

Routing Primitives for Wireless Mesh Networks: Design, Analysis and Experiments

Stanislav Miskovic and Edward W. Knightly

Abstract—In this paper, we consider routing in multi-hop wireless mesh networks. We analyze three standardized and commonly deployed routing mechanisms that we term “node-pair discovery” primitives. We show that use of these primitives inherently yields inferior route selection, irrespective of the protocol that implements them. This behavior originates due to overhead reduction actions that systematically yield insufficient distribution of routing information, effectively hiding available paths from nodes. To address this problem, we propose a set of “deter and rescue” routing primitives that enable nodes to discover their hidden paths by exploiting already available historic routing information. We use extensive measurements on a large operational wireless mesh network to show that with node-pair discovery primitives, inferior route selections occur regularly and cause long-term throughput degradations for network users. In contrast, the deter and rescue primitives largely identify and prevent selection of inferior paths. Moreover, even when inferior paths are selected, the new primitives reduce their duration by several orders of magnitude, often to sub-second time scales.

I. INTRODUCTION

Deployed mesh networks employ routing protocols based on the IEEE 802.11s standard, proprietary protocols such as those developed by Motorola [1], Mikrotik [2], Cisco [3], and Coauthored [4], and research routing protocols such as AODV-ST [5] and HOVER [6]. Unfortunately, we will show that common elements of such routing protocols can yield severely inferior routes that persist for long time scales.

In this paper, we first analyze these common routing elements, referring to them as *node-pair discovery* primitives. These primitives are: (1) constrained flooding, (2) unicast feedback, and (3) temporal ordering of route discovery information. We localize the general problem of *inferior route selection* to one of the inherently incomplete distribution of routing information. Specifically, node-pair discovery primitives can systematically suppress the distribution of information about the best paths for many nodes participating in route discovery. Such *participating* nodes are then forced to re-route to inferior paths based on other received routing information, without even being aware that better paths exist. Consequently, this inconsistent routing state causes nodes to perceive their inferior paths as the optimal ones, thus preventing them from trying to restore their true best paths until a subsequent instance of route discovery.

Second, we develop a *historic ranking principle* targeted towards prevention of inferior route selections and restoration of the true best paths. In particular, this principle provides route selection with valuable information that may otherwise be systematically hidden by the node-pair discovery primitives.

To this end, our principle does not induce any additional traffic overhead, but instead relies on historically persistent network properties and readily available routing information from previous route discoveries. Based on this information, we rank all paths previously reported to the node, thus enabling identification of a subset of node’s candidate paths that are likely to be the true best path. Therefore, the node infers potentially inferior route selection whenever it fails to receive route discovery reports about this subset of best-ranked paths. Note that this inference also addresses the problem of physically lost routing information, which significantly improves the robustness of least-cost route selection in inherently lossy wireless networks.

Third, we apply the historic ranking principle towards the design of two low-overhead routing primitives that help prevention of inferior route selection and restoration of true best paths. Specifically, while route selection is still based *only* on the presently reported routing information, our primitives ensure that no route selection is finalized *until* a node receives route-discovery updates from all of its historically best-ranked paths. The DETER primitive enables a node itself to ensure selection of its best paths, while the RESCUE primitive employs a node’s neighbors to initiate recovery from a node’s inferior paths by offering it better paths. Recoveries initiated by the DETER primitive have complete information about the node’s best-ranked paths, and can therefore make an informed query about the unreported metric costs of specific paths. On the other hand, while the RESCUE primitive does not have such precise ranking information available, it helps address problems that cannot be solved by the DETER primitive, e.g., losses of the DETER recovery packets, a node’s insufficiently trained historical ranking, network re-configurations, etc.

Finally, we preform an experimental and simulation-based evaluation of both currently-employed “node-pair routing” and our historically-assisted routing. We first evaluate the operational behavior of node-pair routing in a large wireless mesh network, Technology For All (TFA) [7]. Our results confirm that current routing primitives indeed fail to consistently select high quality network paths. We also show that poorly selected paths can have significantly higher routing-metric costs, and their duration can extend to minute time scales. Next, we show by measurements that despite a large-scale network having many variable properties (channel state, traffic load, etc.), a number of key properties are largely persistent, e.g., throughputs of isolated paths and throughput rankings of fully backlogged contending links. Having validated these premises of our historically-assisted designs, we use simulations to conduct a per-packet evaluation of the deter and rescue primitives. Our results show that these primitives largely enable avoidance of inferior route selections. Moreover, when inferior selections do occur, the deter and rescue primitives reduce the duration

Stanislav Miskovic (misko@rice.edu) and Edward W. Knightly (knightly@rice.edu) are with the ECE Department of Rice University, Houston, TX (<http://www.ece.rice.edu/networks>).

of such inferior paths by several orders of magnitude, often to sub-second time scales.

The remainder of this paper is organized as follows. In Section II, we analyze node-pair discovery primitives and show that they inherently cause inferior route selections. In Section III, we introduce the historic ranking principle, while in Section IV we design our deter and rescue primitives. In Section V, we extensively evaluate routing based on the node-pair discovery primitives, the historically-persistent network properties, and routing assisted by our deter and rescue primitives. Finally, we present related work and conclude in Sections VI and VII.

II. ORIGINS OF INFERIOR ROUTE SELECTION

In this section, we show that inferior route selections occur because insufficient routing information is distributed for nodes to identify their least-cost paths. Although packet loss is sufficient to cause this problem, we here focus on a more critical issue. We show that node-pair discovery primitives themselves may systematically prevent distribution of the least-cost routing information for many nodes. To this end, we first overview node-pair discovery primitives, and then we show how their actions can produce inferior paths.

A. Overview of the node-pair discovery primitives

The main goal of any node-pair route discovery is to identify a single *a priori* unknown least-cost path requested by a source node to some specific destination node. To this end, the first phase of route discovery identifies a path leading to the source and the second phase identifies a path leading to the destination.

Constrained Flooding Primitive. The path leading to the source node is discovered by constrained flooding. The source initiates flooding of the route discovery packets that accumulate metric costs of candidate paths while traversing each node. Based on reception of such packets, each traversed node identifies potential best-path candidates leading the source and *selects* the least cost candidate for its own path to the source. Each node also *constrains* the flooded discovery by forwarding *only* the discovery packets that report better accumulated path-metric costs. All other discovery packets are silently discarded. Note that such constrained flooding indeed enables the destination to select its least-cost path to the source.

