Wireless Netw (2010) 16:2263–2279
DOI 10.1007/s11276-010-0257-z

Joint multi-cost routing and power control in wireless ad hoc networks

Nikos Karagiorgas · Panagiotis Kokkinos · Christos Papageorgiou · Emmanouel Varvarigos

Published online: 17 June 2010

Abstract In this work we study the combination of multi-cost routing and adjustable transmission power in wireless ad hoc networks, so as to obtain dynamic energy- and interference-efficient routes to optimize network performance. In multi-cost routing, a vector of cost parameters is assigned to each network link, from which the cost vectors of candidate paths are calculated. Only at the end these parameters are combined in various optimization functions, corresponding to different routing algorithms, for selecting the optimal path. The multi-cost routing problem is a generalization of the multi-constrained problem, where no constraints exist, and is also significantly more powerful than single-cost routing. Since energy is an important limitation of wireless communications, the cost parameters considered are the number of hops, the interference caused, the residual energy and the transmission power of the nodes on the path; other parameters could also be included, as desired. We assume that nodes can use power control to adjust their transmission power to the desired level. The experiments conducted show that the combination of multi-cost routing and adjustable transmission power can lead to reduced interference and energy consumption, improving network performance and lifetime.

N. Karagiorgas · C. Papageorgiou
Department of Computer Engineering and Informatics, University of Patras, Patras, Greece
e-mail: karagion@ceid.upatras.gr

C. Papageorgiou
e-mail: xpapageo@ceid.upatras.gr

P. Kokkinos (✉) · E. Varvarigos
Research Academic Computer Technology Institute, Patras, Greece
e-mail: kokkinop@ceid.upatras.gr

E. Varvarigos
e-mail: manos@ceid.upatras.gr

Keywords Ad hoc · Multi-cost · Power control · Energy · Interference

1 Introduction

An ad hoc network is a set of nodes that have the ability to communicate wirelessly without the existence of any fixed infrastructure. Nodes in an ad hoc network use other nodes as intermediate relays to transmit packets to their destinations. Since nodes are usually battery operated, energy conservation is an important issue. Furthermore, because of the broadcast nature of the wireless medium, ad hoc networks are also limited by interference/capacity considerations.

We distinguish between two routing approaches: the *single-cost* and the *multi-cost* approach. Most routing protocols proposed to date are based on the single-cost (shortest path) idea, where a single metric is used to represent the cost of using a link. This link metric can be a function of several network parameters (including load, energy and interference related parameters), but it is still a *scalar*. Routing algorithms of this kind calculate the path that has the minimum cost for each source-destination pair. Single-cost routing algorithms cannot optimize performance with respect to general cost functions, and they do not easily support Quality of Service (QoS) differentiation. Also, they usually yield only one path per source-destination pair, leading to non-uniform traffic distribution and possible instability problems [1].

In multi-cost routing, a *vector* of cost parameters is assigned to each link, and the cost vector of a path is defined based on the cost vectors of the links that comprise it. A set of candidate non-dominated paths is calculated for each source-destination pair, and an optimization function is used to choose the optimal among them. The cost parameters of interest to us are the hop count, the interference caused by a packet transmission, the residual energy of the nodes, and the transmission power used. We assume that nodes can adjust their transmission power to the minimum required level for coherent reception at the recipient node; in contrast, networks that use static transmission power consume more power than necessary, leading to energy squander and increased interference. As we will show, multi-cost routing makes application-specific routing possible, and permits the use of metrics that could not be considered in single-cost routing.

It is commonly accepted that in ad hoc networks there is a strong coupling among the performances of the different traditional layers of the ISO/OSI network model [1]. Decisions made at one layer affect decisions and performance at other layers, making multilayer optimization necessary. In our work we show how the combination of power control and multi-cost routing can help alleviate the energy and interference limitations of ad hoc networks. To this end, we propose and evaluate a number of energy- and interference-aware multi-cost routing algorithms that use the power adjustment capability of the nodes.

The context in which our energy- and interference-aware routing strategies are evaluated is that of the *evacuation problem* (Fig. 1). In this setting, the network starts with a certain number of packets that have to be routed and a certain amount of energy per node, and the objective is to serve the packets in the smallest number of steps, or to serve as many packets as possible before the energy at the nodes is depleted. We assume that the nodes are capable of dynamically adjusting their transmission power, so that they can communicate directly with any node they wish. Hence the network is fully connected, but, depending on the routing algorithm employed, a node may choose not to use the direct link to the destination, and use a multihop path instead. This assumption, which may not be valid in practice, is made so as to study the extent to which power control along with multi-cost routing can help optimize network performance. We are interested in measuring the average packet delay, the mean and the variance of the residual energy at the nodes after all data transfer has been completed, the average number of hops on the paths taken, the fraction of packets delivered to their destination, the frequency of packet collisions and the network throughput. Our simulation results show that the proposed multi-cost routing algorithms reduce interference and energy consumption, spread energy consumption more evenly across the network, and improve network performance and lifetime, when compared to more traditional algorithms that use a single cost criterion for making routing decisions.

The remainder of the paper is organized as follows. In Section 2 we report previous work. In Section 3 we present multi-cost routing for a general network. We also describe link cost parameters and optimization functions that are appropriate for the special case of wireless ad hoc networks. Section 4 outlines some of the differences between multi-cost and single-cost routing. In Section 5 we evaluate the performance of the multi-cost routing algorithms proposed under the network evacuation model. Finally, in Section 6 we present our conclusions.

2 Previous work

A great deal of work on wireless ad hoc networks has focused on the design of efficient routing protocols, where efficiency is interpreted using various performance criteria. Some works have designed routing protocols that exhibit small end-to-end delay, adaptiveness to the mobility of the nodes, or efficient use of the bandwidth or the energy. Most of these algorithms are single-cost, in the sense that they assign a scalar cost parameter to each link and compute the path that has minimum cost. Multipath single-cost routing has also been investigated, where a set of paths, instead of one optimal path, is found. Some works have also examined multi-constrained routing algorithms for ad hoc networks, by concentrating more on providing polynomial time multi-constrained heuristic routing algorithms and less on the cost parameters used and their effects on network performance. The present work differs from earlier works, by introducing multi-cost routing, as a generalization of both single-cost and multi-constrained routing, and using this approach to perform efficient energy-aware routing in ad hoc networks, using the nodes' power adjustment capabilities.

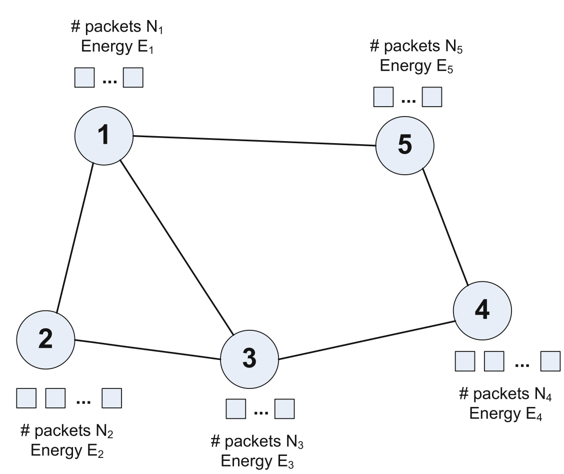

Fig. 1 The evacuation problem. Initially, each node i has energy E_i and holds N_i packets that have to be evacuated from the network

In [2–4] some well known routing algorithms for ad hoc networks are presented, where the emphasis is on finding ways to deal with mobility, and the metric optimized is the hop count or the delay. In [5], which is one of the first works proposing energy-aware routing for ad hoc networks, five energy-related metrics are investigated. In [6] a distributed protocol for finding the minimum power topology is presented. In [7] and [8] link costs are defined based on the energy expenditure for unit flow transmission and the initial and residual energy at the transmitting nodes. In [9] a new cost metric is used for routing, which is a function of the remaining battery level and the number of neighbours of a node. Other works have focused on the discovery of energy efficient routes under the constraint of a fixed end-to-end bit error rate, or by considering the expected number of retransmissions for reliable packet delivery [10, 11]. The selection of multiple energy efficient paths for a given source-destination pair has been proposed in [12].

Transmission power control for energy efficiency is also investigated in various works. In [13] two algorithms are proposed for selecting the node transmission power. In [7] the authors incorporates power control in the routing of packets, and try to increase energy consumption at nodes with plenty of energy, while reducing consumption at nodes with small energy reserves. In [14] power control is incorporated in the MAC layer by using the RTS-CTS-DATA-ACK sequence to reach an agreement on the transmission power to be used. In [15] and [16], joint power control, scheduling and routing algorithms are presented. In [17] the Slow Start MAC Protocol is proposed, where a slow start mechanism is used for the transmission of RTS/CTS and DATA packets, so as to save energy and decrease interference. Moreover, a number of devices exist capable of adjusting dynamically their transmission power, such as the Sun SPOT devices [18]. The Sun SPOT device is a small, wireless, battery powered experimental platform that includes a range of built-in sensors (e.g., temperature sensor) and it is developed by Sun.

A number of works [9, 19, 20] have studied the effects capacity and interference limitations have on the maximum achievable throughput of an ad hoc network, under a variety of assumptions on the network topology, the routing algorithm used, and the traffic pattern. In [21] the expected transmission count (ETX) metric is used, which incorporates the link loss ratios and the interference among successive links of a path. These metrics, however, ignore energy limitations, and tend to negatively impact network lifetime by overusing the energy reserves of a small set of nodes. Other works propose interference-aware routing using different definitions for the interference metric [22, 23].

