

PCQoS: power controlled QoS tuning for wireless ad hoc networks

Javier Gomez · Luis A. Mendez · Victor Rangel ·
Andrew T. Campbell

© Springer Science+Business Media, LLC 2010

Abstract Mobile ad hoc networks typically use a common transmission power approach for the discovery of routes and the transmission of data packets. In this paper we present PCQoS; a power-controlled Quality of Service (QoS) scheme for wireless ad hoc networks which builds QoS mechanisms for specific applications that wish to trade-off better QoS performance for sub-optimal paths. PCQoS allows selected flows to modify their transmit power as a way to add and remove relay nodes from their paths in order to coarsely modify their observed application QoS performance. We present simulation results and show that PCQoS can be used to provide coarse control over traditional QoS metrics (e.g., delay, throughput). To the best of our knowledge the PCQoS protocol represents the first attempt to use variable-range transmission control as a means to provide QoS differentiation to applications in wireless ad hoc networks.

Keywords Ad hoc networks · Power control · WLAN · QoS

J. Gomez (✉) · L.A. Mendez · V. Rangel
Department of Electrical Engineering, National University of Mexico, Ciudad Universitaria, C.P. 04510, D.F., Mexico City, Mexico
e-mail: javierg@fi-b.unam.mx

L.A. Mendez
e-mail: alonso@fi-b.unam.mx

V. Rangel
e-mail: victor@fi-b.unam.mx

A.T. Campbell
Department of Computer Science, Dartmouth College, Hanover, NH 03755, USA
e-mail: campbell@cs.dartmouth.edu

1 Introduction

The impact of transmission power control on network throughput has been widely studied in the literature in the context of cellular networks [12, 25], and more recently in the case of Wireless LAN (WLAN) and wireless ad hoc networks [2, 13]. The later analysis focuses on the maximum capacity of the network as a function of the transmission range, node density, and average distance between source-destination pairs. In [13] the authors show that the end-to-end throughput available to each node is $O(\frac{1}{\sqrt{n}})$ for random traffic pattern where n is the number of nodes.

One of the main QoS trade-offs involved in a wireless ad hoc network is related to the average number of times a packet is forwarded versus the average number of interfering nodes per attempted transmission. Increasing the transmission range reduces the number of times a packet needs to be forwarded by intermediate nodes to its final destination. However, increasing the transmission range increases channel contention every time a node attempts to transmit, thus, increasing transmission delays. An inverse trade-off applies when the transmission range is reduced. In [13] it is shown that reducing the transmission range is a better solution in terms of increasing the traffic carrying capacity of wireless ad hoc networks. The analysis presented in [13] and [28] only considers the physical capacity of the network, and not, the inefficiency of the MAC protocol used to transport data on top of the physical network. Unfortunately, MAC protocols used in wireless ad hoc networks provide only limited performance in particular those protocols developed for shared medium access like CSMA [27].

It is widely known that IEEE 802.11 is not the best MAC protocol for multihop wireless ad hoc networks [27], and this results in lower throughput and increased end-to-end delays experienced by applications [2, 8]. This problem is

emphasized by the fact that an approach where all nodes use the same common-range for all control and data transmissions exhibits a poor spectral reuse footprint (i.e., number of simultaneous transmissions that can take place in the network). A node transmitting a packet to another node in close proximity must transmit Request To Send/Clear To Send (RTS/CTS) packets with the common agreed transmission range (usually a higher power than the minimum power necessary to reach the target destination) for correct operation of the MAC layer. However, this transmission will “lock” an area limited by the sensing range where no other transmission can take place. As a result, there is an inherent space wasted in each transmission by holding to a common-range transmission approach.

Power control has had a limited use in existing IEEE 802.11 radios which has become a de-facto standard in wireless ad hoc networks. In fact, most IEEE 802.11 radios are usually configured to use the maximum transmit power (e.g., maximum transmission range) available to them. Recently, there has been a push by the research community to explore power-controlled IEEE 802.11 based MACs. In these works, however, performance metrics such as throughput decreases when reducing the transmission power. These results are in contrast to theoretical results found in [13] and [9]. While the basic IEEE 802.11 standard does not exhibit good QoS performance in wireless ad hoc networks, there are several proposals around the IEEE 802.11 standard that are customized for higher spectral reuse, and therefore, increased performance in single and multihop ad hoc networks. In the next section, we review these proposals and show how they can provide the foundations for power controlled differentiated services in wireless ad hoc networks.

The goal of this paper is to study the interplay between power control and the observed QoS delivered to applications for wireless ad hoc networks. Based on the results from this study we propose the Power Controlled QoS tuning (*PCQoS*) protocol to capture this power/QoS trade-off for applications that want to tradeoff better QoS performance at the expense of adding intermediate hops in their paths/routes.

The specific contributions of this paper are as follows. We review several proposals around the IEEE 802.11 standard targeted to achieve higher spectral reuse [22, 23, 26] for single and multihop wireless ad hoc operations, and show how these protocols can be used as the foundation for power controlled differentiated services in wireless ad hoc networks. Next, we propose, design, implement, and evaluate *PCQoS*, which is capable of trading off application QoS and power control in wireless ad hoc networks. To the best of our knowledge, PCQoS represents the first routing scheme that integrates control algorithms to realize this Power/QoS trade-off in wireless ad hoc networks. PCQoS can also be used to establish a set of differentiated

service classes in wireless ad hoc networks. For example, wireless ad hoc networks could offer two types of service classes to devices/applications: (i) a gold class, which attempts to improve the throughput and delay observed by applications/devices; and (ii) a best effort class with potentially poorer throughput and delay. PCQoS offers a number of strategies and policies that make different service classes simple to implement. Under such a regime, applications that need preferential throughput or have delay constraints would subscribe to the gold class, while all other applications would use the default best-effort class. Such a partitioning of applications in power controlled wireless ad hoc networks represents a new direction. We argue that future wireless ad hoc networks would need to provide service differentiation to possibly different classes of applications. These applications are yet to emerge but we anticipate that existing applications such as real-time streaming and transactional data applications would benefit from wireless ad hoc networks built on PCQoS techniques.

The structure of this paper is as follows. Section 2 presents an overview of MAC protocols targeted to achieve higher spectral reuse in IEEE 802.11 based networks. A detailed discussion of the motivation behind PCQoS is presented in Sect. 3. In addition, the detailed design of PCQoS is also presented. Section 4 explains how PCQoS can dynamically add and remove relay nodes in selected flows. Following this, we study the performance of PCQoS using the ns-2 simulator in Sect. 5. Related work is discussed in Sect. 6. Finally, we present our conclusion in Sect. 7.

2 Higher spectral reuse in IEEE 802.11

One of the main drawbacks of the IEEE 802.11 MAC is that it requires that all nodes in the network use a common agreed transmission power for transmission of control and data packets. In what follows, we describe several proposals made around the IEEE 802.11 standard that removes this limitation providing an increased spatial reuse. We refer to these MAC proposals as Space-Reuse CSMA (SR-CSMA) in the rest of the paper.

The two main principles governing the design of SR-CSMA MAC protocols are [23]:

- (i) *power conservation principle*, which dictates that each source must transmit using the minimum transmission power necessary to reach the intended receiver; and
- (ii) *cooperation principle*, which dictates that no source that initiates a new transmission can disrupt on-going transmissions by transmitting too “loud”. An example of the operation of SR-CSMA is presented in Fig. 1.