Unicast Feedback Primitive. The path leading to the destination is discovered by unicasting the discovery-feedback information. Once flooded discovery reaches a node allowed to inform the source about a candidate path leading to the destination¹, a node will stop forwarding flooded route discovery packets, select the path leading to the source, and unicast the feedback about its path leading to the destination in the reverse direction of its selected path leading to the source. Unicasting such feedback is done under a common assumption that best paths traverse similar nodes in both directions of communication. If this assumption is true, the source would indeed be able to identify its least-cost path to the destination. However, such feedback also forms a *coupling* property of the

¹Depending on a configuration of a routing protocol, the discovery feedback can be generated either only by the destination or it can also be generated by other participating nodes on behalf of the destination.

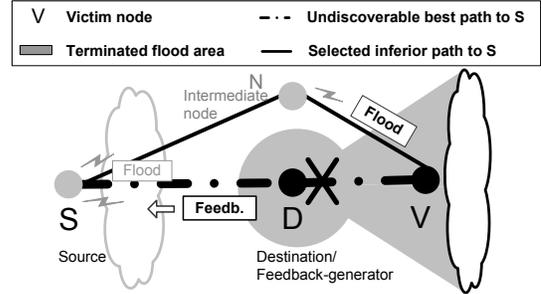


Fig. 1. Partial termination of route discovery at the node FG creating a terminated flood area over the victim node V .

forward and reverse directions of route discovery that will be crucial in our analysis of inferior route selections.

Sequencing Primitive. Routing protocols in general rely on the routing sequence information to temporally order routing packets, constrain flooding of route discovery, and prevent occurrence of routing loops (e.g., see [8]). According to the temporal ordering of routing information, each node has to re-select its path whenever it receives more up-to-date information about an end-point of that path. To this end, the source announces its most up-to-date routing sequence information via the flooded discovery packets, while the feedback packets carry such information about the destination. Finally, flooding reduction and loop prevention are provided by discarding the packets that do not carry sufficiently updated routing sequence information to help further route discovery.

B. Inferior route selections caused by routing primitives

We start our analysis of inferior route selections with an observation that routing based on node-pair discovery primitives originates from algorithms that provably identify least-cost paths for *all* nodes in wired networks, e.g., the distributed Bellman-Ford algorithm. However, in order to adapt to wireless environment, “node-pair routing” implements additional overhead reduction techniques that can excessively suppress distribution of routing information, thus forcing inferior route selections. In fact, the only nodes that can generally identify their least-cost paths during the node-pair route discovery are: (1) the source, (2) its requested destination, and (3) the nodes connecting them via the least-cost path. All other participating nodes may *not* be informed about their best paths, while still being informed about their inferior paths. Next, we present an analysis of inferior route selections that are systematically caused by such inconsistent routing state.

Cause 1: Partial termination of flooded route discovery by feedback generation. Feedback generation is the core routing primitive that informs the source about its candidate paths to the destination. However, partial termination of flooded route discovery at feedback-generating nodes is sufficient to force *many* nodes to select inferior paths leading to the source.

Specifically, partial flooding termination stops informing nodes that are deployed in the downstream direction of flooding, i.e., in the *terminated flood areas*, about entire subtrees of paths leading to the source. For such nodes, route selection problems occur when “hidden” paths are prefixes of their best paths leading to the source. In fact, such *victim nodes* then select inferior paths due to the unavoidable reception of other flooded routing information. In Figure 1, we illustrate such onset of inferior route selection in which feedback-generating

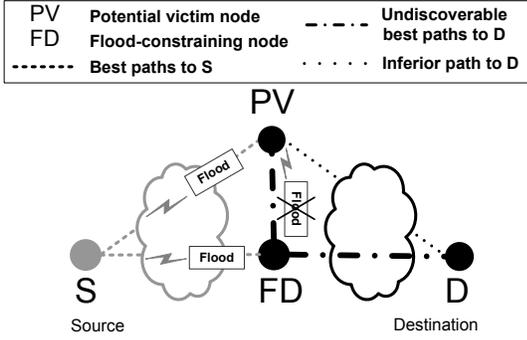


Fig. 2. Initialization of a potential inferior route selection of the node PV to the destination D , resulting from the overhead reduction at the node FD .

destination D hides the best path leading to the source S from the victim node V . Consequently, node V selects an inferior path whenever it receives any other flooded discovery information, in this example the one sent by node N .

Finally, note that the number of victim nodes can be arbitrarily large, depending on the impact of the terminated flood areas. Also, note that poorly selected paths can have any inferior metric-costs as well as arbitrarily long activity duration, because the best paths of potential victims remain hidden until a subsequent route discovery.

Cause 2: Direction-coupling property of route discovery.

Direction-coupling of route discovery is an overhead reduction property that unicasts the discovery feedback *only* over the paths that were previously selected during the flooded route discovery. Such a coupling property can suppress sufficient routing information to force *many* nodes to select inferior paths leading to the *destination*. Next, we describe two stages of such inferior selections: (1) initialization by the constrained flooding primitive, and (2) finalization by the unicast feedback primitive.

Exclusion of paths leading to the destination initializes inferior route selection during the constrained flooding phase of route discovery. The constrained flooding primitive excludes from route discovery any paths that cannot help the destination to identify a better path to the source. However, this primitive ignores the fact that these “excluded” paths can be best paths of other nodes leading to the destination. Then, due to the coupling property of route discovery, the “excluded” best paths are also not reported in the feedback phase of route discovery and consequently the nodes will never be able to identify them. Note that such exclusion of paths does *not* force any inferior route selection in itself, but it does initiate such selections. Therefore, we refer to the nodes having their best paths excluded as the *potential victim* nodes. In Figure 2, we illustrate an initialization of an inferior route selection for the potential victim PV whose flooded discovery packet gets discarded at PV 's best next-hop FD . Consequently, PV 's best path to the destination D , i.e., the path $PV-FD-\dots-D$, becomes excluded from further route discovery.

