

TAF: A Temporal Adaptation Framework for Hybrid Routing in Mobile Ad Hoc Networks

Venugopalan Ramasubramanian

Emin Gün Sirer*

Department of Computer Science,
Cornell University,
Ithaca, NY 14853.
{ramasv,egs}@cs.cornell.edu

Abstract

A central challenge in ad hoc networks is the design of routing protocols that can adapt their behavior to frequent and rapid changes at the network level. Choosing between reactive, proactive, or hybrid routing regimes and selecting appropriate configuration parameters for a chosen protocol are difficult tasks. This paper introduces a framework, called TAF, for seamlessly adapting between proactive and reactive routing protocols. This general framework enables a proactive and reactive protocol to coexist on the same network, provides a low-overhead mechanism by which these two routing strategies can be combined at fine grain and proposes an analytical model for automatically adjusting protocol parameters. Combined, this mechanism and model enable a protocol within our framework to find a near-optimal mix of proactive and reactive routing strategies for the mobility rate and traffic patterns observed on the network. We examine the application of this temporal adaptation framework to the construction of three specialized ad hoc routing

protocols. These protocols minimize packet overhead, achieve a targeted loss rate, and minimize routing latency using the TAF framework. In all three cases, hybrid protocols based on the TAF framework perform as well as or better than a proactive (TORA) and a reactive (AODV) protocol.

1 Introduction

Mobile networks are characterized by change[14]. Many of the diverse application areas for ad hoc networks, including emergency relief operations, battle-front applications and environmental data collection, exhibit a high degree of temporal or spatial variation. Nodes may join the network at any time, get disconnected as they run out of power and alter the physical network topology by moving to a new location. Link characteristics, such as bit error rates and bandwidth, might change due to external factors such as interference. And traffic patterns in the network might shift drastically as applications modify their behavior and redistribute load within the network. Consequently, a primary challenge in ad hoc networks is the design of routing protocols that can adapt their behavior to rapid and frequent changes seen at the network level.

*This work is supported in part by ONR grant N00014-01-0968. The views and conclusions contained herein are those of authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of these organizations or of the Government.

Many routing protocols have been proposed to address these challenges. Ad hoc routing protocols proposed to date fall between two extremes based on their mode of operation. *Proactive protocols* exchange routing information periodically between hosts, and constantly maintain a set of available routes for all nodes in the network. *Reactive protocols*, on the other hand, delay route discovery until a particular route is required, and propagate routing information on demand in response to requests. Both proactive and reactive protocols have inherent advantages depending on the characteristics of the network and the observed traffic patterns. Proactive protocols can provide good reliability and low latency in the presence of high mobility in the network. However, they entail a high overhead and scale poorly with increasing numbers of participating nodes. In contrast, reactive protocols can achieve low routing overhead, but may also lead to increased packet loss when the topology changes frequently and may suffer from increased latency due to on-demand route discovery and route maintenance. Since the characteristics of a real-world network vary dynamically with time, choosing an appropriate routing protocol is a difficult deployment decision. A protocol suited for a given mobility rate and traffic pattern may behave inefficiently as the mobility and communication patterns change. A fixed routing strategy represents a brittle decision embodied in the network, making it difficult to adapt to changing conditions.

In this paper, we present TAF (Temporally Adaptive Framework), a general, unified hybridization framework for seamlessly switching between proactive and reactive routing regimes. TAF enables both a reactive and a proactive routing protocol to coexist on the same network. TAF uses the proactive protocol to pre-calculate routes for a common destination at all nodes within the *proactive zone* of that host. The proactive zone is simply the set of surrounding nodes that are reachable within a given, and destination specific, number of hops of the destination node. This

zone enables destinations to create an area around them with constantly updated, available routes. Nodes outside this zone use a traditional reactive ad hoc routing algorithm to discover routes on demand. Unlike traditional reactive protocols, however, route requests need not be propagated all the way to a given destination under TAF. Any node at the boundary of the destination's proactive zone can respond to a route request and curtail a costly route request from propagating through the proactive zone.