The routing protocols mentioned above follow the single-cost approach, in the sense that they base their decisions on a single, scalar metric (which maybe a function of several metrics). Multi-constrained routing algorithms have also been investigated, especially for wired networks [24–28]. Finding paths subject to two or more cost parameters/constraints is in the general case an NP-complete problem [24, 29]. As a result, most algorithms proposed in this area concentrate on solving the Multi-Constrained Path (MCP) problem or the Multi-Constrained Optimal Path (MCOP) problem in a heuristic and approximate way with polynomial and pseudo-polynomial-time complexities, paying little attention to the parameters/costs used and their effects on network performance. The multi-constrained problem has been less studied in the context of wireless ad hoc networks, even though these networks have important reliability, energy, and capacity constraints that are not present in wired networks. In [30] the authors propose a probabilistic modeling of the link state for wireless sensor networks, and propose an approximation of a local multipath routing algorithm to provide soft-QoS under delay and reliability constraints. In [31] a multi-constrained QoS routing algorithm for mobile ad hoc networks is proposed that uses simulated annealing. In [32] the authors present an algorithm based on depth-first-search that solves the general k-constrained MCP problem with pseudo-polynomial time complexity. In [33] and [34] well-known routing algorithms for ad hoc networks, are extended to support QoS through the usage of multiple constraints. These algorithms focus on the bandwidth and delay constrained routing problem. In [35] a QoS routing scheme for ad hoc networks that uses flooding is proposed.

In the present work we examine how multi-cost routing, which is a generalization of multi-constrained routing, can be used to improve the performance of energy-and capacity/interference-constrained ad hoc networks. Despite the potential of multi-cost routing, the research activity in this field has been limited. The idea of multi-cost routing was presented in [36], where it was applied to wireline max-min fair share networks. Some of the ideas concerning multi-cost routing in wireless ad hoc networks and power control are scattered over our previous papers [37, 38] and are presented here coherently and more structured.

3 Interference/energy-aware multi-cost routing algorithms

3.1 Multi-cost routing

In multi-cost routing [36], each link of the network is assigned a cost vector consisting of several cost parameters. The cost vector of a path is obtained from the cost vectors of the links that comprise it by applying, componentwise, a monotonic associative operator to each cost vector parameter. The parameters that may be included in the path cost vector are categorized by the way they are obtained from the link cost vectors, that is, by the associative operator used for each cost vector component, and by the criterion applied to them (maximization or minimization) to select the optimal path. To be more specific, we denote by $v_l = (v_{1l}, v_{2l}, \ldots, v_{kl})$ the link cost vector of link l, by $V(P) = (V_1, V_2, \ldots, V_k)$ the cost vector of the path P that consists of links $l = 1, 2, \ldots, L$, and by $f(V)$ the optimization function that has to be minimized in order to select the optimal path. The cost vector $V(P) = (V_1, V_2, \ldots, V_k)$ of a path P consisting of links $l = 1, 2, \ldots, L$, is then obtained from the cost vectors of the links that comprise it by applying component-wise a monotonic associative operator $\odot$ to each cost vector parameter:

$$V_m = \odot_{l=1}^{L} v_{ml}$$

The associative operator may be different for different cost vector components. Generally, the parameters in the path cost vector, are categorized by the way they are obtained from the link cost vectors, that is by the associative operator used for each cost vector component, and by the criterion that is applied to them (maximization or minimization) to select the optimal path. For example, the m-th parameter of the cost vector may be of one of the following types:

– *additive*, where

$$V_m = \sum_{l=1}^{L} v_{ml}, \ \ v_{ml} \geq 0$$

and f is monotonically increasing in V_m (so our objective is to minimize V_m),

– *restrictive*, where

$$V_m = min_{l=1,\ldots,L}\{v_{ml}\}$$

and f is monotonically decreasing in V_m (so our objective is to maximize V_m), and

– *maximum representative*, where

$$V_m = max_{l=1,\ldots,L}\{v_{ml}\}$$

and f is monotonically increasing in V_m (so our objective is to minimize V_m).

A multi-cost routing algorithm [36] consists of two phases. In the first phase a set of candidate paths, called non-dominated paths, is obtained for a given source-destination pair using a multi-dimensional Dijkstra-like algorithm, executed at predefined intervals. Non-dominated are the paths for which it is impossible to find other paths that are better with respect to one of the cost criteria without being worse with respect to some other cost criterion. The algorithm that computes the non-dominated paths is a generalization of Dijkstra's algorithm with the basic difference that a set of non-dominated paths between the source and a destination node is obtained, instead of a single path. A detailed description of the algorithm is given in [36].

In the second phase of a multi-cost routing algorithm, when the origin node wishes to route a packet or a session to a given destination, a scalar cost function f is applied to the cost vectors of the non-dominated paths leading to that destination, and the path that gives the minimum cost is chosen. The optimization function f used depends on the QoS requirements of the session and may be different for different sessions. Note that the first phase reduces significantly the algorithm's total computational effort, since the optimization function does not need to be applied to every possible path for a given source-destination pair, but only to the set of non-dominated paths; this was proven in [36].

3.2 Cost parameters for ad hoc networks

The cost parameters used in the proposed interference/energy-aware multi-cost routing algorithms for ad hoc networks are the following.

– The number of hops h. The associative operator $\odot$ used in this case is the addition:

$$h = \sum_{l=1}^{j} h_l,$$

where $h_l = 1$ for all links l. Paths with a small number of hops are generally preferable to longer paths.

– The minimum residual energy R of a path. Here we use the residual energy R_l at the transmitting node of link l as the link cost metric. The minimum residual energy on the path is then obtained by applying the minimization (restrictive) operator to the link cost metrics to obtain:

$$R = \min_{l=1,\ldots,j} R_l.$$

The minimum residual energy R indicates the degree to which a path is energy-critical. Paths with large minimum residual energy are generally preferable.

– The sum T_1, or the maximum T_∞ of the transmission powers used by the nodes on a path. If we denote by T_l the transmission power required for correct reception over link l, then T_1 is obtained by combining the link metrics using the additive operator, while T_∞ is obtained by combining the link metrics using the maximization (maximum representative) operator:

$$T_1 = \sum_{l=1}^{j} T_l$$

or

$$T_\infty = \max_{l=1,\ldots,j} T_l.$$

Paths with small values for T_1 consume little total energy, and are therefore preferable. Similarly, paths with small values of T_∞ avoid energy-critical nodes and are also preferable.

– The total interference I_1, or the maximum interference I_∞ caused by using a path. As in [22], we define the interference I_l caused by using link l as the number of nodes (other than the transmitter and the receiver) that are within the transmission range of the end nodes of link l. If we denote the distance between the transmitter a and the receiver b of link $l = (a, b)$ as $|a, b|$, then:

$$I_l = I_{(a,b)} = |\{c \in V, |b,c| \le |a,b|\} \cup \{c \in V, |a,c| \le |a,b|\}| - 2$$

Of course one could suggest that some of the neighbours of nodes a and b may be inactive; however the metric assumes of the worst case scenario were all nodes are active and can transmit at any time. Also, in our metric we count the nodes that are within the transmission range of both the transmitter and the receiver and not only those of the transmitter. This is consistent with the RTS/CST mechanism of the 802.11 protocol we use, which informs and deters from transmission the neighbouring nodes of both ends of a link. Note that the interference metric is defined in our paper as an indicator of the number of other nodes that hear a transmission (and not as the interference power).

The total interference I_1 or the maximum interference I_∞ caused by using a path is obtained by employing the additive or the maximization (maximum representative) operator, respectively, for combining the interference metrics of the links on the path:

$$I_1 = \sum_{l=1}^{j} I_l$$

or

$$I_\infty = \max_{l=1,\ldots,j} I_l$$

Paths that create little total interference I_1 or little maximum interference I_∞ are generally preferable.

To clarify the notion of domination, consider the hop count h, the minimum residual energy R, the (total or maximum) transmission power T, and the (total or maximum) interference I, as the parameters of interest. Then a path p_1, with cost vector $V_1 = \{ h_1, R_1, T_1, I_1 \}$, is said to dominate a path p_2, with cost vector $V_2 = \{h_2, R_2, T_2, I_2\}$, when $h_1 < h_2$, $R_1 > R_2$, $T_1 < T_2$ and $I_1 < I_2$, In other words path p_1 dominates path p_2, if p_1 uses a smaller number of hops h, its nodes have larger energy reserves R, uses smaller transmission power T and causes less interference I. If p_1 is not dominated by any other path, then we say that p_1 is a non-dominated path.

Note that the parameters used in multi-cost routing do *not* have to be independent. The case where the parameters are independent is a special case, and it is not desirable to limit the algorithms to consider only this case. For example, when a node uses large transmission power the interference caused to its neighboring nodes will also tend to be large. However, the interference caused by a transmission also depends on the nodes' locations, and not only on the transmission power, and this is captured in the interference metric we use. Large transmission power also results in large energy expenditure by the transmitting nodes. Since both energy and interference are important, we include both the transmission power and interference metrics in the cost vectors and in the optimization functions. Also, the residual energy parameter only captures the current energy state of the network and as a result we also use the transmission power parameter as an indicator of the energy that will be consumed when using a candidate path.

The freedom that exists in the choice of the optimization function (Sect. 3.3) is used to obtain algorithms that give different emphasis on the different parameters (which do not have to be independent). The best choice of the optimization function is found through performance comparisons (Sect. 5). Even if the parameters were independent, this would not help us know a priori which optimization function would perform better.