Reception of the discovery feedback over any inferior paths finalizes inferior route selection for the potential victims whose best paths to the destination were initially excluded from route discovery. Ideally, this would never occur if all feedback was forwarded *only* over the true best path of the source-destination node pair. However, in operational networks this is often not the case, because nodes often have to generate

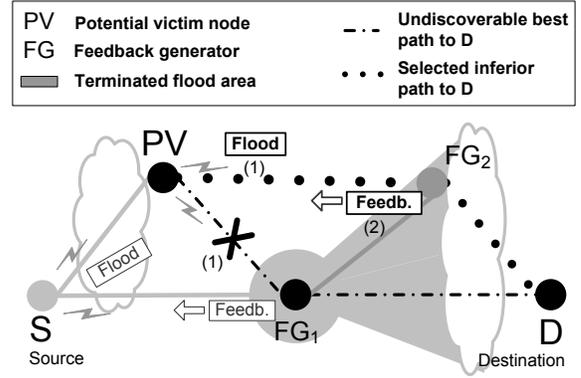


Fig. 3. Finalization of inferior route selection of the potential victim PV occurring due to reception of an inferior discovery feedback.

feedback over inferior paths before being informed about the true best path. This occurs because best-path information may be delayed (e.g., due to longer hop-count distance), lost (e.g., due to packet collisions), or systematically hidden (e.g., due to creation of the terminated flood areas).

While inferior route selection induced by delays and losses occurs in any routing configuration, in Figure 3 we illustrate an example of such selection caused by the systematic hiding of paths in the configuration in which participating nodes generate feedback on behalf of the destination. In the illustrated scenario, the feedback-generating node FG_1 first *initializes* inferior route selection for the potential victim PV by excluding PV 's best path to the destination D , i.e., the path $PV-FG_1-D$. FG_1 also prevents the node FG_2 from identifying its best path to the source, i.e., the path FG_2-FG_1-S , by covering that path with the terminated flood area (see Cause 1). Second, not knowing about its best path to the source, node FG_2 forwards discovery feedback over its inferior path FG_2-PV-S . This path coincides with an inferior path for potential victim PV , i.e., the path $PV-FG_2-D$. Therefore, reception of this feedback *finalizes* inferior route selection for PV .

Finally, note that the source is also a potential victim, because it also receives feedback over its inferior paths. However, barring packet losses, the source will be eventually informed about its best path to the destination, because the node-pair discovery primitives do not prevent propagation of best-path information for the source-destination node-pair. On the other hand, any other potential victims may not be informed about their best paths. Inferior route selections for these nodes can have any inferior metric-costs and arbitrarily long durations.

Cause 3: Routing sequence information as a supporting property of inferior route selection.

Inherently unequal updating of nodes with the most up-to-date routing sequence information occurring during the unicast feedback propagation supports inferior route selections. Due to space limitations, we present details in [9].

Discussion. In our analysis, we showed that inferior route selections can systematically occur for any nodes that participate in route discovery, except for the source, the destination, and the nodes connecting them via the least-cost path. Inferior paths can be systematically selected both to the source and to the destination. Therefore, problems of inferior route selection would be significant in any environments in which many nodes share similar end-points of communication, e.g. in the gateway-centric wireless mesh networks. Finally, note that

poorly selected paths can have any inferior metric cost and any duration that is only limited by an onset of a subsequent route discovery.

III. HISTORICALLY-ASSISTED IDENTIFICATION OF INFERIOR ROUTE SELECTION

In the previous section, we showed that inferior route selection results primarily from inherently insufficient distribution of best-routing information. Therefore, to target selection of best paths, extended route discovery with alternative information sources is needed to compensate for the presently missing information. To this end, the simplest solution would be to continually exchange best-path information for all nodes as is done in the distance-vector and link-state routing protocols. However, such a solution would induce a prohibitively large amount of overhead, thus causing disproportionate reductions of data throughput as shown in [10]. Instead, we propose a historic ranking principle enabling zero-overhead identification of inferior route selections. We employ this ranking as a core mechanism of our primitives for preservation of least-cost routing.

A. Historic ranking principle

Next, we define historic routing information, show how it can be used to identify inferior route selections by our historic ranking principle, and discuss necessary conditions for identification of such selection.

Definition 3.1: A node’s historic routing information is the collection of paths to each destination and their related metric costs reported during all prior route discoveries in which the node participated.

Maintaining historic routing information enables a node to gain an extensive view of candidate paths that may be hidden during individual route discoveries. Moreover, having an extensive view of candidate paths enables a node to rank the paths according to their previously reported metric costs. Such a *historical ranking* would help identify paths that are likely to become the current least-cost path. Consequently, during route discovery, failure to receive reports from such “likely optimal” paths provides an indication of a potentially inferior route selection.

For this historic ranking principle to correctly identify exposure to inferior route selection, two conditions must be met: (1) historically reported costs of a path must preserve the path’s high ranking, i.e., previously best paths must remain feasible candidates for the present best path, and (2) the present best path must have been previously reported to the node.

These conditions are justified for the hop-count path metric merely by the static nature of the network, i.e., the topology is largely unchanging. Moreover, for performance-based path metrics such as [11], [12], [13], the fixed topology should also yield persistence in path performance. For example, if a link suffers from low throughput due to having a long inter-node distance, this condition will not change. Similarly, low throughput can also be caused by persistent contention with a (non-mobile) hidden terminal. Nonetheless, we evaluate and validate these conditions via experiments on an operational wireless mesh network (see Section V).

B. Implementation of the historic ranking principle

The concept of historically-based path ranking can be applied to any underlying routing algorithm. Here, we propose a ranking implementation for wireless distance-vector protocols such as IEEE 802.11s HWMP [14]. Specifically, to reconstruct best-ranked paths, each node ranks its next hops according to the *end-to-end* path costs reported previously for each destination.

Next, we formally introduce historic ranking principle adapted to the wireless distance-vector routing as illustrated in Figure 4.

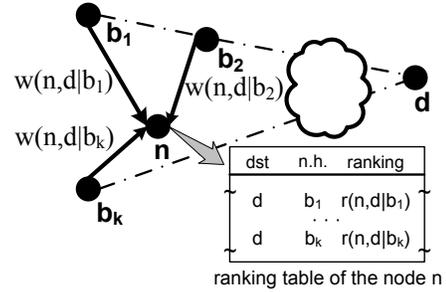


Fig. 4. Illustration of the historically-based path ranking that enables a node (node n) to infer which of its neighbors $b_i \in B(n)$ is a potentially best next-hop to a given destination (destination d).