One of such protocols is the power controlled media access protocol (PCMAP) [23]. PCMAP uses two separate frequency channels for its operation. One channel is used for




Fig. 1 We show the operation of SR-CSMA MACs by a way of an example. In this case nodes A and C are going to transmit to nodes B and D, respectively. When a fixed common transmission power is used (*dotted circles*), only one transmission can take place at a time since the other transmission will sense the medium busy and wait for another opportunity to transmit. When nodes reduce the power to just the minimum necessary to reach the intended destination (*solid circles*), both transmissions can take place simultaneously

data traffic while the other channel is used for signaling. The packet exchange on the data channel uses a request-power-to-send (RPTS), acceptable-power-to-send (APTS), DATA-ACK packet handshake, which is similar to the RTS-CTS-DATA-ACK sequence used in IEEE 802.11. The purpose of the RPTS-APTS exchange is similar to the RTS-CTS, except that its purpose is not to force hidden terminal to back off. Rather, it is to let source and destination nodes compute the minimum transmission power to communicate with each other (the power conservation principle). In PCMAP, active receivers advertise a periodic busy tone on a signaling channel to other potential transmitters including their maximum tolerance to admit extra noise (i.e., interference). A node intending to transmit a packet must first sense the busy tone signal on the signaling channel. If a busy tone exists, then the node adjusts its transmission power such that it does not disrupt on-going transmissions prior to establishing a communication with its intended receiver (the cooperation principle).

Another example of SR-CSMA is the Power Saving MAC Protocol presented in [26]. This MAC takes advantage of power control techniques to reduce interferences among transmitters in order to increase the spatial reuse in the network. Based on the concept of Maximum Independent Set (MIS), this MAC allows for as many simultaneous transmission pairs as possible. The Interference Aware (IA)-MAC presented in [22] is yet another example of SR-MACs. IA-MAC is quite similar to the PCMAP protocol except that IA-MAC does not use an additional control channel for signaling. A different approach to increase the number of simulta-

neous transmissions in the network can be achieved by dynamically controlling the carrier sensing threshold [17, 29].

Performance results shown in [22, 23, 26] indicate that these MAC protocols allow for a greater number of simultaneous transmissions than IEEE 802.11 (i.e., higher capacity) by reducing the transmission power to the minimum levels necessary to guarantee a successful reception by the intended destination. The benefits of using these protocols over IEEE 802.11 increase as the traffic becomes more localized (e.g., when nodes communicate with other nodes in their neighborhood only).

A negative property of SR-CSMA MAC protocols is that they favor short-range transmissions over long-range ones under high traffic loads. We highlight this observation because it is this unfairness what we use to our advantage in the PCQoS discussed in Sect. 3 in support of QoS differentiation. As an example, we implemented PCMAP in a network simulator in order to first understand this unfair behavior, and then experiment with PCQoS.

Figure 2 shows the performance of PCMAP. There are 400 nodes in a 500×500 meter network with 100 flows, each one of them sending 512-byte packets for 10000 seconds of simulation time. Only 1 hop exists between source and destination nodes for a connectivity range of 250 meters (complete details of the simulation settings can be found in Sect. 5). Each source selects a destination at random within its 250 meters maximum range. Figure 2 shows the fraction of total packets received by destinations over five distance ranges (viz. 0–50, 50–100, 100–150, 150–200, and 200–250 meters, respectively) from their associated sources (we use the same 5 intervals used in [23] for comparison), which transmit either 1 or 16 packets per second. A fair MAC protocol would result in a linearly increasing number of packets transmitted at each range since the number of receivers at each range increases by $2\pi r$, where r is the distance separating source-destination pairs. In Fig. 2, we can observe that for a lightly loaded network (i.e., 1 packet per second) PCMAP supports fair behavior because the fraction of packets received increases linearly with range. In the case where the network operates under heavier traffic conditions, the fraction of the packets sent over longer distances decreases due to the unfair behavior of SR-CSMA MACs. This unfairness is the result of applying rule number 2 of SR-CSMA MACs, which dictates that no source that initiates a new transmission can disrupt on-going transmissions by transmitting too “loud”. For long-range transmissions it is unlikely that they get an opportunity to transmit in the presence of short-range transmissions in their neighborhood.

In Table 1 we show the number of flows and the unfairness factor (normalized to 1 for flows in the 200–250 meter range) over each range used for the same network setup as shown in Fig. 2. The unfairness factor in this case expresses the transmission opportunities to destinations located within

Fig. 2 PCMAP: fraction of packets sent to each range and traffic load




Table 1 Throughput unfairness of PCMAP

Range [meters]	Number of flows	Unfairness factor
0–50	3	×23
50–100	14	×10
100–150	19	×7
150–200	28	×3
200–250	36	×1

different transmission ranges. For example, each of the 19 destination nodes located within the 100–150 meter range from their respective sources has 7 times more transmission opportunities than any of the 36 destinations located within the 200–250 meter range from their respective sources. An extreme example of this unfairness phenomena exhibited by PCMAP is reflected in that any of the 3 destinations located within the 0–50 meter range have 23 times more transmission opportunities than any of the 36 destinations located in the 200–250 meter range. These results best illustrate the inherent unfairness of SR-CSMA protocols.


The inherent unfairness toward long-range transmission is not specific to PCMAP, but is a common behavior of SR-CSMA MAC protocols that provide higher spectral reuse in the network. Counter-intuitively, we use this unfairness as the basis for providing QoS service differentiation in wireless ad hoc networks. The intuition is as follows: *if we break a long-range transmission into shorter-range trans-*

missions, then we can likely increase the transmission opportunity of the resulting shorter-range transmissions, improving the end-to-end QoS observed by a particular flow. This goal can be achieved by adding relay nodes between source-destination pairs. Such approach, however, could be detrimental to other flows and to the overall capacity of the network to carry traffic as we will show later in Sect. 5. In what follows, we study this tradeoff that we call “PCQoS” and discuss its benefits and drawbacks in detail in the next section.

Without this unfairness toward long-range transmissions exhibited by SR-CSMA MACs, adding relay nodes to a flow/path would simply degrade the throughput and delay performance observed, impacting other flows in the network [10]. We will use the term *redirector* instead of relay node to differentiate while adding intermediate hops in links that otherwise can communicate directly. The availability of SR-CSMA MACs on the other hand provides a window of research opportunities into the design of systems that can target higher spatial reuse, QoS differentiation, and possibly energy-savings in multihop networks. We study these issues and open questions in the following sections.

3 PCQoS: realizing the QoS-power trade-off

In the previous discussion, we hinted as the possibility of breaking long-range links into various shorter-range links

Fig. 3 Packets received vs number of redirectors

as a way to tune application layer QoS. Now we consider building QoS mechanisms for specific applications that wish to trade-off better QoS performance. This tradeoff, to the best of our knowledge, has not been discussed in the wireless ad hoc literature.

Figure 3 shows the advantages and drawbacks of adding redirectors to a route while using SR-CSMA MACs. The simulation settings are the same as in Fig. 2 except that here one of the 36 flows in the 200–250 meter range was broken into shorter range links by adding redirectors between the end nodes. For the one redirector case, the forwarding node was chosen to be in the mid-point between the end nodes. For the 3 redirector case, the 3 forwarding nodes were positioned approximately at equal distances between the end nodes and so on. At first, adding one redirector increases end-to-end throughput by almost 1000% compared when no redirectors are added. Having three redirectors increase performance by only 21% compared with the one redirector case. For 5 and 7 redirectors performance begins to get worse. After observing this behavior we can identify two operational zones of PCQoS in Fig. 3, a stable left-hand zone where the addition of redirectors translates into better QoS performance due to the unfairness of SR-CSMA MACs toward long-range flows, and an unstable right-hand zone, where performance gets worse. The later behavior happens because of the presence of many short-range links competing against each other for accessing the channel, which

translates into larger contention delays. It is important to mention that the shape of Fig. 3, and even whether or not the addition of redirectors results beneficial depends on the particular settings of the network. As it is explained later in the evaluation section, factors such as node density, connection density and offered load per connection actually shape the behavior of Fig. 3.