3.3 Optimization functions: proposed algorithms

We combine the aforementioned cost parameters in different ways to produce various multi-cost routing algorithms. The following table contains the optimization functions examined, each corresponding to a different routing algorithm for selecting the paths. All of them select the path P with the minimum cost returned from the corresponding function. The optimization functions examined generally try to select paths that have a small number of hops, cause little interference, consume little energy and pass through nodes that have large residual energy, but they differ in the relative importance each of them gives to these metrics. We examined a number of different optimization functions (e.g., with and without square roots), the most interesting of which are the following:

– *Minimum Interference algorithm*: The criterion optimized is the sum of the interference of all the links on the path:

$$\min_P I_1(P).$$

This cost function actually results in a single-cost algorithm.

– *Minimum Transmission Power algorithm*: The criterion optimized is the sum of the transmission powers of the nodes on the path:

$$\min_P T_1(P).$$

This cost function also results in a single-cost algorithm.

– *SUM/MIN Energy-Interference algorithm*: The optimization that takes place is:

$$\min_P \frac{T_1(P) \cdot I_1(P)}{R(P)},$$

which tends to select paths that cause little total interference, use little total transmission power, and pass through nodes that have large residual energies.

– *SUM/MIN Energy-Half-Interference algorithm*: The function optimized is similar to the one used in the SUM/MIN Energy-Interference algorithm, but has a smaller dependence on the interference metric:

$$\min_P \frac{T_1(P) \cdot \sqrt{I_1(P)}}{R(P)}.$$

– *SUM/MIN Energy-Interference-Half Hop algorithm*: The optimization function is equal to the SUM/MIN Energy-Interference function, multiplied by $\sqrt{h(P)}$ so as to discourage, to a certain extent, the use of long paths:

$$\min_P \frac{\sqrt{h(P)} \cdot T_1(P) \cdot I_1(P)}{R(P)}$$

– *SUM/MIN Energy-Half-Interference-Half Hop algorithm*: The optimization function used is equal to that in the SUM/MIN Energy-Half-Interference algorithm, multiplied by $\sqrt{h(P)}$:

$$\min_P \frac{\sqrt{h(P)} \cdot T_1(P) \cdot \sqrt{I_1(P)}}{R(P)}.$$

– *MAX Interference algorithm*: The function optimized is the maximum of the interferences of the links on the path:

$$\min_P I_\infty(P).$$

– *MAX/MIN Energy-Half-Interference algorithm*: The optimization function is similar to that in the SUM/MIN Energy-Half-Interference algorithm, except that the transmission power and the interference are used as *maximum representative* instead of *additive* cost metrics:

$$\min_P \frac{T_\infty(P) \cdot \sqrt{I_\infty(P)}}{R(P)}.$$

– *MAX/MIN Energy-Half-Interference-Half Hop algorithm*: The optimization function is similar to that in the SUM/MIN Energy-Half-Interference-Half Hop algorithm, except that the transmission power and the interference are used as *maximum representative* instead of *additive* cost metrics:

$$\min_P \frac{\sqrt{h(P)} \cdot T_\infty(P) \cdot \sqrt{I_\infty(P)}}{R(P)}.$$

In what follows, we will refer to the functions SUM/MIN Energy- Interference and SUM/MIN Energy- Half-Interference as *Energy-Interference functions*, and to the corresponding routing algorithms as Energy-Interfe- rence algorithms. For the sake of brevity, we will also refer to the functions SUM/MIN Energy- Interference- Half- Hop and SUM/MIN Energy- Half- Interference- Half Hop as *Mixed functions*, and to the functions MAX/ MIN Energy- Half-Interference and MAX/MIN Energy- Half- Interference-Half Hop as *MAX/MIN functions*.

In general, the number of different non-dominated (and candidate) paths depends on the number of parameters in the cost vector, and on the type of operators used for calculating a path's cost vector from the constitutes links' cost vectors. In this context, it is possible that in some algorithms the number of non-dominated paths calculated is in the worst case exponential and it is not guaranteed that these algorithms run in polynomial-time. The cost parameters h, T_1 and I_1 are additive metrics, while the R, T_∞ and I_∞ are concave (restrictive or maximum representative). Based on [24], if the cost vector contains at most one additive metric (other than the hop count), then the algorithm is polynomial, independently of the number of the restrictive (that use the minimization operator) and maximum representative (that use the maximization operator) metrics. If the cost vector contains two or more additive metrics (other than the hop count), then the algorithm is exponential. The complexity considerations make some (polynomial) algorithms interesting even though they underperform some other (exponential) algorithms. As a result the SUM/MIN (Energy-Interference and Mixed) algorithms are exponential, while all the other algorithms (e.g., MAX/MIN) are polynomial. One of the main reasons we examine additive and concave metrics is that additive metrics usually result, as we will see, in somewhat better performance but longer algorithmic complexity, while the opposite is true for concave metrics. However, in practice, we found that the running times of the non-polynomial algorithms were also acceptable, at least for the network sizes used in the simulations.

In all cases, the algorithms first find the set of non-dominated paths with cost parameters (h, T, I, R), and then use the corresponding optimization function $f(h, T, I, R)$ to select the optimal path. In other words the computation of the set of non-dominated paths is common to all algorithms and the selection of the optimal path is done at the end in a way that is different for each of the algorithms proposed. The function to be optimized at the last step may depend on the QoS requirements of the user. The optimization functions considered penalize paths that use a large number of hops, or consume a large amount of energy, or pass through nodes that have little energy left, differentiating however from each other by giving different importance to each of these factors.

3.4 Power control

The ability of the nodes to adjust their transmission power levels is very important. With static transmission power case, where this ability is not present, a node may expend an unnecessarily large amount of energy and cause unwarranted interference to other nodes, when the desired recipient is at a smaller distance than the static transmission range used.

The ability of a node to adjust its transmission power results in interesting dilemmas and tradeoffs, which is basically what the multi-cost routing algorithms try to resolve. For example, in Fig. 2 node 1 can communicate directly with node 5, or it can use nodes 2, 3 and 4 as intermediate relays. In the first case the number of hops is equal to one, but the transmission power used is large. In the second case the number of the hops is four, but the total transmission power used is smaller than in the first case. A question that arises is which of the two approaches is better when all factors are taken into account. Another example indicating the benefits obtained from the flexibility provided by the variable transmission power is presented in Fig. 3. In case (a) nodes C and D cannot communicate, because they are both within the transmission floor reserved by the 802.11 MAC protocol for the communication of A with node B, while in case (b) node A adjusts its transmission power to the minimum required so that both pairs of nodes can communicate simultaneously.

Since nodes have the ability to control their transmission power, the topology of the network is not fixed and depends on the transmission powers chosen. During the execution of the algorithm and the candidate paths formulation phase, all the different combinations of nodes' transmission level and the resulted network to- pologies are evaluated. Due to the domination relations applied some of these combinations are discarded from further

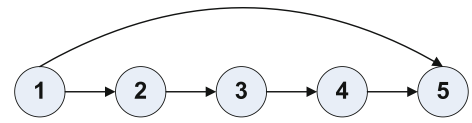

Fig. 2 First example of variable transmission power

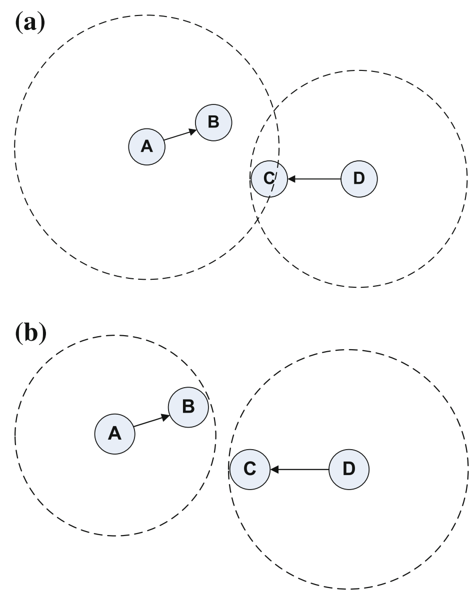

Fig. 3 Second example of variable transmission power

consideration, reducing this way the solution space. For example, in Fig. 2 our algorithm, during the path formulation phase, can find two paths connecting nodes 1 and 5: (1) (1, 5) and (2) (1, 2, 3, 4, 5), where node 1 uses two different transmission power levels T_1^{large} and T_1^{small} correspondingly. This way there are two cost vectors for paths (1) and (2):

$$(T_1^{large}, R_1, I_1)$$

and

$$(T_1^{small} + T_2 + T_3 + T_4, \min\{R_1, R_2, R_3, R_4\}, \max\{I_1, I_2, I_3, I_4\}),$$

assuming additive, minimization and maximization operators for the transmission power T, the residual energy R and the interference I parameters, correspondingly.

Next we can check the dominance relations of these paths. In case, no path dominates then we can either proceed with applying an optimization function in order to select one of them (and selecting the corresponding transmission power for node 1) or in case the topology of Fig. 2 is part of a larger topology (not presented in the figure), then we can continue extending these two sub-paths according to the multi-cost approach (and considering both transmission power levels for node 1). In case one of the paths dominates the other then the algorithm continues as described previously, using only the dominant path and the corresponding transmission power for node 1.

4 Multi-cost vs. single-cost routing

Multi-cost routing should not be confused with and is more general and powerful than single-cost routing. In single-cost routing one or more parameters (link characteristics) are combined in a single-cost metric characterizing the link. Multi-cost routing optimizes a function $f(V_1, V_2, \ldots, V_k)$, where $V_m = \odot_{l=1}^{L} v_{ml}$, which is different than optimizing $\sum_{l=1}^{L} f(v_{1l}, v_{2l}, \ldots, v_{kl})$, as single cost routing does. The paths discovered by multi-cost routing are optimal for the specific optimization function *f*, which is not always the case with single-cost routing, except for linear optimization functions. Also, multi-cost routing better captures the meaningfulness of each cost parameter by considering it for the whole path than for a single link. In the same context, in the multi-cost approach operators like minimization and maximization can be used, something that is not possible in the single-cost approach.

