global optimum since it works on a per call basis. 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 certain call and given the current network state, no heuristic can provide a better solution than the ExSrch but sometimes more than one valid frequency assignments along the minimum cost path may exist; every valid assignment though, has the same effect on blocking of neighboring nodes but may have different effect on future calls, depending on the future traffic characteristics. Hence, it is possible 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 (ie in a way worse than that of the heuristic). Of course, all the heuristics work on a per call basis, so on the average scale such situations are not very likely to happen and as our results indicate they do happen only in few cases, which is consistent with the above explanation.

Table 1: Comparison of frequency allocation mechanisms for $m = 3$ channels and MPM link metric

	P_b		E_s		
λ	$d_{max} : 35$	$d_{max} : 40$	$d_{max} : 35$	$d_{max} : 40$	Alg.
0.1	0.193	0.129	1.409	1.472	ExSrch
	0.220	0.190	1.329	1.326	ESMP
	0.224	0.187	1.336	1.306	MCLF
	0.226	0.197	1.315	1.309	LLG
0.3	0.442	0.373	1.364	1.489	ExSrch
	0.442	0.399	1.252	1.287	ESMP
	0.446	0.409	1.276	1.278	MCLF
	0.457	0.413	1.236	1.237	LLG
0.5	0.556	0.506	1.290	1.454	ExSrch
	0.555	0.519	1.194	1.274	ESMP
	0.559	0.521	1.170	1.252	MCLF
	0.565	0.528	1.177	1.235	LLG
0.7	0.634	0.593	1.232	1.457	ExSrch
	0.627	0.594	1.148	1.250	ESMP
	0.626	0.589	1.143	1.236	MCLF
	0.629	0.595	1.108	1.189	LLG

Comparison of LLG versus MCLF

The frequency allocation heuristics LLG and MCLF are evaluated by considering the MPM link metric. We have simulated 100 random networks of $N = 20$ and $d_{max} = 50$ and in figure 1 we plot P_b versus the average arrival rate λ for the cases of $m = 6$ and $m = 9$ frequency channels. As was expected, MCLF provides consistently a slight improvement versus 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 accounted for is the need for centralized operation by MCLF versus the fully distributed nature of LLG. Similar remarks can be drawn by considering other values of N and d_{max} ([5]).

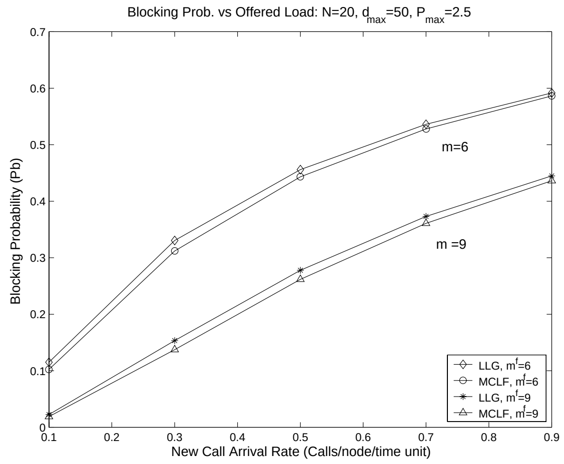

Figure 1: Comparison of LLG versus MCLF

Comparison of MPM versus PIM

We compare the performance of MPM and PIM metrics when the channel allocation is given by the LLG algorithm. For $N = 10$ and 20 and for $d_{max} = 50$ we have run simulations for 100 randomly generated topologies assuming $m = 6$ frequencies. Figures 2 and 3 illustrate graphically the relative performance of MPM and PIM in terms of P_b and E_s respectively. Use of PIM provides better performance in terms of P_b. Note also that when the network becomes denser (larger N), P_b increases. By contrast, MPM performs better when the performance metric is the average energy per session. 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.

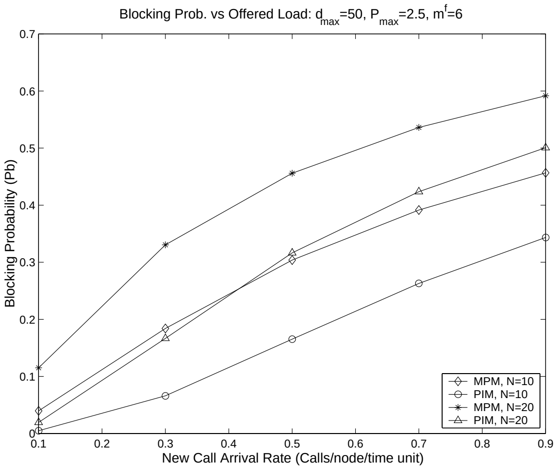

Figure 2: Blocking probabilities for MPM and PIM

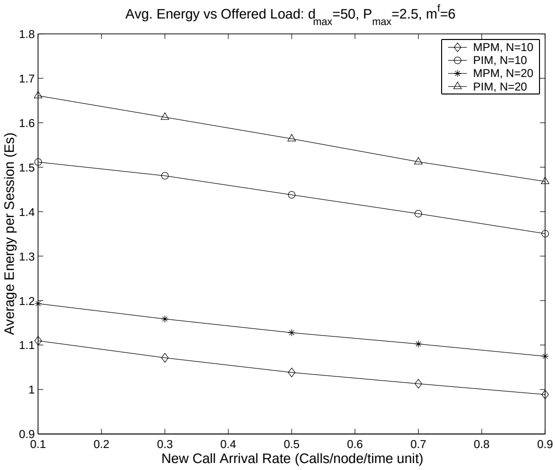

Figure 3: Energy consumption for MPM and PIM

CONCLUSIONS

We addressed the problem of session routing under energy and bandwidth limitations and proposed a set of algorithms that first identify candidate routing paths and then search for interference-free frequency allocations. Our results indicate that improved performance can be obtained by jointly considering the transmission power and the bandwidth allocation selection. We also demonstrated that even with a greedy channel allocation scheme, performance is comparable to that of exhaustive search mechanisms, whereas implementation complexity is extremely lower. [3]

REFERENCES

[1] A. Michail and A. Ephremides. Energy-efficient routing of connection-oriented traffic in ad-hoc wireless networks. In *Proceedings of the 4th Annual FEDLAB Symposium*, March 2000.

[2] 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.

[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] A. Michail and A. Ephremides. Routing algorithms in all-mobile wireless networks. In *Proceedings of the 1st Annual FEDLAB Symposium*, January 1997.

[5] A. Michail and A. Ephremides. Energy-efficient routing of connection-oriented traffic in ad-hoc wireless networks. *ISR Technical Report*, 2001.

---

[3] 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.