When enabling the addition or removal of redirectors to achieve some coarse QoS control, we need to pay particular attention to which flows add or remove redirectors in order to assure “stable” and meaningful operations for the wireless network as a whole. Allowing all flows to add or remove redirectors may result in an unstable solution (right-hand zone of Fig. 3) where each flow attempts to optimize its own QoS constraints at the same time. We call this phenomenon the *domino effect*. The domino effect can be seen as the global impact of a local greedy strategy by a node/application/user.

In order to control the impact of the domino effect in the network, it is necessary to limit the number of adding or removing redirector operations in the network. The simplest way to accomplish this objective is to limit the number of flows that are allowed to add or remove redirectors. For example, *gold* flows can have such control to optimize their application performance while *normal* users cannot. This policy essentially differentiates between the populations of nodes/users/applications in the network. Such

a policy would help to limit the number of gold service users by an ISP in order to support the differentiated service quality over the normal users. PCQoS is motivated by this model. In PCQoS, we propose that only a subset of flows/applications/users is given the capability of adding or removing redirectors. Flows with the flexibility of adding or removing redirectors in this manner would be more sensitive than other flows in terms of their QoS requirements. For example, some applications may be transmitting rate-sensitive information such as low-rate audio or important alarm messages, while other applications may be transmitting delay-insensitive information such as local temperature measurements as in the case of sensor networks. More specifically, lets define “gold” for flows (high priority) that are QoS sensitive and “normal” (low-priority) for flows that tolerate best effort QoS. Separating flows using different priorities is not a limitation of PCQoS, but a common property of protocols that attempt to improve the average performance of a certain set of flows in detriment to others, as is the case of the DiffServ model discussed in the IETF [7]. As we will see later in the evaluation section, in case all flows become gold flows in PCQoS, performance continues to be optimal, but without QoS differentiation, similar to DiffServ.

3.1 Protocol description


PCQoS is defined by the *monitoring* and *control* phases. During monitoring periods, gold flows monitor the continuous flow of packets from their respective sources and may decide to take QoS Power-control actions or not based on a user/application specific policy. During the control phase, redirectors can be dynamically added or removed

from routes of gold flows. Positioning redirectors is concerned not only with adding or removing redirectors from the network path, but also with the location where redirectors are positioned in relation to gold flows.

Figure 4 illustrates the operational cycle of PCQoS. In this figure, we show an example trace of the performance behavior for a QoS metric (e.g., throughput, delay, etc.) for a “hypothetical flow” over time. The PCQoS cycle has *active* and *passive* operational periods. During active periods, gold flows can add or remove redirectors from their paths in order to coarsely modify their QoS/power performance trade-off. Different gold flows may have different QoS/power policy objectives. However, there are several base policies that gold flows must obey while adding or removing redirectors in order to assure the stable operation of the wireless network (we will explain these baseline policies later). By a stable network we mean a situation where a user/flow/application may trigger the addition of one more redirector to its original link, only if by doing it, the performance of a certain metric improves by a certain minimum margin. After a gold flow finishes adding or removing redirectors from its path, it moves into a “passive” operational mode for an interval when no redirectors can be either added or removed even if during that interval the observed QoS performance changes. The motivation for having active and passive periods in PCQoS is to make unlikely that two gold flows in the same neighborhood add or remove redirectors from their paths at the same time. The duration of active and passive intervals is discussed below.

Active intervals are composed of several monitoring and control periods. Figure 4 focuses on one active interval for

Fig. 4 PCQoS operational cycle



further elaboration. A destination node monitors the performance of a metric (e.g., end-to-end packet delay, throughput, etc.) for a short period before a specific policy used triggers the addition or removal of redirectors. The duration of monitoring periods should allow for the reception of multiple packets to compute the average value of the metric being measured or controlled. The duration of active periods depends on the specific policy being used and may extend over several monitoring/control intervals.

After a successful active period, a gold flow gets into passive mode and remains in that mode for some time before moving into active mode again. In case a gold flow fails to meet its QoS goals during an active period, it may be tempting for that flow to immediately get into active mode again. However, doing so may create interference with other gold flows in the neighborhood having a similar problem. When two or more gold flows in active mode overlap in space, a gold flow no longer controls its own QoS, as this is affected by other gold flows adding or removing redirectors of their own. This situation makes QoS metric readings to be unstable, thus making difficult for a gold flow to find the right number of redirectors to meet its QoS goals. As a result, there is a need to control the minimum interval gold flows spend in passive mode. Short passive intervals may create interference among neighboring gold flows, while long passive intervals may reduce the ability of a gold flow to respond to changing network conditions. The number of overlapping, gold flows (O_{GF}) can be approximated by:

$$O_{GF} \approx \left\lceil \frac{X_{gold} A_{gold}}{A_{total}} \right\rceil \quad (1)$$

where X_{gold} is the number of gold flows in the network, A_{gold} is the area of the network affected by a gold flow in active mode (i.e., the area of the network limited by the sensing range of the nodes belonging to a gold flow), and A_{total} is the total area of the network. Let us define the average duration of an active period as T_{act} . We model the duration of passive intervals T_{pas} as a random variable uniformly distributed between $[O_{GF} T_{act}, 2^{(O_{GF}+1)} T_{act}]$. This helps to reduce the probability of two or more gold flows in the same neighborhood having overlapping active periods. In case a gold flow in active mode senses another gold flow adding redirectors (this can be achieved by overhearing route-request messages by other nodes), the gold flow passes to passive mode immediately, then waits for a new random interval as before.

3.2 Monitoring-control phase

In the design of PCQoS we consider the following metrics: packet delay, packet throughput and transmission power. However, other metrics could also be monitored depending on a particular application/policy. Based on the monitoring

of one or more metrics, the receiver decides whether the observed QoS/power performance is satisfactory based on the user-specific policy being used, and may take further action to modify the number of redirectors in its path during this active period.

3.3 User policy

In PCQoS, gold users have no performance goals restrictions. What PCQoS does restrict on the other hand, are the policies (e.g., mechanisms or rules) that gold users should obey while attempting to reach their individual QoS and energy savings goals. These policies are necessary to limit the inherent QoS degradation in the network resulting from the addition of redirectors by gold users. In PCQoS we identify two stable operational points or policies that are feasible:

- *Normal*: This is the default behavior of IEEE 802.11 or SR-CSMA based networks without the addition of redirectors (e.g., packets are transmitted directly between source-destination pairs). However, applying no power control means that long range flows in the SR-CSMA MAC case will suffer degraded performance due to the unfairness of the protocol.
- *Metric Saturation Point (MSP)*: Under this policy gold users are allowed to actively add or remove redirectors. For instance, when bigger is better (i.e., throughput), we define the metric saturation point, as the point where the action of adding one more redirector to a path would not provide any significant improvement in the performance of a particular metric being controlled.

Definition Let M_k be the value of the performance metric being controlled after adding k redirectors to the route. Assuming bigger is better, redirector $k + 1$ will be added to the route only if:

$$M_{k+1} > M_k(1 + \delta) \quad (2)$$

where δ is this predefined margin that makes worth the addition of one more redirector. The idea behind limiting the number of redirectors is to limit the potential negative effect of adding more redirectors in terms of additional QoS degradation observed by other flows (both gold and best effort flows) in the wireless network.