Moreover, multi-cost routing supports service differentiation for sessions with different QoS requirements, where each optimization function can be thought of as representing a different QoS class. For example, if both delay and energy are important for the session, then we may use the Energy-Hop algorithm, while if only energy is important, we may choose the Energy or the Energy-Half-Hop algorithms. Finding the optimal path for a different QoS class does not require the recalculation of the sets of non-dominated paths, but only the application of the corresponding optimization function on the sets of non-dominated paths already found. In the single-cost approach each path is characterized by a single scalar, which is the sum of the scalar costs that characterize each link of the path. In order to find the optimal path for a different QoS class, a different cost metric has to be applied to every link of the network and the shortest paths need to be recalculated.

For example, consider the SUM/MIN Energy-Hop multi-cost algorithm that uses the optimization function

$$f(h,T,R) = h(P)\cdot\frac{T_1(P)}{R(P)} = h\cdot\frac{\sum_{i\in P} T_i}{\min_{i\in P} R_i},$$

and compare it to the single-cost algorithm that uses the cost metric

$$h_i\cdot\frac{T_i}{R_i},$$

for link (i, j), where h_i equals 1. Having a single-cost algorithm corresponding to the above multi-cost optimization function is not possible, since the hop count per link h_i is 1 (and would therefore disappear from the link cost metric), the path residual energy $R(P)$ is obtained through a minimization operation (maximization and minimization cannot be captured by the single-cost approach) and division or multiplication cannot be combined with addition (used for example in the calculation of $T_1(P)$). Using the network of Fig. 4 it can be shown that the single-cost and multi-cost algorithms select different paths, for a given source-destination pair, even though they use the same parameters and in similar way.

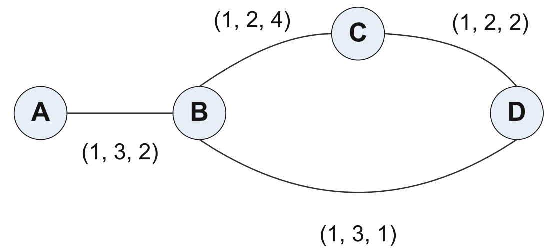

Fig. 4 Example network

When the multi-cost algorithm is applied in order to find a path between nodes A and B, then the (A, B, C, D) and (A, B, D) candidate paths (assuming that only these two are candidates) are calculated with the following cost vectors: (3, 3 + 2 + 2, min{2, 4, 2}) = (3, 7, 2) and (2, 3 + 3, min{2, 1}) = (2, 6, 1). It is is evident, that by using these two cost-vectors we can make more conscientious choices regarding which path to choose. If we are interested in minimizing the number of hops then we can use the second path, if we are interested in the path with the largest energy reserves then we can use the first path, or we can apply an optimization function (as the one mentioned previously) for finding a path that optimizes more than one parameters. On the other hand the single-cost approach (using the function mentioned above) would result in a single value/metric for each path: $1\cdot\frac{3}{2}+1\cdot\frac{2}{4}+1\cdot\frac{2}{2}=3$ and $1\cdot\frac{3}{2}+1\cdot\frac{3}{1}=4.5$. These values however do not give any insight about the characteristics of the paths. The only way would be to execute the single-cost algorithm many times, using different link cost metrics each time.

The example of Fig. 4 also shows that the inclusion property of single-cost routing does not hold for multi-cost routing. This property states that every subpath of an optimal (shortest) path is an optimal (shortest) path. In the network of Fig. 4 and the SUM/MIN Energy-Hop algorithm the optimal path for the A, B source-destination pair is path (A, B, C, D), while for the B, D source-destination pair it is path (B, D) and not (B, C, D); therefore, the inclusion property, which is true for single-cost routing, does not hold in general for the optimal paths produced by multi-cost routing. This also shows that the distributed implementation of multi-cost routing is not possible for some optimization functions, since intermediate nodes may not choose the paths originally intended by the source node.

5 Performance results

We evaluated the performance of the proposed energy-aware multi-cost routing algorithms under the network evacuation model. Under this model, the network starts with a certain number of packets that have to be routed and a certain amount of energy per node, and the objective is to serve the packets in the smallest number of steps, or to serve as many packets as possible before the energy at the nodes is depleted. We implemented the proposed algorithms and carried out corresponding experiments using the Network Simulator [39]. The routing agent running on each node calculates the set of non-dominated paths to all destinations at periodic time intervals. Generally, the routing process involves two levels: the routing information exchange level and the routing algorithm level. Routing information protocols deal with collecting and disseminating network state information, while routing algorithms compute the optimal-best path(s) using this information. Our focus is on the routing algorithm level and thus assume that each node has global knowledge of the network topology and all other information it needs for making routing decisions.

We assume that source routing is used, since, as discussed earlier, for some choices of the optimization function multi-cost routing is not amenable to distributed implementation (the inclusion property may not hold). When a data packet is generated at a node, the node applies the optimization function to the cost vectors of the corresponding non-dominated paths to select the optimal path, and the packet is sent on that path. If no route to the destination can be found, the packet is discarded.

The ad hoc network simulated consists of 16 stationary nodes randomly placed in a two-dimensional $350 \times 350\ \mathrm{m}^2$ area. The threshold of the received signal's power required for correct reception (that is, the receiver's sensitivity) is the same for all nodes. In our experiments the nodes are capable of dynamically adjusting their transmission power. Specifically, we assume that all nodes can communicate directly with each other and hence the network is fully connected, but, depending on the routing algorithm employed, a node may choose not to use the direct link to the destination, and use a multihop path instead. We assume nodes know the topology of the network and the physical distances between the nodes, so that they can adjust their transmission power to the minimum value needed for coherent reception at the receiving end, in order to consume only the minimum required energy and create minimal interference. In reality most of the existing devices with transmission power control capabilities, provide a finite set of possible power levels (e.g., the SunSpots [18]). Our assumption that the transmission power can take continuous values is mainly made for the simulations and is not required by the algorithms themselves. The minimum transmission power required for the communication between two nodes with distance d is calculated based on the following relation:

$$P_r(d) = \frac{P_t \cdot G_t \cdot G_r \cdot \lambda^2}{(4\pi)^2 \cdot d^a \cdot L},$$

where P_r is the power of the received signal, P_t is the power of the transmitted signal, G_t and G_r the gains of the senders' and receivers' antennas, respectively, $L \geq 1$ the system loss, and λ the wavelength used. In our calculations we assume $G_t = 1$, $G_r = 1$ and $L = 1$. The parameter a is the path loss constant, and is typically between 2 and 4 depending on the wireless channel. In our calculations we assume $a = 2$, corresponding to the Free Space propagation model.

It is obvious that the value calculated is optimal and both in our experiments and in reality does not always guarantees successful reception at the destination node. This is due to the interference limited environment, where successful reception is based on the signal to interference and noise ratio (SINR) and the fact that the line of sight communication ($a = 2$) between any pair of nodes in any network is usually rare. Also, packet collisions can also cause transmission failures. Alternatively to the procedure described, a protocol such as the Slow Start power control protocol of [17] can be used, to enable the transmitter and the receiver to agree on the transmission power to be used.

In our experiments, the number of packets that have to be evacuated varies from 100 to 1,000 (at steps of 100) packets per node. Packet destinations are taken to be uniformly distributed over all remaining nodes of the network. The packet sizes are fixed and equal to 500 bytes, and the transmission rate is equal to 0.1 packets/s. The energy of the nodes was taken to be either practically infinite (100 J) or finite (5 J), corresponding to a network with sufficient energy reserves and an energy-constrained network, respectively. The amount of energy expended for a packet transmission is equal to the transmission power multiplied by the duration of the packet transmission. A constant (independent of the distance) energy is also consumed for packet reception. When a node is idle we assume that it consumes no energy.

5.1 Performance metrics

We conducted a number of experiments to evaluate the performance of the proposed interference/energy-aware multi-cost routing algorithms. The performance measures of interest to us were:

- The average residual energy E remaining at the nodes at the end of each experiment, that is, when all packets have been evacuated from the network.

- The variance σ_E^2 of the node residual energies.
- The average number of hops h on the paths followed by the packets.
- The received-to-sent packets ratio, denoted by RS. Packets are dropped when a node runs out of energy before transmitting all the packets it was supposed to forward.
- The number of data packet collisions C due to the MAC protocol and the hidden terminal problem.
- The average packet delay D, defined as the average time that elapses between the beginning of an evacuation instance and the time a packet reaches its destination, averaged over all packets delivered to their destinations.
- The network throughput T, defined as the amount of information (in bytes) sent by the nodes during an evacuation interval, over the corresponding time duration.

The first two performance measures are related to energy considerations, while the remaining five are directly related to network performance. In our presentation of the results, we distinguish two cases: the case where the network has sufficient energy resources (nodes have practically infinite initial energy) and the case where the network is energy-constrained (nodes have finite initial energy).

5.2 Networks with sufficient energy (infinite initial energy)

In this section we present the performance results for the case where the nodes have very large initial energy, so that the network is not energy-constrained.