Let $n \in N$ be a routing node, and $B(n)$ be a set of its potential next-hops. Let $w(n, d|b)$ be the path metric cost, i.e., the total cost of on-path links from node n through its next-hop $b \in B(n)$ to the destination $d \in N$. To rank likelihood that the next-hop b can provide the best path to the destination d and to smooth ranking indications, we utilize an Exponentially Weighted Moving Average (EWMA) filter to determine the ranking value $r_i(n, d|b)$:

$$r_i(n, d|b) \equiv (1 - \alpha) w_i(n, d|b) + \alpha r_{i-1}(n, d|b)$$

$$r_1(n, d|b) \equiv w_1(n, d|b)$$

In this definition, $i \geq 2$ represents the total number of route discoveries in which neighbor b reported a path metric cost $w(n, d|b)$, and α is a configurable coefficient.

By this ranking policy, neighbor $b_o(n, d)$ is ranked as a likely *best* next-hop to the destination $d \in N$ if $b_o(n, d) = \arg \max_{b \in B(n)} (r(n, d|b))$. The ranking of other nodes is performed accordingly. Finally, determining a threshold $T_H(r(n, d))$ that limits the number of next-hops nodes belonging to the highly-ranked set $H(r(n, d)) \supseteq \{b_o\}$ is an implementation decision.

IV. DETER AND RESCUE: HISTORICALLY-ASSISTED AVOIDANCE OF INFERIOR PATHS

Having the historically-based assessments of paths that are likely to become the present metric optimal paths, we next introduce two routing primitives: the DETER primitive that enables *prevention* of and *recovery* from inferior paths, and the RESCUE primitive that enables collaborative *restoration* of best paths.

A. Historically-assisted DETER primitive

The DETER primitive is a set of historically-assisted mechanisms a node itself employs during the route discovery in order to ensure selection of its least-cost path. To this end, the

primitive addresses two main causes of inferior route selection (see Section II): (1) delayed arrivals of best routing information inducing potentially irrecoverable inferior routes due to the direction-coupling property, and (2) lost or systematically undelivered best routing information preventing any least-cost routing.

“Wait for the historically-best neighbors” is an informed zero-overhead prevention mechanism that counters delay of best routing information. Relying on the historic ranking, the mechanism *adaptively* delays forwarding of any route-discovery reports until a node receives reports from its best-ranked next-hops for a given destination. Such delayed forwarding not only prevents inferior route selection, but also reduces routing overhead and prevents pollution of a node’s historical ranking at its neighbors, which would otherwise occur due to propagating inferior routing information.

During the adaptive delay interval, a node $n \in N$ identifies its presently reported best next hop $b(best) = \arg \max_{b_i \in B(n)} (w(n, d|b_i))$, and verifies that route-discovery reports are received from the best-ranked next-hops $b_k \in H(r(n, d))$. If all such reports are received, the node infers that it has sufficient information to perform least-cost route selection. It then selects a path via its *presently* best next-hop $b(best)$, and it advertises a report about that path to the route discovery process. Otherwise, the node waits for such reports until the expiration of a predetermined *threshold* interval $W(t)$, and subsequently initiates the inquiry mechanism.

“Historically-best neighbors inquiry” is a historically-informed low-overhead prevention and recovery mechanism that counters undelivered best routing information. Relying on the historic ranking, a node $n \in N$ would employ this mechanism to send a query to a subset of its best-ranked next-hops, asking them for their undelivered route-discovery reports. The node would then potentially re-select its path based on the comparison of its presently best path-metric cost reported by $b(best)$ and the path-metric costs reported via recovery responses. If a better path is received, the node would advertise it to the route discovery process. This ensures that the best path is indeed determined by the *present* best path-metric costs.

B. Historically-assisted RESCUE primitive

The RESCUE primitive is a historically-assisted effort of *neighboring* nodes to offer better paths to the node exposed to a potentially inferior route selection. The primitive is activated during or immediately after route discovery when all nodes have selected their paths. While ideally RESCUE would be unnecessary, it becomes crucial when a node cannot recover itself to the best path due to effects such as losses of recovery packets or a node’s insufficiently trained historical ranking. Next, we describe the RESCUE primitive and address several new challenges arising from its application: (1) each node’s general incapacity to know paths of other nodes, (2) potentially significant generation of rescue-related overhead, and (3) potential routing instabilities caused by rescue.

We address incapacity of nodes to know each others selected paths by opportunistically *initiating* the rescue based only on the *local* avoidance of bottleneck links. Accordingly, a neighbor $b_j \in B(s_i)$ (see Figure 5) initiates rescue when it infers that it can offer less constrictive link to the node

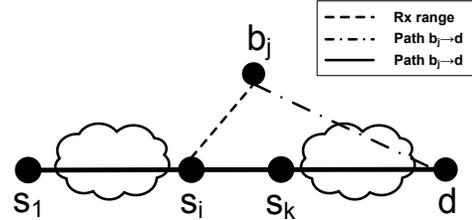


Fig. 5. Recovery from inferior path selections. Neighboring node b_j offers re-routing to node s_i upon its inference of path qualities.

$s_i \in N$, i.e., when $w(b_j, s_i) \leq w(b_j, s_k)$. We experimentally verified this method of recovery initiation in the TFA mesh network [7]. Our results (not presented in this paper) indicate that such initiation *does not* prevent better paths from being identified. However, this may produce unnecessary overhead when the total cost of an offered path is higher, i.e., when $w(s_i, d|b_j) \geq w(s_i, d|s_k)$. To constrain such overhead, we limit the number and the interval of rescue retries.

Finally, to prevent routing instabilities that may occur due to fluctuations of present path metric costs during the rescue interval, we further rely on historic ranking of paths. Specifically, the node decides to continue processing a rescue offer only when the offer is sent by historically higher ranked next-hop, i.e., if $r(s_i, d|b_j) > r(s_i, d|s_k)$. Note that this decision does not require an offering node to belong to the highly ranked set of neighbors $H(r(n, d))$, enabling the RESCUE primitive to overcome limitations of the over-constrained or insufficiently trained highly-ranked sets. Subsequently, the node reselects its current path if the offered path has *presently* lower metric-cost, i.e., if $w(s_i, d|b_j) > w(s_i, d|s_k)$.

V. EVALUATION

In this section, we use experiments and measurements on an operational mesh network, as well as simulations to (1) explore the severity of inferior route selection under the node-pair discovery primitives, (2) study the persistence of routing metrics, a key requirement for historically-assisted route selection, and (3) evaluate the ability of the deter and rescue primitives to overcome inferior route selection.