In PCQoS, gold flows are capable of adding and removing redirectors in order to achieve their QoS/power performance tradeoff in a greedy fashion. We define the targeted performance of such a flow as $Metric^{target}$. This target could be application specific, service class specific or a default for all gold flows in the network. We define monitored performance of supporting N redirectors in a path as $Metric_N^{measured}$. During the monitoring-positioning periods,

a gold flow will add or remove redirectors in order to bring the observed performance $Metric_N^{measured}$ closer to the target performance $Metric^{target}$. In all cases gold flows can add as many redirectors as long as the metric saturation point policy described before has not been reached.

The performance of QoS metrics such as throughput or delay could be improved by either adding or removing redirectors, depending on the specific operational conditions in the network. Under certain conditions the throughput and delay performance may improve by adding redirectors due to the unfair behavior of power controlled MAC (left-hand of Fig. 3). However, in other network conditions removing redirectors could improve the throughput and delay performance because less costly packet-forwarding (i.e., less packet contention among redirectors in the path) takes place (right-hand zone in Fig. 3). As a result, gold flows may need to determine experimentally whether adding (*adding-search*) or removing (*removing-search*) redirectors leads to better QoS performance or not as the case may be. Considering that bigger is better (i.e., throughput), the algorithms shown in Table 2 control the addition and removal of redirectors during an active period determining this tradeoff point.

It is important to note that even if a flow is able to reach its target performance level during an active period, PCQoS cannot guarantee that the performance level can be maintained during preceding passive operational periods. This is because during these periods, other gold flows may attempt to optimize their own performance metrics, thereby affecting by some magnitude the QoS performance observed by all other flows in the network, as is the case with the domino effect.

4 Adding and removing redirectors

Until now we have been adding and removing redirectors without actually explaining how these two operations can take place in wireless ad hoc networks. In PCQoS we use the Power Aware Routing Optimization (*PARO*) protocol to perform these operations [10]. PARO is a routing protocol that operates above the link layer but below the network layer capable of adding redirectors to split longer-range links into various shorter-range links. Originally we tested PARO as a way to reduce the overall transmission power consumption in wireless ad hoc networks, however the very same protocol can be used in PCQoS to split long-range links/routes to improve throughput and delay performance.

4.1 PARO

At first, the operation of PARO may seem counter-intuitive because in the first iteration of PARO the source node com-

Table 2 PCQoS operation

```

# Beginning of active period
# Currently N redirectors in the path
Adding-search {
    if( $Metric_N^{measured} > Metric^{target}$ )
        begin passive interval
    else {
        ⊙ add redirector
        N ++
        if( $Metric_N^{measured} > Metric^{target}$ )
            begin passive interval
        elseif( $Metric_N^{measured} > Metric_{N-1}^{measured}(1 + \delta)$ ) {
            goto ⊙ }
        else {
            remove redirector
            begin passive interval }
    }
}

Removing-search {
# Currently N redirectors in the path
    if ( $Metric_N^{measured} < Metric^{target}$ ) {
        ⊙ remove redirector
        N --
        if( $Metric_N^{measured} > Metric^{target}$ )
            begin passive interval
        elseif( $Metric_N^{measured} > Metric_{N+1}^{measured}$ )
            goto ⊙
        else {
            add redirector
            begin passive interval }
    }
}

```

municates with the destination node directly without involving any packet forwarding by redirectors. Any node capable of overhearing both source and destination nodes can compute whether packet forwarding can reduce the transmission power in comparison to the original direct exchange between source and destination nodes. When this is the case, an intermediate node may want to become a *redirector* and send a *route-redirect* message to the source and destination nodes to inform them about the existence of a better route in terms of power efficiency to communicate with each other.

Definition Let SIR_{min} be the minimum signal to interference ratio (SIR) at which a packet can still be received properly. If $R_{i,j}$ is the measured received signal power at node i from a packet transmitted by node j at power T_j , and I_i is the local interference measured by node i then the minimum transmission power for node j to communicate with node i , $T_{j,i}^{min}$, is such that $\frac{R_{i,j}}{I_i} \geq SIR_{min}$.




Fig. 5 Redirect operation

Figure 5 illustrates how PARO operates. In this example, nodes A, B and C are located within maximum transmission range of each other and, initially, node A communicates directly with node B. Because node C is capable of overhearing packets from both A and B nodes, it can compute whether the new route $A \leftrightarrow C \leftrightarrow B$ has a lower transmission power than the original route $A \leftrightarrow B$. More precisely, node C computes that a route optimization between nodes A and B is feasible if:

$$T_{A,B}^{min} > \alpha(T_{C,A}^{min} + T_{C,B}^{min}) \tag{3}$$

The factor α in (3) and Fig. 5 restricts the area between two communicating nodes where a potential redirector node can be selected from. In Fig. 5, we show the equivalent region where a potential redirector can be located for $\alpha = 1$ and $\alpha = 2$. Similarly, we define the optimization percentage of adding a redirector between two other communicating nodes in a route, η , as:

$$\eta = 1 - \frac{T_{C,A}^{min} + T_{C,B}^{min}}{T_{A,B}^{min}} \tag{4}$$

We have shown the case where only one intermediate redirector node is added to a route between a source-destination pair. The same procedure can be applied repeatedly to further optimize a route into smaller links with the result of adding more redirectors between source-destination nodes. Figure 6 illustrates an example of a source-destination route comprised of five links with four redirectors requiring four iterations for route convergence. Figure 6 shows the route taken by data packets after each iteration and the intermediate nodes selected as redirectors after transmitting successful route-redirect requests. Figure 6 illustrates an initial route with one hop only, clearly the same procedure can be applied between an arbitrary pair of links, including the case of a route composed of multiple links or hops which is the case of multihop ad hoc networks.

PARO optimizes routes one step at a time, thus it requires several iterations to add more and more redirectors. The word “iteration” refers to the event in which a data packet triggers a node to transmit a route-redirect request for the first time. As a result PARO will add redirectors as fast as the transmission rate of data packets.

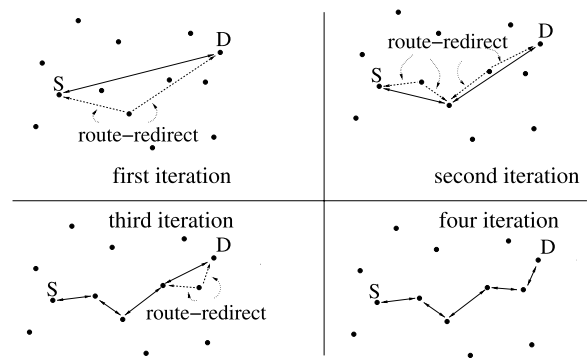


Fig. 6 PARO convergence

4.2 PCQoS and PARO

In order to support PCQoS redirector positioning we made a modification of the baseline PARO protocol. The basic PARO protocol adds as many redirectors to a route as possible. As a result, for PCQoS it would be necessary to add control over the specific number of redirectors introduced into the route.

Since more than one redirector can be added to a route during one iteration, it is insufficient to send a signaling packet requesting the addition of one redirector to the current path. This would lead to ambiguous behavior because it would not be clear which redirector, among all the potential redirectors found along a path in one iteration, offers the best QoS performance. In the case of iteration 2 (in Fig. 6), the redirector on the right-hand side is the one that should be selected because it achieves a higher η compared to the redirector on the left-hand side. We modify the baseline PARO protocol in a manner where all potential redirectors found in one iteration are first evaluated at either the source or destination points before a decision is made about which specific redirector to select.