The central insight behind TAF is that judicious adjustment of the proactive zone enables TAF-based protocols to find near-optimal trade-off between proactive route propagation and on-demand route discovery in an ad hoc network. This inherent trade-off is one of increased overhead for proactive information dissemination versus reduced latencies and loss rates stemming from pre-computed partial routes within a zone. TAF provides a natural integration between proactive and reactive regimes by adjusting the size of the proactive zone. A proactive zone of size zero corresponds naturally to a purely reactive protocol, while a zone whose radius equals the network diameter corresponds to a purely proactive protocol. TAF provides an analytical model and a low-overhead mechanism for determining the size of this zone, and thus finds a near-optimal combination of proactive and reactive routing for the given network topology, link characteristics, and traffic pattern. It constantly measures these metrics, and adapts the proactive zone of each node to reflect the best trade-off.

Ideally, A framework for hybrid routing protocol construction would exhibit the following properties:

- **General-purpose:** The framework should accommodate many different kinds of reactive and proactive routing protocols. The framework should enable the construction of protocols for optimizing diverse network metrics.

- **Effective:** Protocols based on the hybridization framework should perform as well as the better of the reactive and proactive routing protocols.
- **Efficient:** The framework should not require excessive communication overhead, latency or power and bandwidth consumption. The framework should enable nodes to make decisions independently, without requiring costly operations such as distributed consensus.
- **Adaptive:** The framework should enable protocols to adapt readily to changing network topologies, link characteristics and traffic patterns.
- **Multiprotocol/Multimetric:** Different nodes in the network should be able to pursue disparate goals. Each node should be able to adjust the routing protocol optimizations to serve its service requirements.
- **Backwards Compatible:** The hybridization framework should use well-studied, off-the-shelf components wherever possible. It should be compatible with existing standards.

This paper describes the TAF framework and makes the following contributions. First, it provides a novel, general-purpose, adaptive technique for hybridizing proactive and reactive routing algorithms desirably over time. The framework embodies a low-overhead mechanism for node management, and an analytical model to guide the fine-grain trade-off between competing routing regimes. It enables multiple nodes in the network to pursue disparate goals of optimization at the routing layer. Second, it describes the application of this framework to the construction of three separate protocols for minimizing packet overhead, reducing latency and achieving a target loss rate, while also optimizing other network parameters. Finally, it describes, through a simulation study and analysis, that the resulting protocols

are as good as or better than both purely proactive and purely reactive protocols. Overall, this paper demonstrates the case for hybrid, adaptive routing protocols, quantitatively showing that the ideal point for achieving an optimal packet overhead, loss rate, and latency resides at a varying point between fixed, purely reactive or purely proactive protocols. It shows that protocols built on top of the TAF framework perform well because they dynamically find configurations very close to that optimal.

The rest of this paper is organized as follows. In the next section, we discuss related work on unicast routing protocols, and place our hybridization approach in context. Section 3 presents our framework, outlines the analytical model that drives adaptation in TAF, and describes three TAF-based protocols for optimizing different, relevant metrics. Section 4 describes our implementation decisions and any changes we had to make to off-the-shelf protocols. Section 5 shows that the TAF framework leads to hybrid protocols that can outperform the better of the fixed routing regimes. We conclude in Section 7.

2 Related Work

While the vast majority of the routing protocols proposed to-date for ad hoc networks are purely reactive or purely proactive, some hybrid protocols have been proposed. We provide a brief overview below, and summarize how they differ from our framework. Overall, while the other hybrid approaches combine proactive routing with reactive routing, few attempt to explore the trade-off between the two, or adapt their parameters to best suit the observed mobility and traffic patterns on the network.

CEDAR [15], Core-Extraction Distributed Routing Algorithm, is a hybrid protocol that uses a core-extraction algorithm to partition the network spatially into neighborhoods around core nodes. These core nodes perform the packet forwarding

tasks in CEDAR, while they also maintain their topology through periodic broadcasts. CEDAR uses a QoS algorithm to compute the shortest widest path between the set of core nodes on a given path. CEDAR periodically invokes a stable distributed agreement algorithm to compute the core, but the core is modified only in response to topology changes.

ZRP [3], Zone Routing Protocol, is a hybrid routing protocol that divides the network into zones around each sender. Proactive routing is used within zones, while a reactive routing algorithm is used to propagate inter-zone packets. Forwarding in ZRP is performed via bordercasting, where each node sends a packet to the nodes at the boundary of its zone. Unlike CEDAR, ZRP nodes are not spatially tiled; the zone decomposition is root-directed (determined relative to senders) and overlapping. Selection of the appropriate zone radius for optimal ZRP performance is a non-trivial task [11].