A. Evaluation platforms

The **TFA network** is an operational wireless mesh network that provides Internet access to a residential area of 3 km². At the time of our measurements, the network consisted of 17 statically deployed backhaul nodes and a single Internet gateway (see Figure 6). Each node is equipped with a 15 dBi omnidirectional antenna elevated on a 10 meter pole to provide both the access point functionality and wireless routing. Additionally, a directional link operating over a separate radio channel between nodes 12 and GW serves as a throughput-increasing and interference-avoiding network resource.

TFA replica is a simulation environment that enables a study of routing effects at much finer granularity than possible in TFA itself, providing per-packet observability and perfectly synchronized time references at all nodes. We employ the replica in the ns-2 simulator, which we significantly extend and configure with parameters comprehensively measured in the TFA network. For example, we replicate the signal-strength coverage of each node, the range of signal-strength variation of each link, capture properties of each pair of links, etc.

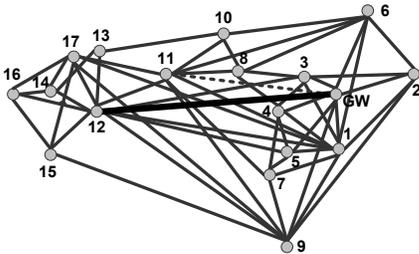


Fig. 6. TFA topology and connectivity map.

Routing protocols. Both in TFA and the replica, we used the routing protocol that employed the node-pair discovery primitives, the SINR threshold of usable link quality, the application-layer gateway announcements, and the distance-vector accumulation of path-metric costs. In our evaluations of the historically-assisted routing performed on the TFA replica, this protocol was enhanced with our deter and rescue primitives.

B. Inferior route selections under node-pair routing

Here, we measure the main aspects of inferior route selection identified by our analysis: (1) the occurrence of inferior paths, (2) the time intervals during which such paths remain selected, and (3) the detrimental effects of inferior route selection on throughputs attained by several nodes.

Methodology. In our first set of experiments, we choose a representative high-quality path (17-12-GW) and evaluate the ability of “node-pair routing” to correctly select this path. In this evaluation, we employ the hop-count metric that provides us with a firm reference of path optimality, because the hop-count cost of individual paths is time-invariant, barring link and node outages. To this end, we ensure that the targeted path 17-12-GW was the least-cost path and that it was available during all our experiments. We extract this information from the connectivity records reported by each node’s wireless card driver. This first set of experiments is performed during the normal TFA operation with the 5 second granularity during three one-hour intervals, each day for 5 days.

In our second set of experiments, we evaluate the severity of inferior route selection by observing achievable throughputs attained by the “node-pair routing” and by the statically configured best paths (17-12-GW, 16-12-GW, 11-12-GW) that we identified by a large set of preliminary measurements. In our evaluation, we measure throughputs of fully-backlogged TCP flows that contend with the network management traffic of all nodes, and with the traffic of in-home networks outside of our administrative control. We gain confidence in our results by conducting numerous 5-minute measurements over several days.

Results: Occurrence, duration, and impact of inferior paths. Figure 7 depicts the hop-count costs and durations of selected paths from the observed node 17 to the GW. The results indicate that only 3 out of 12 selected paths are metric (hop-count) optimal, even though the targeted best path 17-12-GW (denoted as $2H_1(Opt)$) was available during all our measurements. Moreover, four out of nine non-metric-optimal (i.e., inferior) paths have metric costs that are more than twice that of the available best path. These results confirm our analytical findings about the arbitrary costs of inferior paths.

Figure 7 also indicates that although the best path $2H_1(Opt)$ has the longest average duration (max. duration being 27.4 minutes), the duration of inferior paths can be arbitrarily long, i.e, spanning from 5 seconds (our measurement granularity) to tens or hundreds of seconds. This result also confirms our analytical finding about the potentially long durations of inferior paths limited only by an onset of new route discoveries. In our case, these discoveries occur when the on-path nodes infer that their existing path is broken, or when other nodes initiate route discovery.

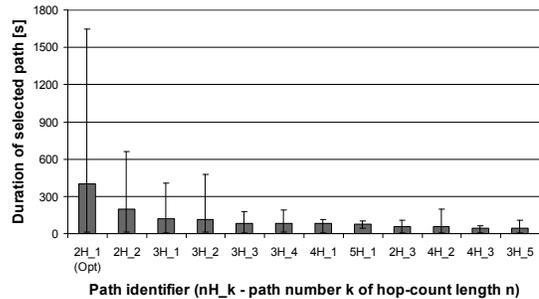


Fig. 7. Durations and hop-count lengths of paths selected by our reference node 17. Presented are average, minimal, and maximal durations.

Table I shows the impact of inferior route selection on the achievable throughputs of nodes. The results indicate that depending on the amount of time a node spends on its inferior paths, poor route selection can reduce achievable throughput significantly, from approximately 20% (node 17) to 80% (node 11). Moreover, note that in the TFA network such inferior paths also negatively impact throughputs of neighboring paths, because the inferior paths fail to include the capacity-injecting link 12-GW and subsequently increase contention with the neighboring paths in the critical gateway area.

| Avg Throughput [kb/s] | 11 to GW | 16 to GW | 17 to GW |
|---------------------------|----------|----------|----------|
| Node-Pair Route Selection | 401.63 | 847.86 | 2075.36 |
| Static Best Path | 2143.11 | 1723.02 | 2570.24 |

TABLE I
AVERAGE THROUGHPUTS ATTAINED BY THE NODE-PAIR ROUTE SELECTION AND BY STATICALLY SET BEST PATHS.

Finally, we note that the targeted best path 17-12-GW is not only optimal in the hop-count metric, but also in the SINR metric and any throughput-related metric. Therefore, similar inferior route selections would be observed if these other metrics were employed.

C. Prerequisites for historically-assisted mesh routing

Here, we measure the existence of persistent network properties that would enable our historic ranking principle to identify potentially inferior route selections. We are particularly interested in the properties that would support the ranking according to performance-based metrics [11], [12], [13].

Methodology. To understand if paths can be ranked according to the performance-based metrics in mesh networks, we evaluate two dominant impacts of the wireless environment on such metrics: (1) the impact of wireless channels, and (2) the impact of contention and interference. We encapsulate evaluation of these two impacts in a single type of throughput measurement by employing fully backlogged TCP traffic. The

throughput of such traffic is known to be highly sensitive to the both analyzed impacts [15].