Figure 5 shows the average residual energy at the nodes when all packets have been evacuated from the network, for the multi-cost routing algorithms proposed in Sect. 3, while Fig. 6 illustrates the corresponding variance of the node

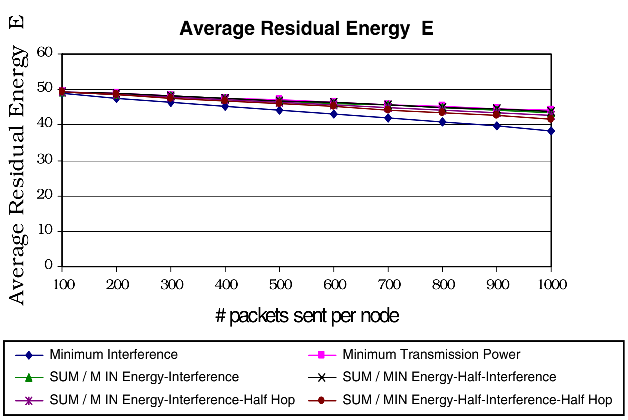

Fig. 5 Illustrates the average residual energy at the end of an evacuation period, as a function of the number of packets evacuated per node, for the case of infinite initial energy, and different choices of the routing algorithms

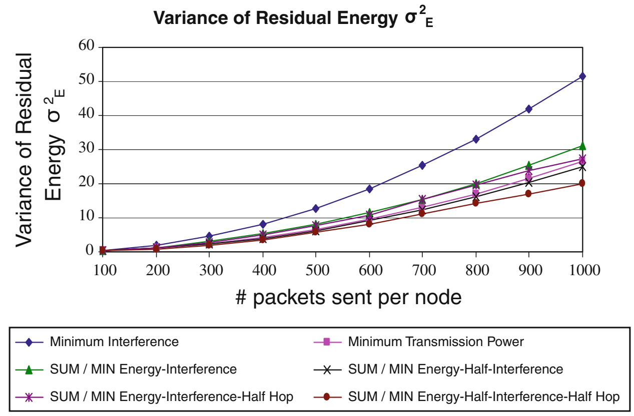

Fig. 6 Illustrates the variance of the node residual energies at the end of an evacuation period, as a function of the number of packets evacuated per node, for the case of infinite initial energy, and different choices of the routing algorithms

residual energies. As expected, the average residual energy decreases as the number of packets that are evacuated increases, since more packet transmissions lead to larger energy consumption. The Minimum Interference algorithm gives the lowest (worst) average node residual energy, since it does not take into account energy considerations. On the other hand the algorithms that incorporate energy related parameters, such as the transmission power or the node residual energies, perform better than the Minimum Interference algorithm. This is also true for the variance of the residual energies, illustrated in Fig. 6, where it can be seen that the Minimum Interference algorithm exhibits the highest variance, indicating that it does not uniformly spread energy consumption across the network. In contrast, algorithms that use the residual energy metric, which changes dynamically over time, distribute traffic and energy consumption more evenly across the network. The SUM/MIN Energy-Half-Interference-Half Hop algorithm seems to achieve the best results, because it considers all the

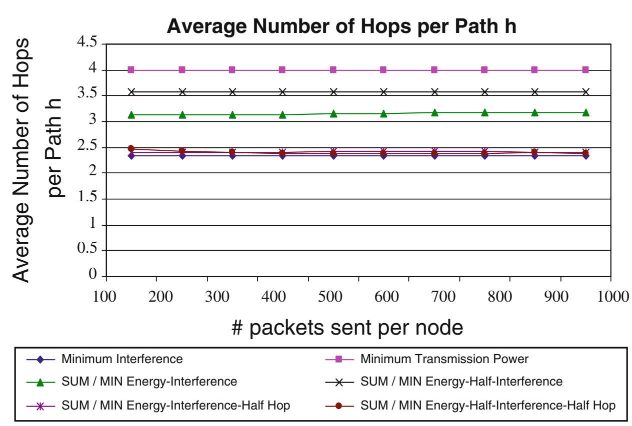

Fig. 7 Illustrates the average number of hops of the paths taken by the routing algorithms examined, as a function of the number of packets evacuated per node, for the case of infinite initial energy

parameters, including the number of hops, in its cost function. We also observe in Fig. 6 that the variance of the residual energy increases with the number of packets that are evacuated, since more packet transmissions intensify the energy variances between the nodes.

Figure 7 illustrates the average number of hops of the paths followed by the packets for the algorithms examined. We observe that the Minimum Transmission Power algorithm results in paths with the largest number of hops, while the Minimum Interference algorithm uses paths with the smallest number of hops. This is because the Minimum Transmission Power algorithm selects paths consisting of many short links to minimize the total transmission power used, while the Minimum Interference algorithm chooses paths with a small number of hops to keep the interference at low levels. The Mixed algorithms results are very close to those obtained by the Minimum Interference algorithm, since they include the interference and the hop metric in their cost functions. On the other hand the Energy-Interference algorithms achieve larger average number of hops, since they include only the interference metric in their cost functions. As expected, since we assume infinite energy per node, the path lengths do not change with the number of packets evacuated from the network.

Regarding the frequency of collisions shown in Fig. 8, the Minimum Interference algorithm exhibits the best results. This was expected since our definition of the interference metric aims at minimizing the number of nodes that are within the transmission range of the transmitter or the receiver, thus minimizing the probability of a packet collision. The worst results regarding this metric were produced by the Minimum Transmission Power algorithm. The Minimum Interference algorithm also outperforms the other algorithms examined with respect to the throughput and the average packet delay, as it can be seen

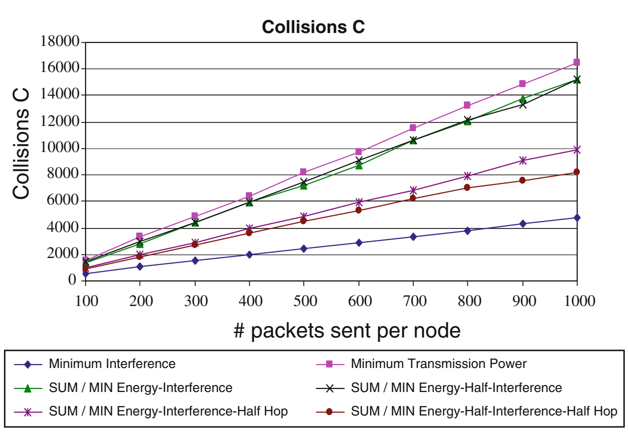

Fig. 8 Illustrates the number of collisions over an evacuation period, as a function of the number of packets evacuated per node, for the case of infinite initial energy, and different choices of the routing algorithms

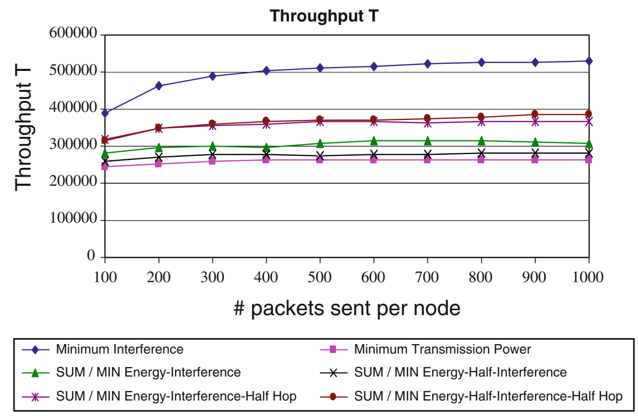

Fig. 9 Illustrates the throughput obtained during the execution of an evacuation problem, as a function of the number of packets evacuated per node, for the case of infinite initial energy, and different choices of the routing algorithms

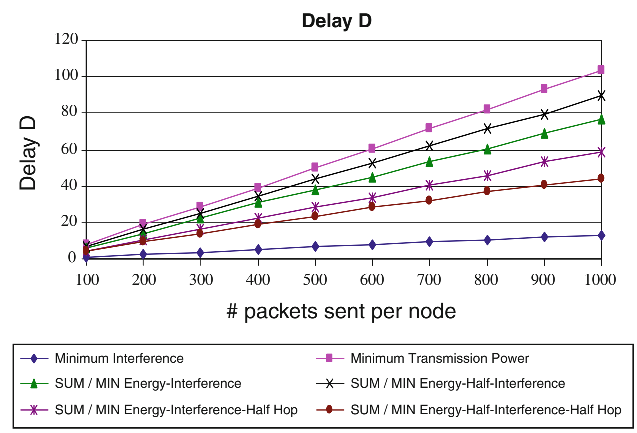

Fig. 10 Illustrates the average packet delay taken over an evacuation period, as a function of the number of packets evacuated per node, for the case of infinite initial energy, and different choices of the routing algorithms

in Figs. 9 and 10, respectively. The fewer packet collisions caused by this algorithm produce proportionally fewer packet retransmissions, higher throughput and smaller delay. On the other hand the Minimum Transmission Power algorithm gives the worst results, since it does not consider the interference caused. Moreover, the performance of the remaining algorithms lies between that of these two algorithms, with the algorithms considering more metrics (Mixed) producing better results than those considering fewer metrics (Energy-Interference).

The results presented so far for the case of infinite initial power suggest that the Minimum Interference algorithm outperforms the other algorithms with respect to most network performance measures of interest. This was expected in the particular case considered in this section, where nodes have sufficient energy, and the use of energy related cost metrics cannot improve routing decisions. The situation is different, however, when the network is energy-constrained, as the results in the following section indicate.

5.3 Energy-constrained networks (finite initial energy)

In this section we present the performance results for the case where the nodes have finite initial energy, so that the network is energy-constrained.