The operation of PCQoS is different to the baseline PARO protocol in the actions taken after the reception of a route-redirect request from potential redirectors. Reception of a route-redirect request by a potential redirector in PCQoS does not trigger the immediate redirection of the flow of packets, as it occurs in the baseline PARO protocol. Rather, PCQoS creates entries in a route-redirect table and marks their state as *dormant* (i.e., not active). Dormant state entries in route-redirect tables remain inactive until a signaling message explicitly changes the state to the active state. When entries are made active, they behave exactly like route-redirect entries in a baseline PARO system. Once either the source or the destination node selects a specific redirector based on some policy decision, a packet can be sent along the path to dynamically activate a selected redirector on-demand. The same procedure applies while removing a redirector from a route except that now the last added redirector is also the first removed one.

Table 3 Simulation parameters

Parameters	Value
Area	500 meters × 500 meters
Nodes	50 and 400
Connections	50 and 100
Traffic load	1, 4, 16, 32, 64, 128 pkt/sec
Traffic type	CBR/UDP
Packet size	512 bytes
Nodes per range	3,14,19,28,36
Ptmin	0.1778e-3 W
Ptmax	0.707945 W
SIR thresh	6 dB
SIR target	10 dB
CS thresh	1.5849E-11 W
Rx thresh	3.981E-10 W
Rx target	4.06062E-10 W

5 PCQoS evaluation

We use the *ns-2* network simulator to analyze the operation of PCQoS. We use the PCMAP MAC protocol as an example of SR-MAC, as defined in [23]. We extend our previous implementation of the baseline PARO protocol [10] to implement PCMAP and the positioning and monitoring components of PCQoS. An overview of the simulation settings is displayed on Table 3.

Communication between two nodes in PCMAP uses RPTS-APTS packet handshake signaling before the actual data transmission takes place. We reuse the same module to compute the minimum transmission power used in the baseline PARO to implement PCMAP (to see how PARO computes the minimum transmission power between two nodes refer to [10]). In the PCMAP implementation, however, we added a local copy of the noise to the packet header of each transmitted packet, as defined in the specifications of PCMAP [23]. This addition is necessary for PCMAP because the noise (or interference) levels are not negligible as is the case with the baseline PARO protocol. PCMAP operations including the conservation and cooperation principles are implemented according to [23].

The propagation model in *ns-2* is based on a Friss model for short ranges and a Two-ray model for longer ranges. This model is appropriate for outdoor environments where a strong line of sight signal exists between the transmitter and receiver nodes, and where the antennas are omnidirectional. The Friss model computes the received power by:

$$R_{j,i} = \frac{T_{i,j} G_t G_r \lambda^2}{(4\pi d)^2 L} \quad (5)$$

where $R_{j,i}$ is the received power at node j when node i transmits with power $T_{i,j}$, d is the distance separating trans-

mitter from the receiver, λ is the signal wavelength, L is an adjustment factor and G_t and G_r are the antenna gain of the transmitter and receiver nodes, respectively. The two-ray propagation model assumes there are two main signal components. This model computes the strength of the received signal at the destination nodes as:

$$R_{j,i} = \frac{T_{i,j} G_t G_r h_t^2 h_r^2}{d^4} \quad (6)$$

where h_t^2 and h_r^2 are the antenna height of the transmitter and receiver nodes, respectively. The cutoff distance d_c where the model changes from Friss to Two-ray is given by:

$$d_c = \frac{4\pi h_t h_r}{\lambda} \quad (7)$$

In the following section we present an evaluation of PC-QoS. We experiment with different operational aspects of PCQoS and show how gold flows can add or remove redirectors to dynamically modify their observed QoS performance. Each point in the presented graphs is the average of 10 experiments, each of them using a different seed number while locating nodes in the network.

We did not consider mobility in the evaluation of PCQoS because mobility adds another dimension and complexity to the problem. The same ideas and solutions presented in the baseline PARO protocol to support mobile nodes, such as keeping a minimum rate of packets flowing between source-destination pairs and increasing the minimum transmission power of each transmission by some margin, are applicable to both IEEE 802.11 and SR-MAC protocols [10].

The signaling overhead of PCQoS is related mainly to the overhead incurred by PARO, which is the underline routing scheme in charge of adding and removing redirectors. PARO needs two signaling packets to add one redirector, and no signaling packets to remove a redirector, so adding X redirectors in a flow requires $2X$ signaling packets. Because many data packets can be transmitted on a route before there is a need to add or remove redirectors, the percentage of overhead control packets incurred by PCQoS is minimum compared with data traffic.


5.1 PCQoS performance

In what follows, we evaluate several issues of the performance and behavior of the proposed PCQoS protocol.

5.1.1 Individual PCQoS behavior

Figure 7 shows traces of the throughput performance achieved for a gold flow in the 200–250 meter range under PCQoS. The dashed line in Fig. 7(a), (b), (c) denotes the targeted performance (desired throughput in this case). Monitoring intervals are set to 5 seconds and the duration

Fig. 7 Throughput performance of a flow operating PCQoS



of passive periods is uniformly distributed between 150–800 seconds. The three graphs shown in Fig. 7 contrast the PCQoS operation of the test flow in different network conditions: In Fig. 7(a) only 10% of the flows in the network are gold flows, Fig. 7(b) same as (a) but with a higher target throughput, and Fig. 7(c) where 30% of the flows are gold flows.

Figure 7(a) shows the PCQoS behavior when the targeted performance is higher than the initial performance. In this case, the test flow uses the adding-search algorithm (detailed in Sect. 3.2) anticipating that there are fewer short-range flows in the neighborhood and thus the unfair behavior of the MAC could improve throughput. Adding 1 redirector after 30 seconds into the trace brings the performance above the target performance, then the flow moves into a passive period. Denying gold flows the capability of adding more redirectors after performance goals are reached is an important property of PCQoS. Active and passive periods are shown above the trace line in addition to the number of redirectors being used in the path.

Figure 7(b) shows the PCQoS behavior when the target throughput is again above the initial performance without redirectors. The test flow uses the adding-search algorithm adding 1, then two and finally a third redirector in its path. Because the performance improvement after the third predictor is below the saturation point ($\delta = 10\%$ in this example), the third redirector is removed and only two redirectors are selected for this active period. Note that during the passive interval the monitored throughput performance changes

as a result of the domino effect created by other gold flows adding redirectors of their own.

Figure 7(c) shows the behavior of PCQoS when there are already 3 redirectors in a test flow after 500 seconds into the trace and the performance is again below the target performance at the beginning on an active period. Again the test flow uses the adding-search algorithm to add a 4th redirector into the path. Since the desired performance is not met, the test flow then uses the removing-search algorithm up to 1 redirector in this case providing the highest performance possible.

5.2 PCQoS aggregate performance

In the previous experiments we have shown the performance of PCQoS for individual flows. Now we analyze the aggregate impact on QoS when a subset of flows in the network is allowed to add redirectors. For these experiments, we disabled the metric saturation point policy of PCQoS in order to see what would happen when certain flows added a specific number of redirectors.