First, as a baseline, we measure whether persistent throughput ranking can be achieved for isolated paths under the inherently variable conditions of wireless channels. To minimize any traffic-related impacts on our results, we prevent traffic generation at all TFA nodes that do not belong to the observed path. Second, we experimentally validate persistent ranking of contending wireless links, which are the core elements of path-metric accumulation. In these measurements, we also expose the observed links to additional contention and interference produced by *all* other TFA nodes that generate regular network management traffic (e.g., SNMP, connectivity maintenance, gateway announcements, etc.). Our sets of measurements consist of numerous 5-minute *iperf* throughput tests conducted during 5 days.

Results: Existence of historically persistent ranking properties. In Figure 8, we illustrate how the inherently variable wireless channel impacts ranking of isolated paths. Specifically, we show throughput properties of five most frequently selected paths by the node 17. The results indicate that although the standard deviations of throughputs can be significant (up to approx. 300 kb/s), the average throughput of paths remains sufficiently persistent to enable historical ranking.

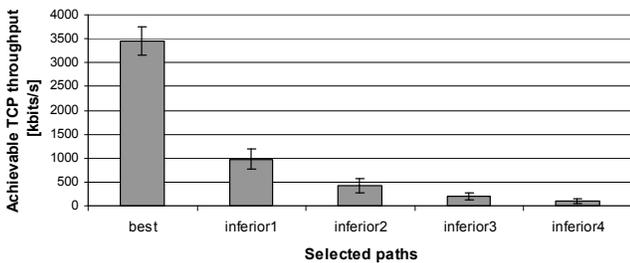


Fig. 8. Persistence of throughputs of isolated paths connecting our reference node 17 to the GW. Presented are average values and standard deviations.

In Figure 9, we show the impact of contention and interference on the historical ranking. Specifically, we measure throughput properties of 4 gateway links in isolation, in contention of all pairs, and in contention of all triplets. Moreover, some of these links act as hidden terminals (e.g., 7-GW and 2-GW; 7-GW and 6-GW), and some can capture over the others (e.g., 1-GW over 7-GW). Our results indeed indicate highly variable throughputs within and between contention scenarios. However, the properties supporting the historical ranking still exist. First, in each experimental scenario, average throughput does indicate the ranking of links. Second, using results from all contention scenarios, a node could identify its links that would best contribute to high ranking of paths in varying contention scenarios of real networks. For example, in the presented measurements the two highest-ranked links are persistently 1-GW and 2-GW. Note that since the metric-costs of paths represent accumulation of the metric-costs of the on-path links, our measurements indicate that paths can also be successfully ranked under contention and interference.

Finally, we note that the network traffic properties would also impact the ranking of paths during regular network operation. Although this may perturb the ranking order of specific paths, our results indicate that key ranking components

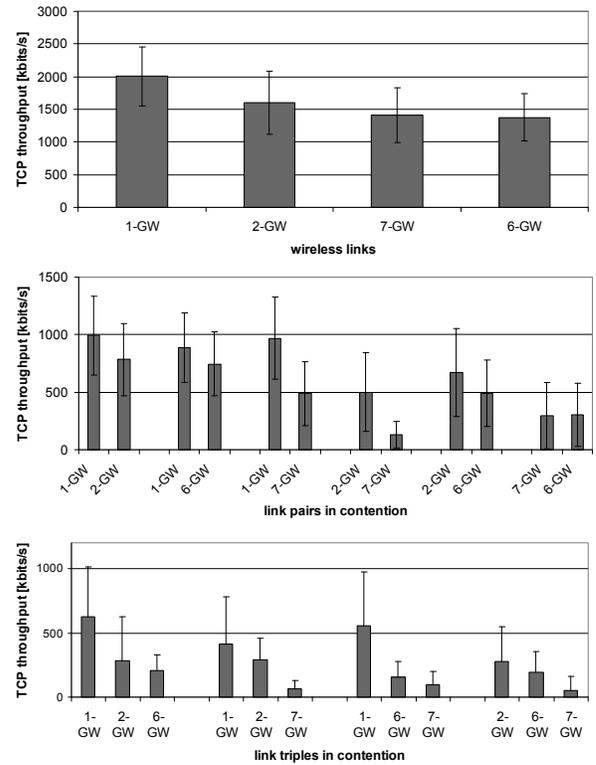


Fig. 9. Persistence of throughputs of contending links measured over 5 days. Presented are average values and standard deviations.

are persistent. Therefore, a subset of highly-ranked paths would exist, which supports application of our historic ranking principle.

D. Evaluation of historically-assisted routing

Here, we evaluate the ability of the deter and rescue primitives to overcome inferior route selection. Specifically, we show: (1) the metric costs of selected paths, (2) the contribution of each primitive towards the restoration of least-cost paths, and (3) the amount of additional overhead induced to enable such restoration.

Evaluation setting. To understand a broad set of factors that are not observable in the operational network, we employ the TFA replica to evaluate the deter and rescue primitives. In the preliminary evaluation, we validated that the least-cost paths in the replica are similar to ones in the TFA network, and that inferior paths occur as predicted by our analysis.

Next, we configure the historical ranking to only identify a single best path per-destination at each node. We allow the RESCUE primitive to attempt restoration of optimal paths in three equally spaced attempts during the 9 seconds interval following each route discovery. Finally, to enable identification of high-throughput paths, we employ a combination of the ETT metric [12], the SINR metric, and the hop-count metric.²

Methodology. We observe route selections of three nodes (11, 16, 17), for which we identified best paths (11-12-GW, 16-12-GW, 17-12-GW) by preliminary simulations. To invoke route re-discoveries, we gradually increase (in 30 sec intervals) the likelihood of packet loss, which is generally

²Detailed discussion of how the employed metrics are coordinated to identify high-throughput paths is beyond the scope of this paper.

used as a re-routing initiator (e.g., see [8], [16]). In particular, we increase the number of contending nodes in sequence $17 \rightarrow 16 \rightarrow 11 \rightarrow 3 \rightarrow 2$, in which some nodes act as hidden terminals (11-16, 16-2, 17-2). Similar to our TFA measurements we employ fully backlogged TCP traffic. We emulate conditions of real networks by initially randomizing historic rankings of each node by *randomly* starting TCP flows of all TFA nodes in the preliminary phase of each simulation. Our results are extracted from 50 experiment runs, each lasting 300s.