Figures 11 and 12 illustrate the average residual energy at the end of an evacuation period, and the variance of the residual energies, respectively, as a function of the number of packets that are evacuated per node. The differences between the algorithms examined with respect to these two energy related measures are more pronounced here than they were in the corresponding graphs of Figs. 5 and 6, which were obtained under the infinite initial energy scenario. This occurs because in the finite energy scenario many nodes energy is depleted. We observe that the Minimum Transmission Power algorithm outperforms the other algorithms examined with respect to the average residual energy, since it selects paths consisting of many short links (chooses long paths) in order to minimize the total transmission power used. This way it also minimizes the energy consumption at the nodes. Similarly, the Mixed algorithms that consider the hop metric (choose short paths), behave worse than the Energy-Interference algorithms that do not consider this metric. The Minimum Interference algorithm on the other hand exhibits the worst performance, since it does not consider any energy related metric.

Regarding the variance of the residual energy, shown in Fig. 12, we observe that the results are quite different than those in the infinite initial energy scenario. For almost all the algorithms examined the variance initially increases, but then starts decreasing rapidly as the number of packets that are evacuated increases. This happens because as more

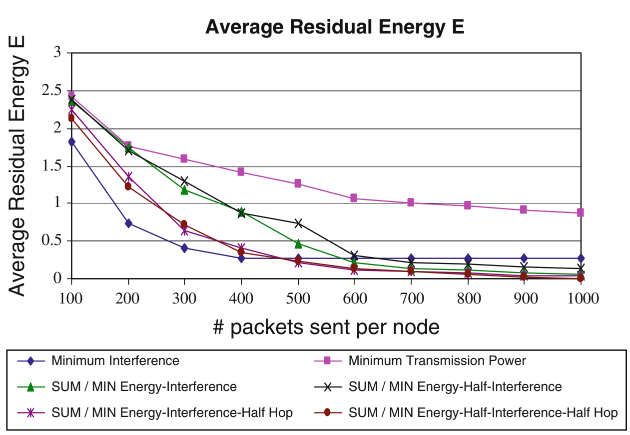

Fig. 11 Illustrates the average residual energy at the end of the evacuation problem, as a function of the number of packets evacuated per node, for the case of finite initial energy, and different choices of the routing algorithms

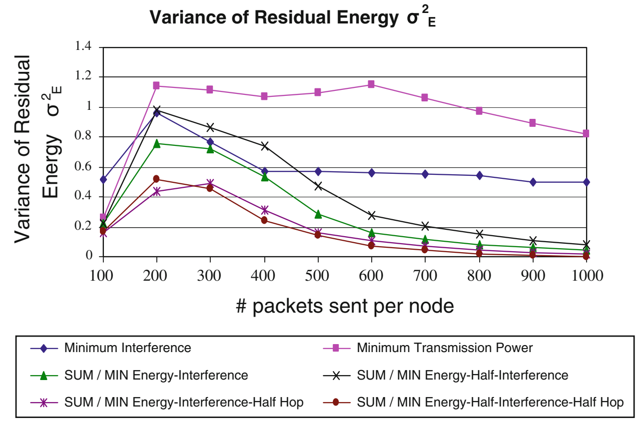

Fig. 12 Illustrates the variance of the residual energy at the end of the evacuation problem, as a function of the number of packets evacuated per node, for the case of finite initial energy, and different choices of the routing algorithms

packets are evacuated, then more nodes energy is depleted, resulting in small residual energy variance. Also, in the experiments of this section, energy is a critical resource and the algorithms that incorporate the path residual energy R in their cost functions (that is, the Mixed and Energy-Interference algorithms) perform better, achieving small variance in the residual energy. This is because the paths these algorithms select are adjusted over time to reflect changes in the energy reserves of the nodes, resulting in a more even distribution of energy consumption across the network. In contrast, the Minimum Interference and the Minimum Transmission Power algorithms do not change their paths when nodes start running out of energy, resulting in earlier depletion of the energy at some nodes, while there are other nodes that still have significant energy reserves.

Figure 13 illustrates the received-to-sent packets ratio RS as a function of the number of packets that are evacuated. The Energy-Interference and Mixed algorithms achieve better RS ratio than the Minimum-Interference and the Minimum-Transmission power algorithms, since they result in a longer network lifetime and, consequently, more packets are delivered to their destinations.

Regarding the average number of hops per path illustrated in Fig. 14, the Minimum Transmission Power algorithm uses paths that have the largest number of hops, for the reasons already discussed for the case of infinite initial energy. The paths selected by the Mixed algorithms, which consider the hop metric, again consist of fewer hops than those selected by the Energy-Interference algorithms. A rather counter-intuitive observation is that the number of hops on the paths decrease when the number of packets that are evacuated increases. This can be explained by considering the depletion of the energy at some of the nodes as more packets are evacuated, which results in the selection of paths with fewer but longer links. For example, in the

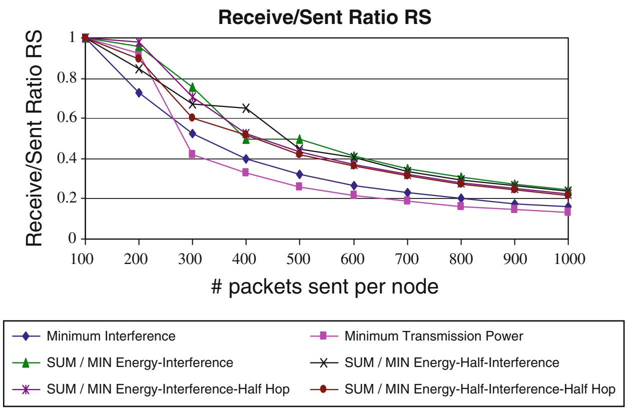

Fig. 13 Illustrates the received to sent ratio at the end of an evacuation problem, as a function of the number of packets evacuated per node, for the case of finite initial energy, and different choices of the routing algorithms

network of Fig. 2 where node 1 wants to communicate with node 5, the Minimum Transmission Power algorithm would initially select the route 1→2→3→4→5. If, however, node 4 exhausts at some time its energy reserves, node 3 will adjust its transmission power, so that the new routing path becomes 1→2→3→5, which contains fewer hops than the original path.

Regarding the number of collisions, illustrated in Fig. 15, the results are similar to those obtained for the infinite initial energy case. However, we should point out that the Mixed algorithms exhibit better behavior in this case, yielding a collision frequency that is very close to that of the Minimum Interference algorithm.

Figures 16 and 17 illustrate the average throughput achieved during an evacuation period, and the average packet delay, respectively, as a function of the number of packets that are evacuated. The results are similar to those obtained under the infinite initial energy model. It is worth noting, however, that the Mixed and the Energy-Interference algorithms achieve better throughput than the

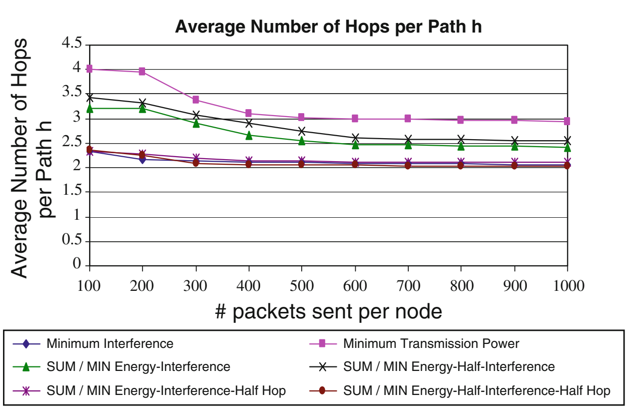

Fig. 14 Illustrates the average number of hops of the paths taken by the routing algorithms examined, as a function of the number of packets evacuated per node, for the case of finite initial energy

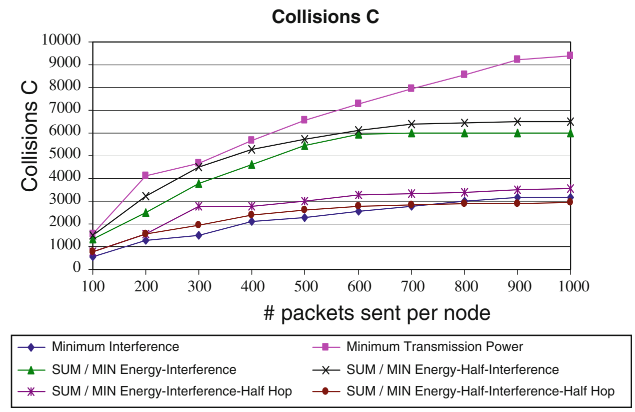

Fig. 15 Illustrates the number of collisions at the end of an evacuation period, as a function of the number of packets evacuated per node, for the case of finite initial energy, and different choices of the routing algorithms

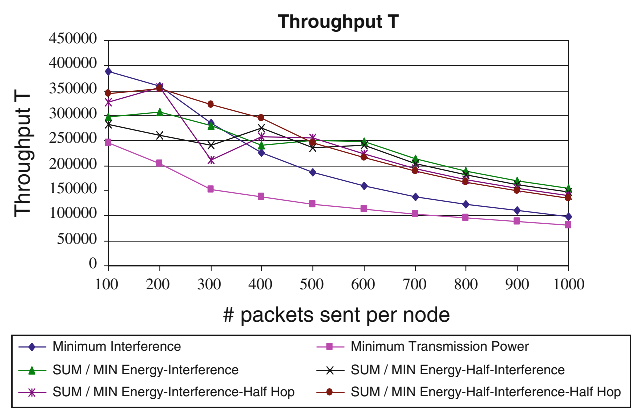

Fig. 16 Illustrates the average throughput over an evacuation period, as a function of the number of packets evacuated per node, for the case of finite initial energy, and different choices of the routing algorithms

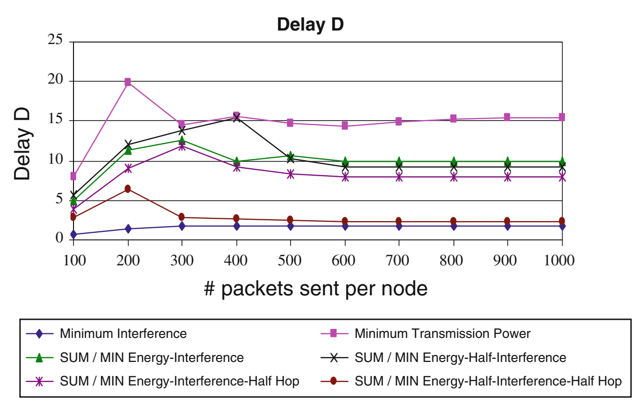

Fig. 17 Illustrates the average packet delay over an evacuation period, as a function of the number of packets evacuated per node, for the case of finite initial energy, and different choices of the routing algorithms

Minimum Interference algorithm, which was not the case under the infinite initial energy model. This is because these algorithms use energy more efficiently, by taking into account the node residual energies, and thus manage to deliver more packets to their destinations before the energy is depleted.