We evaluate a network of 400 nodes in a 500×500 meter network with 100 flows each sending sixteen 512-byte packets per second. Each source picks a destination at random within its 250 meter range. For these experiments we select the following 5 scenarios shown in Table 4. The term $X(N)$ in Table 4 means X gold flows in this range added N redirectors to their paths. We selected these 5 scenarios

Fig. 8 PCQoS: fraction of packets received versus range




Table 4 Simulation scenarios for aggregate performance analysis

	0–50 3 flows	50–100 14 flows	100–150 19 flows	150–200 28 flows	200–250 36 flows
S1	0	0	0	0	0
S2	0	0	0	0	10(1)
S3	0	0	0	0	10(3)
S4	0	0	6(1)	9(1)	12(1)
S5	0	14(2)	19(3)	28(5)	36(7)

because we think they better show the advantages and drawbacks of PCQoS: Scenario 1 corresponds to a SR-CSMA network without PCQoS (i.e., no redirectors are added to any path). Scenario 2 corresponds to the case where PCQoS is applied randomly in 10 of the 36 flows in the 200–250 meter range, and there is one redirector between end points only. Scenario 3 is the same as Scenario 2 except that now 3 redirectors are positioned between end points. In Scenario 4 one third of the flows in the 100–150, 150–200, and 200–250 meter range added one redirector only to their paths. Finally, in Scenario 5 all flows in the network added as many redirectors to their paths as necessary so that all resulting links were in the 0–50 meter range.

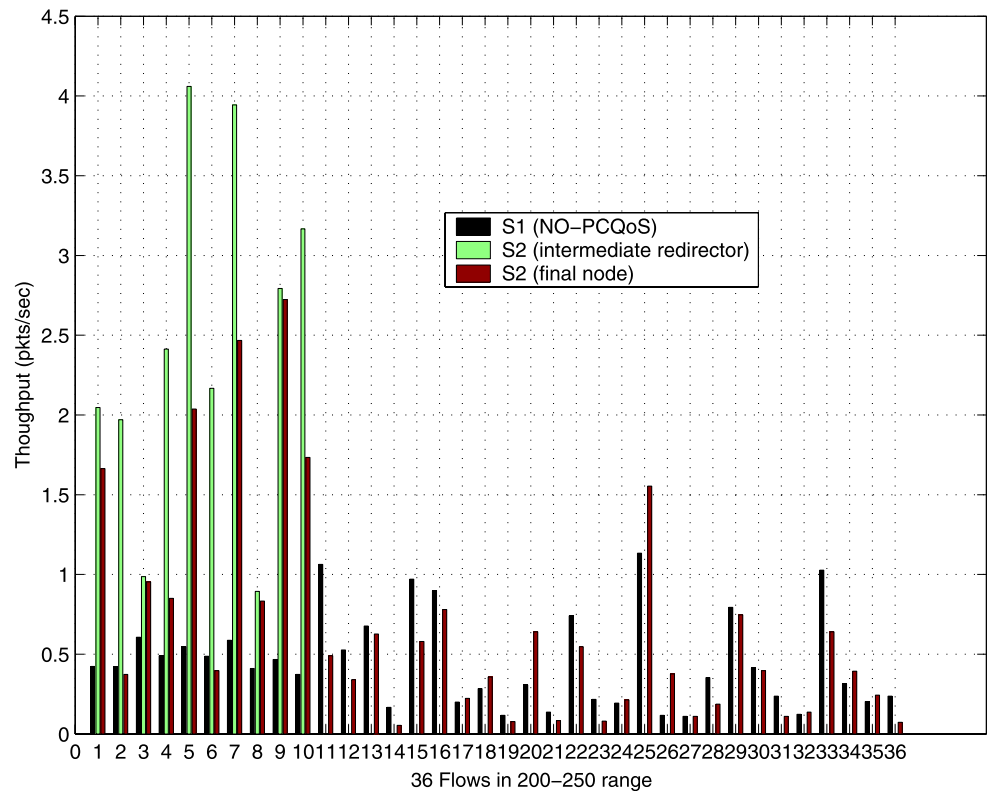
Figure 8 shows the fraction of the total packets received by destinations for each scenario over five distance ranges (0–50, 50–100, 100–150, 150–200, and 200–250 meters, re-

spectively) from their sources. The fact that we got 3, 14, 19, 28 and 36 flows for 0–50, 50–100, 100–150, 150–200, and 200–250 meter range, respectively in Table 4, is a direct result of letting each source pick a destination at random within its 250 meter range (see Sect. 2).

Let first look at Scenario 2 where 10 out of 36 flows in the 200–250 meter range added 1 redirector to their paths only, and compare its performance with Scenario 1 (no PCQoS). As we can see in Fig. 8 the fraction of the total packets received by destinations in the 200–250 meter range improves slightly compared with Scenario 1. This is a direct result of the increased throughput obtained by the 10 gold flows in this range. The negative side is that now we have less packets received for flows in the 150–200 meter range compared with Scenario 1. Flows in the 150–200 meter range obtained lower throughput in Scenario 2, because there are now 20 more “links” (due to the 10 flows in the 200–250 meter range split into 2 100–150 links). This is a clear example of the domino effect where a local greedy decision impacts the performance seen by others.

Figure 9 shows the actual throughput obtained by the 36 flows in the 200–250 meter range of Scenario 2. The first 10 flows in Fig. 9 correspond to the selected 10 gold flows of Scenario 2. The results in Fig. 9 clearly show the increased throughput obtained by selected gold flows compared with the rest of the flows in the 200–250 meters range that did not add redirectors to their paths. Figure 9 also shows the average throughput received after each redirector of selected

Fig. 9 Throughput performance of 10 flows in the 200–250 meter range having 1 redirector



gold flows. The common trend is that the average throughput received after redirector i in the path from source to destination is higher than in redirector $i + 1$. This is a pattern already reported in the literature [8] and created by traffic sources with no rate control (e.g., CBR). It is likely that this anomaly will disappear if a congestion control mechanism is used at the transport layer (e.g., TCP).

Scenario 3 is similar to Scenario 2 except that now the 10 flows selected added 3 redirectors instead of one redirector. Figure 10 shows the average throughput received by each of the 10 selected flows in Scenario 3. For comparison we also show the throughput for Scenario 1. As we can see in Fig. 10, in most cases the addition of 3 redirectors translated into higher throughput for these flows compared with Scenario 1 (flow 7 is the exception). As for Fig. 8 respects, the fraction of the total packets received by destinations in the 150–200 meter range in Scenario 3 is lower compared with Scenario 2. This again is because the appearance of multiple shorter-range links created by the addition of three redirectors to the 10 selected flows of Scenario 3.

In Scenario 4 we have that one third of the total flows in the 100–150, 150–200, and 200–250 meter ranges added one redirector to their paths. In Fig. 8 we can observe that selected flows in the 100–150 and 150–200 meter range obtained higher throughput compared with selected flows in the 200–250 meter range. Gold flows in the 200–250 meter range benefit little by adding one redirector in this scenario because there is a higher number of shorter range links in-

troduced by selected flows in the 100–150 and 150–200 meter ranges which added increased unfairness toward longer range flows.

Finally, in Scenario 5 all 100 flows in the network added as many redirectors to their paths as necessary so that all resulting links were in the 0–50 meter range. This scenario is similar to a IEEE 802.11 network with a common transmission range of 50 meters. In this case the throughput obtained by flows in the 200–250 meter range is still better than Scenario 1, however, as we will see later, the overall throughput of the network degrades quickly in this scenario. We consider this scenario in order to show the potential danger if no control is taken on the number of redirectors added by gold flows in PCQoS. Let's remember that for these experiments we disable the metric saturation policy of PCQoS. It is expected that this scenario should never occur in PCQoS since flows will not add as many redirectors if doing so translates into lower performance.