Result: Least-cost property of selected paths. In Table II, we list *all* selected inferior paths and their durations. The results indicate that our routing primitives enable dominant selection of least-cost paths. Specifically, during the thousands of seconds of the simulation time, two observed nodes (nodes 17 and 16) selected only few sub-second lasting inferior paths. The third observed node (node 11), which was always activated during a severe exposure to its readily active and fully backlogged hidden-terminal (node 16), did select a number of inferior paths. However, although this node experienced severe losses of its best-path information, our primitives did enable it to restore its best path, often in sub-second intervals. Finally, note that all inferior paths listed in Table II would be irrecoverable and arbitrarily long-lasting under the node-pair route selection, because they were all caused by undelivered best routing information.

| Path | min(T)[s] | max(T)[s] | avg(T)[s] | N | # flows |
|------------|-----------|-----------|-----------|----|---------|
| 17 11 0 | 0.0043 | 0.0194 | 0.0102 | 6 | 1 |
| 17 1 0 | 0.0106 | 0.0736 | 0.0230 | 5 | 1 |
| 17 1 0 | 0.0062 | 0.0127 | 0.0119 | 4 | 2 |
| 17 11 0 | 0.0038 | 0.0092 | 0.0073 | 3 | 2 |
| 16 15 9 0 | 0.0447 | 0.0447 | 0.0447 | 1 | 1 |
| 16 17 1 0 | 0.0444 | 0.0444 | 0.0444 | 1 | 2 |
| 11 0 | 0.0120 | 65.0204 | 8.0253 | 35 | 1 |
| 11 17 12 0 | 0.0032 | 7.0011 | 2.0346 | 10 | 1 |
| 11 1 0 | 0.1983 | 18.7345 | 9.4664 | 2 | 1 |
| 11 3 0 | 0.0165 | 12.0243 | 4.2336 | 4 | 1 |
| 11 0 | 0.0164 | 6.0462 | 2.7022 | 10 | 2 |
| 11 17 12 0 | 3.0105 | 3.0105 | 3.0105 | 1 | 2 |

TABLE II

DURATION (T) AND NUMBER (N) OF INFERIOR ROUTE SELECTIONS FOR A GIVEN NUMBER OF ACTIVATED ON-PATH TCP FLOWS (#FLOWS).

Result: Individual contribution of the historically-assisted primitives toward suppression of inferior paths. Although the DETER primitive was itself sufficient to prevent any occurrence of inferior paths in many route selections, here, we only analyze suppression of inferior paths that could not be avoided (see Table II). Our per-packet analysis revealed that exposure to packet losses significantly impacts the ability of each primitive to suppress inferior paths.

We identified that nodes *not* exposed to a severe impact of loss-related problems generally rely on the DETER primitive to re-route themselves from inferior paths. For example, node 17 relied on this primitive in 12 out of 18 successful best-path recoveries (see Table II). Moreover, for such nodes, we identified that the delay of best-path information causes most inferior route selections. However, our results indicate that all these problems can be resolved within sub-second intervals.

On the other hand, a node exposed to severe packet loss does not generally rely on the DETER primitive, because its requests for recovery are likely to experience collisions. For such

node, the neighbor-initiated RESCUE primitive dominantly helps identification of best paths, because the exposure of neighbors to loss-related problems (such as hidden-terminals) is generally different than the one of the node itself. In fact, the RESCUE primitive helped the loss-prone node 11 to recover from 57 of its 62 inferior route selections. In Table III, we provide the timing profile of these 57 recoveries indicating a significant percent of fast (sub-second) restorations of optimal paths.

| Recovery interval T | Percent of recovered paths |
|---------------------|----------------------------|
| $T < 1s$ | 43.9% |
| $T \in [1s, 3s)$ | 15.8% |
| $T \in [3s, 6s)$ | 14.0% |
| $T \in [6s, 9s)$ | 10.5% |
| $T > 9s$ | 15.8% |

TABLE III

TIMING-PROFILE OF RESCUES FROM INFERIOR PATHS OF THE NODE 11, WHICH WAS EXPOSED TO THE SEVERE LOSSES OF ROUTING INFORMATION.

However, inferior paths lasting longer than the 9 seconds rescue interval (see Table III) challenge the ability of our primitives to limit the duration of poorly selected paths. We analyzed all 9 such selections, discovering that their long duration is caused by the well known effect of TCP traffic outages [15]. These outages prevented the neighbors from monitoring poorly selected paths, thus also preventing them from sending their rescue offers. To confirm that our primitives are indeed able to identify inferior route selections, we performed additional experiments using 2 TCP flows per node, which reduced the likelihood of on-path traffic outages. In this setting, our results (see Table II) confirm that rescue always occurs within the configured 9 seconds rescue interval.

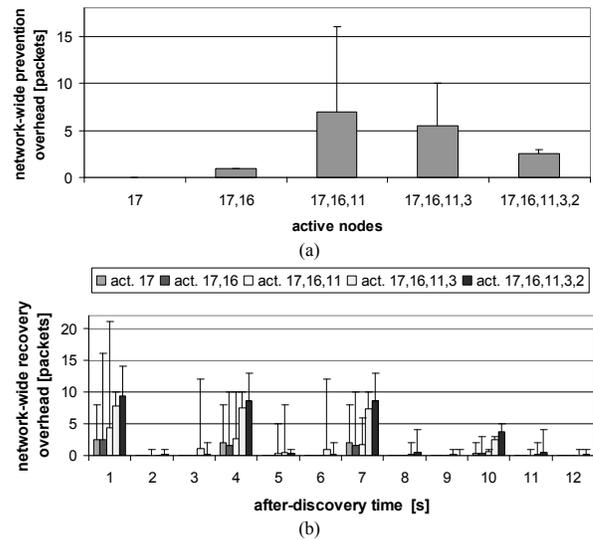


Fig. 10. Total overhead of historically-assisted primitives generated by all nodes: (a) DETER primitive, and (b) RESCUE primitive. Presented are average and maximal values.

Result: Overhead cost of historically-assisted routing. The results in Figure 10 indicate that the total generated network overhead adapts to the increased likelihood of inferior route selection. In particular, the DETER primitive generates

most recovery requests when a node is most likely to lose its route discovery packets, e.g., due to an exposure to hidden-terminals. The average generated overhead is at most 7 pkts for the entire network. The RESCUE primitive also adaptively reduces its overhead as best paths are being recovered. The rescue overhead is somewhat larger because it is sent opportunistically.

VI. RELATED WORK

Node-pair discovery primitives. The node-pair route discovery primitives have been employed and studied in both mobile ad hoc networks (MANETs) (e.g., see [17], [18] and the references therein) and static mesh networks (e.g., see [6], [5], [19], [20] and the references therein).