5.4 Performance of MAX/MIN algorithms

In this section we present the performance results for the MAX/MIN algorithms, which were obtained using the same network topology and parameters as in previous section. As already mentioned, the complexity of any optimization function using at least two additive metrics is exponential, except for the case where one of the two metrics is the hop count. Also, when one additive and one concave (restrictive or maximum representative) metric is used then the complexity of the corresponding optimization function is polynomial. As a result the MAX/MIN algorithms are polynomial, while the SUM/MIN algorithms examined earlier are exponential. However, in practice, we found that the running times of all the algorithms were acceptable, at least for the network sizes used in the simulations.

Indicatively, we present the figures concerning the variance of the residual energy (Fig. 18) and the received-to-sent packets ratio (Fig. 19) in the case of finite initial energy, which is the most realistic and significant case. We compare the MAX/MIN Energy- Half- Interference and the MAX/MIN Energy- Half- Interference- Half Hop algorithms with the corresponding SUM/MIN algorithms (Mixed). We observe that in all occasions the SUM/MIN algorithms are better than the corresponding MAX/MIN algorithms. In other words, the T_1 and I_1 metrics, are more appropriate than the T_∞ and I_∞ metrics, respectively, in making routing decisions: the summing up of the values of the transmission

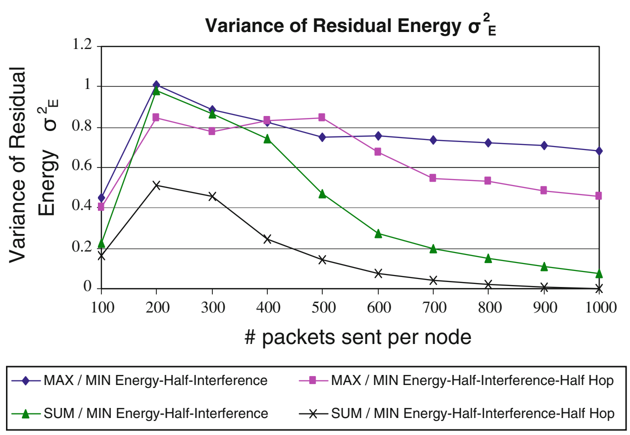

Fig. 18 Illustrates the variance of the residual energy at the end of an evacuation problem, as a function of the number of packets evacuated per node, for the case of finite initial energy, and different choices of the routing algorithms

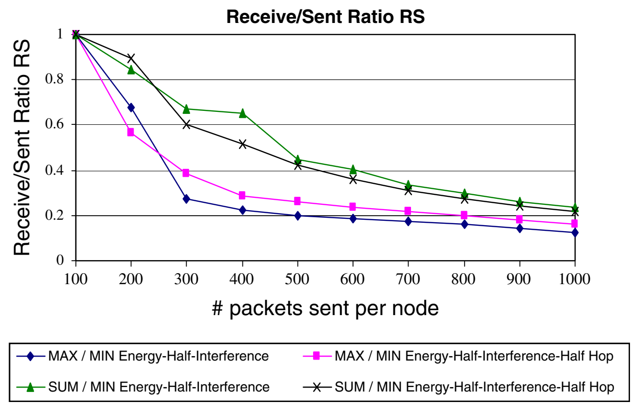

Fig. 19 Illustrates the received to sent ratio at the end of an evacuation proble m, as a function of the number of packets evacuated per node, for the case of finite initial energy

powers of the nodes on a path and the interferences on this path seems to be a more representative metric of the cost of using this path, than taking their maximum value. Note, however, that if both the T_1 and I_1 metrics are used, the algorithm (SUM/MIN) has exponential complexity, while when one them is replaced by the T_∞ and I_∞ metrics, respectively, the complexity becomes polynomial.

5.5 Comparing the SUM/MIN Energy-Interference with a single-cost algorithm

We performed a number of experiments comparing a multi-cost algorithm, namely the SUM/MIN Energy-Interference, with the corresponding single-cost algorithm. In this single-cost algorithm the cost of each link is equal to $\frac{T \cdot I}{R}$, where T is the link's start node transmission power, I is the interference value/metric of the link and R is the residual energy of the link's end node. Figure 20a illustrates the received to sent ratio at the end of an evacuation problem for the case of finite initial energy. We observe that by using the multi-cost SUM/MIN Energy-Interference more packets are successfully delivered to their destination. This is because the multi-cost algorithm makes better use of the nodes' energy reserves and this is also confirmed by Fig. 20b. In particular, Fig. 20b illustrates the time (measured in simulation time units) when nodes run out of energy. We observe that by using the SUM/MIN Energy-Interference algorithm nodes start running out of energy later in time than when using the corresponding single-cost algorithm. Also, we observe that the time period when this happens is between time instances 10 and 70, while in the case of the single-cost algorithm this happens between time instances 10 and 120. This indicates that the multi-cost algorithm tends to spread the energy consumption uniformly in the network, so that when one node is at the point of first running out of energy, most other nodes are at the same energy-critical situation. Also, we should note that after time instance 60, where 10 nodes have zero energy for both algorithms, the network can be considered as disconnected. The nodes after this period, consume energy by trying, unsuccessfully, to transmit new packets.

The main reason the multi-cost SUM/MIN Energy-Interference algorithm behaves better than the corresponding single-cost one, is that it better captures the meaningfulness of each cost parameter by considering it for the whole path than for a single link. However, as we already stated in Sect. 4, the multi-cost approach has a number of benefits that cannot be (at least directly—quantitatively) measured. One of these benefits is that it is possible to run the algorithm once and then select paths using multiple different optimization criteria, based on the QoS requirements. For example, a node can route new packets (probably belonging to different applications) using the SUM/MIN Energy-Interference and (at the same time) the SUM/MIN Energy-Half-Interference optimization functions. To have similar results using the single-cost approach one would need to re-execute the single-cost algorithm many times, using different link cost metrics each time. Also, we should note that in the evaluated single-cost algorithm there is no way to account for the hop count of the selected path. On the other hand in the corresponding multi-cost algorithm, this is indirectly, at least, considered through the use of the transmission power parameter. Also, it can be directly accounted for by including the hop parameter in the link's cost vector (as is the case in the SUM/MIN Energy-Interference-Hop algorithm).

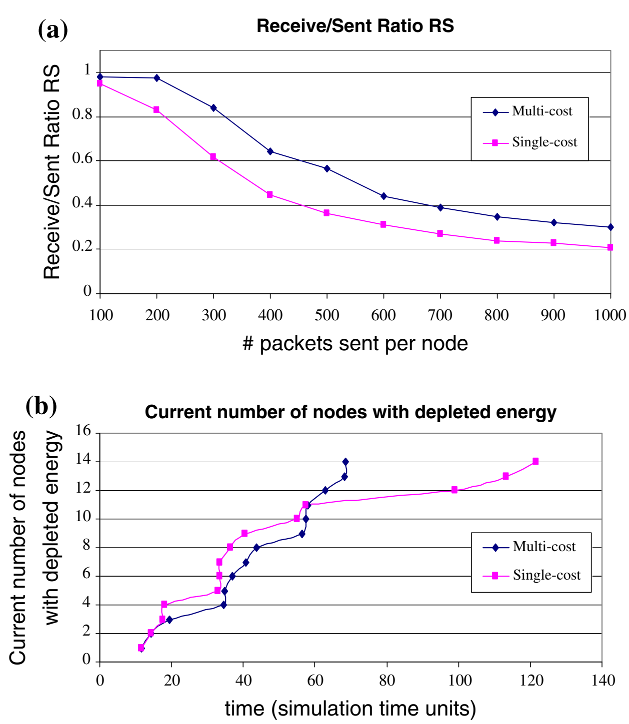

Fig. 20 Illustrates **a** the received to sent ratio at the end of an evacuation problem as a function of the number of packets evacuated per node and **b** the current number of nodes with depleted energy for the multi-cost SUM/MIN Energy-Interference and the corresponding single-cost algorithms, for the case of finite initial energy

6 Conclusions

In our work we showed that the combination of multi-cost routing and variable transmission power can optimize energy constrained ad hoc networks performance, by obtaining dynamic energy- and interference-efficient routes. The multi-cost routing problem is a generalization of the multi-constrained problem, where no constraints exist, and is also significantly more powerful than single-cost routing. We proposed and evaluated a number of energy- and interference-aware multi-cost routing algorithms that use the power adjustment capability of the nodes. A number of cost parameters, including hop count, interference caused, residual energy of the nodes and transmission power, were considered and their impact was evaluated.