5.2.1 PCQoS versus offered load

Figure 11 shows the aggregate average throughput for all 100 flows in the network for Scenarios 1–4 as a function of the offered load. In these experiments we did not consider Scenario 5 because we already discussed such scenario should not occur once the metric saturation policy of PCQoS is enabled. From this figure we can observe that for any

Fig. 10 Throughput performance of 10 flows in the 200–250 meter range having 3 redirectors





Fig. 11 PCQoS versus offered load (all flows)

offered load, the highest throughput is always achieved by Scenario 1, where no PCQoS was implemented. It is important to note, however, that for scenarios S2, S3 and S4, their throughput is within 10% of the throughput obtained by scenario S1, meaning that PCQoS is capable of providing some coarse QoS control to selected flows without significantly reducing the original capacity of the network. As it can also be observed in this figure, for low loads (e.g., 1 pkt/sec) the performance of all 4 scenarios is similar. This behavior is expected since for this traffic conditions long-range flows rarely find the channel busy (i.e., there is no unfairness towards long-range flows).

Figure 12 shows the average received throughput of selected gold flows in the 200–250 meter range only for Scenarios 1–4 as a function of the offered load. Again, in this figure we observe that for 1 packet per second per flow performance of all 4 scenarios is similar. The 4 packets per second case is interesting because performance improves by adding one redirector, but then performance decreases while adding 3 redirectors. Moving from 1 to 4 packets per second increases the unfairness factor towards long-range flows, but not enough to the point where adding 3 redirectors resulted beneficial. It is only for 8 packets per second and above that there is enough unfairness, so it results beneficial the addition of 3 redirectors for long-range flows.

5.2.2 PCQoS versus node density

In the previous experiments we considered a network of 400 nodes located in a 500 m × 500 m area. In order to study the behavior of PCQoS on a less dense network we performed a new series of experiments considering 50 nodes in the 500 m × 500 m area, and 50 connections only. Having 50 connections in the network reduces by half the number of connections per transmission range. This is now we have 2, 7, 9, 14 and 18 connections in the 0–50, 50–100, 100–150, 150–200, and 200–250 range, respectively. However, the number of gold flows per range for scenarios S1–S4 remains exactly as before (e.g., S2 consists of 10 gold flows in the 200–250 range adding 3 redirectors each).

In Fig. 13 we show the average received throughput per connection considering all flows. While throughput in the denser network remains higher for all 4 scenarios, this is mostly the result of having a higher offered load. Similarly, the aggregate throughput considering all flows for Scenario 5 is higher compared with Scenario 1 for the denser network. This is an important result because it shows even if all flows become gold flows; throughput is at least as good as the scenario where there is no QoS differentiation.

Figure 14 shows the average received throughput of selected gold flows in the 200–250 range. In this figure gold




Fig. 12 PCQoS versus offered load (gold flows in the 200–250 range)




Fig. 13 PCQoS performance for 400 and 50 nodes (all flows)




Fig. 14 PCQoS performance for 400 and 50 nodes (gold flows in the 200–250 range)

flows in the 200–250 meters range are able to obtain a higher throughput by adding redirectors even for the 50 nodes network. However, the presence of fewer nodes for the 50 nodes network reduces the availability of potential redirectors for PCQoS, thus limiting its full potential. In fact, with 50 nodes in the network only we could not run scenario S5 because there were not enough nodes available in the network to split long-range flows into 0–50 meters links.

6 Related work

The state of the art in QoS control for wireless ad hoc networks is best represented by the COWPOW system [24]. In [24], the authors present a system where mobile nodes are capable of switching the value of the common-range transmission power they use. Mobile nodes in this system periodically reduce this value and stop right before the first partition of the network occurs. Another important difference of PCQoS and the COMPOW proposal is that in contrast to common-range transmission based proposals where users get a similar QoS performance, PCQoS supports service differentiation with multiple policies. We believe that such a system is better suited to support different types of emerging applications that may require different QoS/power trade-off being supported by the network.

A signaling system supporting QoS in mobile ad hoc networks is discussed in the INSIGNIA project [19]. The INSIGNIA system creates QoS reservation at intermediate nodes visited by the data packets/flows en-route toward destinations. The transmission range used by the INSIGNIA protocol is based on the maximum common-range (like other MANET systems). An advantage of this system is that it is capable of locally restoring reservations in cases where intermediate hops move out of the route, thus it avoids costly end-to-end QoS re-adaptation.

Another example of QoS provisioning in ad hoc networks is the SWAN system [1]. In [1], the authors present a service differentiation system for stateless wireless ad hoc networks. This system is based on the notion that provisioning QoS to applications inside ad hoc networks is rather difficult and end-to-end QoS adaptation results more appealing in such environments. In [1], selected flows monitor end-to-end performance and adjust their transmitting rates according to the service class they belong to. This system uses common-range transmission principles.

In the work described in [2], the authors discuss the impact of TCP throughput on the number of forwarding nodes, or the equivalent common-range transmission value used in static wireless ad hoc networks for unreliable links. Results presented in [2] show that there is an optimum transmission range that maximizes TCP throughput. Other examples

of TCP behavior over wireless links, not necessarily related to wireless ad hoc networks but wireless networks in general include [3, 4, 16]. The use of relay nodes in multi-rate IEEE 802.11 based networks is presented in [20]. In this work nodes located far away from the AP take advantage of intermediate relay nodes to connect to the AP at higher data rates.

The work in [18] presents a QoS topology control for wireless ad hoc networks targeted at reducing the transmit power while meeting some QoS demands, while the work in [6] finds the best of all possible routes to meet some QoS requirements.

Examples of QoS adaptation systems for wireless links (not necessarily related to wireless ad hoc networks) include modifications to the link schedulers [5, 11, 21]. The main feature of these adaptation systems is that they react to link errors and compensate affected flows when links conditions improve.

The performance of IEEE 802.11 over wireless ad hoc networks is studied in [27]. Results from [27] show that one of the main reasons for the poor utilization of IEEE 802.11 over wireless ad hoc networks is its long sensing range. This issue is also studied in [8].

Finally, there are various works studying the advantage of long hops over short hops in wireless ad hoc networks. In [14, 15], the authors give various reasons not to use many short hops over few long hops. The arguments given by the authors consider various aspects including mobility, power consumption, link reliability, fading conditions, etc. The actual advantage or disadvantage of long versus short hops will finally depend on the particular settings of the network being considered.

7 Conclusion

In this paper we studied the impact of adding or removing redirectors in a multihop wireless network on traditional QoS metrics. We first study the unfair performance towards long-range flows exhibited by space-reuse CSMA MACs. We showed how this unfair behavior can be used as a foundation for QoS service differentiation in wireless ad hoc networks. We proposed PCQoS, which builds QoS mechanisms into the baseline PARO system for specific applications that wish to modify their observed QoS performance. In PCQoS, selected flows add or remove redirectors from their paths in order to coarsely modify their QoS. We show through simulations that PCQoS is capable of tuning the QoS observed by selected flows, without significantly degrading the overall capacity of the network. To the best of our knowledge, PCQoS represents the first power-controlled QoS routing protocol for wireless ad hoc networks that is based on the foundation of variable-range transmission control.

Acknowledgements This work is supported in part by research funds from UNAM/PAPIIT Grants' IN106609, IN108910 and CONACYT Grants 105117, 105279.