Fortunately, *MANETs* are largely immune to the effects of inferior route selection that we considered in this paper. First, most MANET studies considered random disjoint pairs of mobile sources and destinations in which nodes mostly act as the end-points of route discoveries. As we showed in Section II, it is not the endpoints that are vulnerable to systematically inferior route selections, but rather the participating nodes. Second, if an inferior path is selected regardless, node mobility will limit the lifetime of the route and thus limit the duration and penalty of an inferior route selection. Indeed, MANET routing protocols were shown to correctly identify minimum-hop paths for a high percentage of route selections [17].

In contrast, studies of routing in *mesh networks* have utilized node-pair discovery primitives but focused on link metrics and high throughput routing [5], [6], [14], [20]. However, while new protocols demonstrated throughput gains, no study assessed whether these protocols selected routes that were actually inferior, i.e., whether a better route existed at the time of route selection. Indeed, we showed that all protocols employing the node-pair discovery primitives will always be systematically prone to inferior route selection irrespective of any individual protocol implementation, or any employed routing metric.

Historic routing information. At one level, all routing protocols employ some use of history and memory in their decisions making, e.g., route caching [21] and time-averaging of routing metrics [11]. In contrast, while route caching utilizes past route selections to reduce overhead by avoiding new route discoveries, we jointly consider current routing information and historical ranking of paths. Consequently, our historically-assisted routing enables both selection of presently best paths and identification of potentially inferior paths.

VII. CONCLUSION

In this work, we showed that widely deployed node-pair routing does not distribute sufficient routing information and prohibits many nodes from selecting their least-cost paths. This leads to selection of paths that have arbitrarily higher metric-costs and arbitrarily long durations. Moreover, we confirmed this result with a large set of measurements on an operational network, showing that inferior paths can last for 10's or 100's of seconds. To solve this problem, we proposed a zero-overhead identification of potentially inferior paths based on the historic ranking of paths. Our measurements showed that such identification is feasible due to static wireless networks having a number of historically persistent properties.

Moreover, we developed a set of historically-assisted routing primitives that help avoid route-selection problems caused by the lack of best-path information. Conducting a per-packet evaluation of the new primitives, we showed that they help nodes to largely avoid inferior routes; when inferior routes are selected regardless, their duration is often reduced to sub-second time scales.

VIII. ACKNOWLEDGEMENTS

This research was supported by NSF grants CNS-0751173, CNS-08102501, and by the Cisco Collaborative Research Initiative. The authors are grateful to Jan Krays (OakTree Wireless Consulting) and Santosh Pandey (Cisco Systems) for insightful comments and advice that significantly improved this paper.

REFERENCES

- [1] Motorola Incorporates Mesh Spec. Available at: <http://www.wi-fiplanet.com/news/article.php/3600221>.
- [2] MikroTik RouterOS. Available at: <http://www.mikrotik.com/>.
- [3] Cisco Aironet 1500 Series Wireless Mesh AP Version 5.0 Design Guide. Available at: <http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP.html>.
- [4] LocustWorld. Available at: <http://www.locustworld.com/>.
- [5] K. Ramachandran, M. Buddhikot, G. Chandranmenon, S. Miller, E. Belding-Royer, and K. Almeroth, "On the design and implementation of infrastructure mesh networks," in *IEEE WiMesh*, Santa Clara, CA, USA, September 2005.
- [6] S. Mir, A. A. Pirzada, and M. Portmann, "Hover: Hybrid on-demand distance vector routing for wireless mesh networks," in *ACSC2008*, Wollongong, Australia, January 2008.
- [7] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *ACM MobiSys*, Uppsala, Sweden, June 2006.
- [8] C. Perkins, E. Belding-Royer, and S. Das, "RFC 3561: Ad hoc on demand distance vector (AODV) routing," July 2003.
- [9] S. Miskovic and E. Knightly, "Routing primitives for wireless mesh networks: Design, analysis and experiments," Rice University, Houston, TX, Tech. Rep. TREE0904, Available: <http://www.ece.rice.edu/~smm1/doc/rtprimitives.pdf>, July 2009.
- [10] J. Camp, V. Mancuso, O. Gurewitz, and E. Knightly, "A measurement study of multiplicative overhead effects in wireless networks," in *IEEE INFOCOM*, Phoenix, AZ, USA, April 2008.
- [11] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *ACM MobiCom*, San Diego, CA, USA, September 2003.
- [12] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *ACM MobiCom*, Philadelphia, PA, USA, September 2004.
- [13] T. Salomidis, M. Garetto, A. Saha, and E. Knightly, "Identifying high throughput paths in 802.11 mesh networks: a model-based approach," in *IEEE ICNP*, Beijing, China, October 2007.
- [14] J. Camp and E. Knightly, "The IEEE 802.11s extended service set mesh networking standard," in *IEEE Comm. Magazine*, August 2008.
- [15] J. Shi, O. Gurewitz, V. Mancuso, J. Camp, and E. Knightly, "Measurement and modeling of the origins of starvation in congestion controlled mesh networks," in *IEEE INFOCOM*, Phoenix, AZ, USA, April 2008.
- [16] D. Johnson, Y. Hu, and D. Maltz, "RFC 4728: The dynamic source routing protocol (DSR) for mobile ad hoc networks for IPv4," February 2007.
- [17] J. Broch, D. A. Maltz, D. B. Johnson, Y. C. Hu, and J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols," in *ACM MobiCom*, Dallas, TX, USA, October 1998.
- [18] S. R. Das, C. E. Perkins, E. M. Royer, and M. K. Marina, "Performance comparison of two on-demand routing protocols for ad hoc networks," in *IEEE INFOCOM*, Tel-Aviv, Israel, March 2000.
- [19] K. Ramachandran and I. Sherif and E. Belding-Royer and K. Almeroth, "Routing Stability in Static Wireless Mesh Networks," in *PAM 2007*, Louvain-la-neuve, Belgium, April 2007.
- [20] A. Subramanian, M. Buddhikot, and S. Miller, "Interference aware routing in multi-radio wireless mesh networks," in *IEEE WiMesh*, Reston, VA, USA, September 2006.
- [21] Y. C. Hu and D. B. Johnson, "Caching strategies in on-demand routing protocols for wireless ad hoc networks," in *ACM MobiCom*, Boston, MA, USA, August 2000.