A large number of experiments were conducted assuming infinite and finite node initial energy reserves. We evaluated the algorithms proposed using both performance and energy related measures. In the infinite energy scenario the single-cost algorithms, namely the Minimum Interference and the Minimum Transmission Power algorithms, produced the best and the worse results, while the multi-cost algorithms produced intermediate results. On the other hand in the finite energy scenario, where the ad hoc network is energy constrained and the energy of many nodes gets eventually depleted, the multi-cost algorithms were superior, exhibiting larger throughput, reduced interference and energy consumption, and increased network lifetime. The SUM/MIN Energy-Half-Interference-Half-Hop multi-cost algorithm, which incorporates hop count, interference and energy related metrics, outperforms, in most cases, all the other algorithms considered. Finally, we showed quantitatively that the multi-cost SUM/MIN Energy-Interference algorithm performs better than the corresponding single-cost algorithm.

Acknowledgments C. Papageorgiou was supported by GSRT through PENED project 03EΔ207, funded 75% by the EC and 25% by the Greek State and the private sector.

References

1. Ephremides, A. (2002). Energy concerns in wireless networks. *IEEE Wireless Communications, 9*(4), 48–59.
2. Perkins C., & Bhagwat, P. (1994). Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers.

ACM SIGCOMM conference on communications architectures, protocols and applications, pp. 234–244.
3. Perkins, C. (1997). *Ad hoc on demand distance vector (aodv) routing*, Internet Draft, draft-ietf-manet-aodv-00.txt.
4. Johnson, D. B., & Maltz, D. A. (1996). *Dynamic source routing in ad hoc wireless networks, mobile computing* (Vol. 353), Dordrecht: Kluwer.
5. Singh, S., Woo, M., & Raghavendra, C. (1998). *Power-aware routing in mobile ad hoc networks*, ACM international conference on mobile computing and networking.
6. Rodoplu, V., & Meng, T. (1999). Minimum energy mobile wireless networks. *IEEE Journal of Selected Areas in Communications, 17*, 1333–1344.
7. Chang, J.-H., & Tassiulas, L. (2000). *Energy conserving routing in wireless ad hoc networks* (pp. 22–31). USA: INFOCOM.
8. Chang, J.-H., & Tassiulas, L. (2004). Maximum lifetime routing in wireless sensor networks. *IEEE/ACM Transactions on Networking, 12*(4), 609–619.
9. Gupta, N., & Das, S. R. (2002). *Energy-aware on-demand routing for mobile ad hoc networks*. International workshop on distributed computing, mobile and wireless computing, (pp. 164–173). New York: Springer.
10. Banerjee, S., & Misra, A. (2002). *Minimum energy paths for reliable communication in multi-hop wireless networks*, ACM international symposium on mobile ad hoc networking and computing.
11. Manohar, R., & Scaglione, A. (2003). *Power optimal routing in wireless networks*, IEEE international conference on communications.
12. Shah, R., & Rabaey, J. (2002). *Energy aware routing for low energy ad hoc sensor networks*, IEEE wireless communications and networking conference.
13. Ramanathan, R., & Rosales-Hain, R. (2000). *Topology control of multihop wireless networks using transmit power adjustment*. IEEE INFOCOM.
14. Agarwal, S., Krishnamurthy, S. V., Katz R. H., & Dao, S. K. (2001). *Distributed power control in ad hoc wireless networks*, IEEE international symposium on Personal, Indoor and Mobile Radio Communications (PIMRC).
15. Elbatt, T., & Ephremides, A. (2002). *Joint scheduling and power control for wireless ad hoc networks*. IEEE INFOCOM.
16. Li, Y., & Ephremides, A. (2007). Joint scheduling, power control, and routing algorithm for ad hoc wireless networks. *Ad Hoc Networks, 5*(7), 959–973.
17. Gkamas, V. S., & Varvarigos, E. A. (2006). A slow start power control MAC protocol for mobile ad hoc networks. *IEEE international symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*.
18. http://www.sunspotworld.com/.
19. Li, J., Blake, C., De Couto, D. S. J., Lee, Hu Imm, & Morris, R. (2001). *Capacity of ad hoc wireless networks*. ACM international conference on mobile computing and networking (pp. 61–69).
20. Gupta, P., & Kumar, P. R. (2000). Capacity of wireless networks. *IEEE Transactions on Information Theory, 46*(2), 388–401.
21. De Couto, D., Aguayo, D., Bicket, J., & Morris, R. (2003). *A high-throughput path metric for multi-hop wireless routing*, IEEE international conference on mobile computing and networking.
22. Burkhart, M., Von Rickenbach, P., Wattenhofer, R., & Zollinger, A. (2004). *Does topology control reduce interference?*, ACM international symposium on mobile ad hoc networking and computing (pp. 9–19).
23. Meyer auf der Heide, F., Schindelhauer, C., Volbert, K., & Grunewal, M. (2002). *Energy, congestion and dilation in radio networks*, ACM symposium on parallel algorithms and architectures.
24. Wang, Z., & Crowcroft, J. (1996). Quality-of-Service routing for supporting multimedia applications. *Journal of Selected Areas in Communications, 14*(7), 1228–1234.
25. Yuan, X., & Liu, X. (2002). Heuristic algorithms for multiconstrained quality-of-service routing. *IEEE/ACM Transactions on Networking, 10*(2), 244–256.
26. Chen, S., & Nahrstedt, K. (1998). *On finding multi-constrained paths*, IEEE international conference on communications.
27. Korkmaz, T., & Krunz, M. (2001). *Multi-constrained optimal path selection*. USA: INFOCOM.
28. Mieghem, P., & Kuipers, F. (2004). Concepts of exact QoS routing algorithms. *IEEE/ACM Transactions on Networking, 12*(5), 851–864.
29. Mieghem, P., & Kuipers, F. (2003). On the complexity of QoS routing. *Computer Communications, 26*(4), 376–387.
30. Huang, X., & Fang, Y. (2006). *Multi-constrained soft-QoS provisioning in wireless sensor networks*. Canada: QSHINE.
31. Liu, L., & Feng, G. (2007). Simulated annealing based multi-constrained QoS routing in mobile ad hoc networks. *International Journal of Wireless Personal Communications, 41*(3), 393–405.
32. Li, Z., & Garcia-Luna-Aceves, J. (2007). Finding multi-constrained feasible paths by using depth-first search. *Wireless Networks, 13*(3), 323–334.
33. Badis, H., & Al Agha, K. (2006). *Quality of Service for Ad hoc Optimized Link State Routing Protocol (QOLSR)*, draft-badis-manet-qolsr-03.txt.
34. Perkins, C., Royer, E., & Das, S. (2000). *Quality of service in ad hoc on-demand distance vector routing*, draft-perkins-manet-aodvqos-00.txt.
35. Chen, S., & Nahrstedt, K. (1999). Distributed quality-of-service routing in ad hoc networks. *Journal on Selected Areas in Communications, 17*(8), 1488–1505.
36. Gutierrez, F. J., Varvarigos, E., & Vassiliadis, S. (2000) *Multicost routing in max-min fair share networks*, Annual Allerton conference on communication, control and computing (pp. 1294–1304).
37. Papageorgiou, C., Kokkinos, P., & Varvarigos, E. (2006). *Multicost routing over an infinite time horizon in energy and capacity constrained wireless ad hoc networks*. Euro-Par.
38. Karagiorgas, N., Kokkinos, P., Papageorgiou, C., & Varvarigos, E. (2007). *Routing in wireless ad hoc networks with variable transmission power*, IEEE international symposium on Personal, Indoor and Mobile Radio Communications (PIMRC).
39. *Ns—network simulator*, http://www.isi.edu/nsnam/.

Author Biographies

Nikolaos Karagiorgas received a Diploma in Computer Engineering and Informatics in 2003 and an MS degree in Hardware and Software Integrated Systems in 2006, both from the University of Patras. His graduate work was on multicost routing in ad hoc networks with variable transmission power. He is a recipient of the Latsis Benefit Foundation scholarship. His scientific interests are in the areas of operating systems, computer networks, communications protocols, and algorithms.

Panagiotis Kokkinos received a Diploma in Computer Engineering and Informatics in 2003 and an MS degree in Integrated Software and Hardware Systems in 2006, both from the University of Patras, Greece. He has worked in the private sector. He is currently a Phd student at the Department of Computer Engineering and Informatics of the University of Patras. His research activities are in the areas of ad hoc networks and grid computing.

Christos Papageorgiou received his PhD and MSc from the Computer Engineering and Informatics Department, University of Patras. Currently, he works as a researcher in the field of wireless communications. His scientific interests are in the areas of computer networks, communications protocols, and algorithms.

Emmanouel Varvarigos received a Diploma in Electrical and Computer Engineering from the National Technical University of Athens in 1988, and the M.S. and Ph.D. degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology in 1990 and 1992, respectively. He has held faculty positions at the University of California, Santa Barbara (1992-1998, as an Assistant and later an Associate Professor) and Delft University of Technology, the Netherlands (1998-2000, as an Associate Professor). In 2000 he became a Professor at the department of Computer Engineering and Informatics at the University of Patras, Greece, where he heads the Communication Networks Lab. He is also the Director of the Network Technologies Sector (NTS) at the Research Academic Computer Technology Institute (RA-CTI), which through its involvement in pioneering research and development projects, has a major role in the development of network technologies and telematic services in Greece. Professor Varvarigos has served in the organizing and program committees of several international conferences, primarily in the networking area, and in national committees. He has also worked as a researcher at Bell Communications Research, and has consulted with several companies in the US and in Europe. His research activities are in the areas of protocols for high-speed networks, ad hoc networks, network services, parallel and distributed computation and grid computing.