References

- Ahn, G.-S., Campbell, A. T., Veres, A., & Sun, L.-H. (2000). SWAN: service differentiation in stateless wireless ad hoc networks. In *Proceedings of IEEE Infocom*, New York, NY.
- Ali, I., Gupta, R., Bansal, S., Misra, A., Razdan, A., & Shorey, R. (2002). Energy efficiency and throughput for TCP traffic in multihop wireless networks. In *Proceedings of Infocom*, New York, NY.
- Balakrishnan, H., Seshan, S., & Katz, R. (1995). Improving reliable transport and handoff performance in cellular wireless networks. *ACM Wireless Networks* 1(4).
- Balakrishnan, H., Padmanabhan, V. N., & Katz, R. H. (1997). The effects of asymmetry on TCP performance. In *Proceedings of ACM Mobicom*, Budapest, Hungary.
- Bhagwat, P., Bhattacharya, P., Krishna, A., & Tripathi, S. (1997). Enhancing throughput over wireless LANs using channel state dependent packet scheduling. In *Proceedings of IEEE Infocom*, Kobe, Japan.
- Chen, Lei, & Heinzelman, W. B. (2005). QoS-aware routing based on bandwidth estimation for mobile ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 23(3).
- Clark, D., & Wroclawski, J. (1997). An approach to service allocation in the Internet. *draft-clark-diff-svc-alloc-00.txt*, work in progress.
- Li, J. et al. (2002). Capacity of ad hoc wireless networks. In *Proceedings of ACM Mobicom*, Atlanta.
- Gomez, J., & Campbell, A. T. (2004). A case for variable-range transmission power control in wireless multihop networks. In *Proceedings of IEEE Infocom*, Hong Kong, China.
- Gomez-Castellanos, J., Campbell, A. T., Naghshineh, M., & Bisdikian, C. (2003). PARO: supporting transmission power control for routing in wireless ad hoc networks. *Journal on Wireless Networks (WINET)*, 9(5), 443–460.
- Gomez-Castellanos, J., & Campbell, A. T. (2003). Supporting application and channel dependent quality of service in wireless networks. *Journal on Wireless Networks (WINET)*, 9(1), 21–35.
- Goodman, D., & Mandayam, N. (1999). Power control for wireless data. In *Proceedings of IEEE Momuc*, San Diego, California.
- Gupta, P., & Kumar, P. R. (2000). The capacity of wireless networks. *IEEE Transaction on Information Theory*, IT-46(2).
- Haenggi, M. (2004). Twelve reasons not to route over many short hops. *IEEE VTC 2004*, Los Angeles.
- Haenggi, M., & Puccinelli, D. (2005). Routing in adhoc networks: a case for long hops. *IEEE Communications Magazine*, 43, 93–101.
- Holland, G., & Vaidya, N. H. (1999). Analysis of TCP performance over mobile ad hoc networks. In *Proceedings of ACM Mobicom*, Seattle, Washington.
- Jeong, K., & Lim, H.: (2008). Experimental approach to adaptive carrier sensing in IEEE 802.15.4 Wireless Networks. *Proceedings of ACM CoNEXT*, Madrid.
- Jia, X., & Li, D. (2004). QoS topology control in ad hoc wireless networks. In *Proceedings of IEEE Infocom*, Hong Kong, China.
- Lee, S., & Campbell, A. (1998). INSIGNIA: in-band signaling support for QOS in mobile ad hoc networks. In *Proceedings of IEEE Momuc*, Berlin.
- Liu, J.-S., & Lin, C.-H. R. (2005). An opportunistic relay method for increasing throughput in multirate IEEE 802.11 wireless LAN. *IEICE Transactions on Communications*, E88-B(6), 2672–2675.
- Luo, H., Songwu, L., & Bharghavan, V. (2000). A new model for packet scheduling in multihop wireless networks. In *Proceedings of ACM Mobicom*, Boston MA.
- Maniezzo, D., Bergamo, P., & Gerla, M. (2003). How to outperform IEEE 802.11: interference aware (IA) MAC. In *Proceedings of MEDHocNet 2003*, Mahdia, Tunisia.
- Monks, J., Bharghavan, V., & Hwu, W. (2001). A power controlled multiple access protocol (PCMAP) for wireless packet networks. In *Proceedings of Infocom*, Anchorage, Alaska.
- Narayanawamy, S., Kawadia, V., Sreenivas, R. S., & Kumar, P. R. (2002). Theory, architecture, algorithm and implementation of the COMPOW protocol. In *Proceedings of the European Wireless Conference*, Florence, Italy.
- Rashid-Farrokhi, F., Liu, K. J. R., & Tassiulas, L. (1996). Transmit beam forming and power control for cellular wireless systems. *IEEE Journal on Selected Areas in Communications*, 16(8).
- Shih, K.-P., Chang, C.-Y., Chou, C.-M., & Chen, S.-M. (2005). A power saving MAC protocol by increasing spatial reuse for IEEE 802.11 ad hoc WLANs. In *International Conference on Advanced Information Networking and Applications*, Taipei, Taiwan.
- Xu, S., & Saadawi, T. (2001). Does the IEEE 802.11 MAC protocol work well in multihop ad hoc networks. *IEEE Communications Magazine*, 39(6) 130–137.
- Xue, F., & Kumar, P. R. (2006). Scaling laws for ad hoc wireless networks: an information theoretic approach. *Foundations and Trends in Networking* 1(2).
- Zhu, J., Guo, X., Yang, L. L., Conner, W. S., Roy, S., & Hazra, M. M. (2004). Adapting physical carrier sensing to maximize spatial reuse in 802.11 mesh networks. *Wireless Communications and Mobile Computing* 4(8), 933–946.



Javier Gomez received the B.S. degree with honors in Electrical Engineering in 1993 from the National Autonomous University of Mexico (UNAM) and the M.S. and Ph.D. degrees in Electrical Engineering in 1996 and 2002, respectively, from Columbia University and its COMET Group. During his Ph.D. studies at Columbia University, he collaborated and worked on several occasions at the IBM T.J. Watson Research Center, Hawthorne, New York. His research interests cover routing, QoS, and MAC design for wireless ad hoc, sensor, and mesh networks. Since 2002, he has been a Professor with the National Autonomous University of Mexico. Javier Gomez is member of the SNI since 2004.



Luis A. Mendez received his B.Sc. in Electrical Engineering at the Institute of Technology in Tuxtla Gutierrez, Chiapas, in 1999, and the M.S. with honors in Electrical Engineering at the National Autonomous University of Mexico (UNAM) in 2005. Since 2006 he works as a System Engineer at Avaya Mexico. His areas of interest include convergence of Voice and Data networks and Enterprise Communications Systems.



Victor Rangel obtained his B.Sc. with honors in Computer Engineering in 1996 from the National Autonomous University of Mexico (UNAM). He received his M.Sc. and Ph.D. in Electrical Engineering in 1998 and 2002, respectively, from the University of Sheffield, England. His areas of interest are routing and MAC modeling for Wireless Access Networks, QoS, Scheduling, handover and traffic modelling. Since 2002, he has been an Associate Professor at the Department of Telecommunications Engineering, School of Engineering (UNAM).



Andrew T. Campbell is a Professor of Computer Science at Dartmouth College where he leads the Mobile Sensing Group. Prior to joining Dartmouth in 2005 Andrew was an Associate Professor of Electrical Engineering at Columbia University (1996–2005) and a member of the COMET Group where he developed a number of mobile networking technologies. His current research focuses on people-centric sensing where he leads the MetroSense project.

Andrew received his Ph.D. in Computer Science (1996) from Lancaster University, England, and the NSF Career Award (1999) for his research in programmable wireless networking. Prior to joining academia he spent 10 years working in industry both in Europe and the USA in product research and development of computer networks and wireless packet networks. Andrew has been a technical program chair for ACM MobiCom and ACM MobiHoc, the general chair for ACM SenSys 2006, and SenSys steering committee chair 2008–2009. He spent his sabbatical year (2003–2004) at the Computer Lab, Cambridge University, as an EPSRC Visiting Fellow, and Fall 2009 as a Visiting Professor at the University of Salamanca, Spain.