ZHLS [6], Zone-based Heierarchical Link State, is similar to ZRP in that it also is a hybrid approach based on the notion of a zone. ZHLS requires physical location information during zone decomposition, keeps the zone connectivity information in each node and once the protocol performs zone assignments, zone sizes do not vary dynamically.

HARP [9], Hybrid Ad-hoc Routing Protocol, is a hybrid protocol that combines proactive and reactive approaches. It relies on a distributed dynamic routing (DDR [8]) protocol for decomposing the network into zones. A set of forwarding nodes in each zone is responsible for communicating with nodes in other zones. HARP uses its own custom protocol for inter-zone routing, whose main goal is to reduce delays through early path maintenance.

ADV [1] is Adaptive Distance Vector algorithm that exhibits on-demand characteristics by varying the frequency and size of routing updates. While comparisons show that it performs better than AODV and DSR under high mobility, its per-

formance characteristics have not been compared to proactive protocols.

Some researchers [7] have examined supplanting reactive protocols with timer-directed route discoveries to produce backup routes prior to losing the primary link. Their protocol uses a fixed timer value across all nodes, which is determined offline from a past history of link failure statistics.

TAF differs from these approaches in several fundamental ways. First, TAF adapts in both the temporal and spatial domain to changing network conditions. In previous work, the regions in which proactive and reactive protocols are executed are specified once and for all at deployment, or computed in a separate, costly topology creation phase. In contrast, TAF actively varies the routing tradeoff in the *temporal* domain based on current network measurements, obviating a separate tuning or self-calibration step. This variation enables TAF to explore the tradeoff between proactive and reactive routing at fine granularity. Second, TAF enables each destination node in the network to pick its own parameters for optimization, and select the tradeoff best suited for its own needs. This support for multiple adaptive protocols in the same network is quite versatile. For instance, one TAF node can adapt the routing layer for reduced latency of access while another targets reliable delivery at a chosen loss rate. Third, previous work relies on explicit messaging for zone construction. In contrast, TAF nodes base their decisions on locally gathered information, and a novel timeout-based zone control scheme allows TAF zones to shrink and grow without excessive control and synchronization overhead. Finally, zone sizes are variable and dynamic in TAF, and depend on network traffic, link characteristics and amount of route reuse. These three metrics effectively capture the benefit to be gained from modifying the zone size. Previous work uses inelastic metrics, such as hop counts, in constructing zones, which limits the responsiveness of the routing layer to changes in the mobility rate and traffic pattern.

3 Approach

In this section, we describe the TAF framework for dynamic adaptation between proactive and reactive protocols based on the characteristics of the network. We discuss an analytical model that provides the insight behind the operation of this framework. We then describe three instances of applying this framework to the construction of specialized protocols. These protocols minimize packet overhead, achieve a target loss rate and reduce network latency, respectively.

The TAF framework adapts between reactive and proactive routing by dynamically varying the amount of routing information shared proactively. It does so by defining a *proactive zone* around each node. All nodes within this zone maintain routes proactively for a given destination. The node-specific *proactivity radius* defines the number of nodes in the proactive zone. Each neighbor at a distance less than or equal to the proactivity radius is a member of the proactive zone for that node. All nodes not in the proactive zone of a given destination use reactive routing protocols to establish routes to that node. The tradeoff and amortization opportunity rests on manipulating this radius appropriately. By increasing the radius, TAF can decrease the loss rate and the latency for route establishment, but will pay more in packet overhead to keep routes fresh in a larger zone. By decreasing the radius, TAF can reduce routing overhead as fewer nodes need to be proactively updated; however, it may pay more in route finding latency and experience higher loss rates. Using this tradeoff, TAF can act as a completely reactive protocol by setting the proactivity radius of all the nodes to zero. Conversely, TAF can emulate a completely proactive protocol by setting the radii to equal the network diameter. In a typical application, TAF would maintain proactive zones only around a few hot destinations.

The primary challenge in the design of a hybrid protocol is how to determine the optimal trade-off between the components of the hybrid. Ide-

ally, a hybrid protocol would achieve fine-grained control over this tradeoff, incur low overhead for adaptation and exploit information locality for maximum efficiency.

TAF achieves these goals by enabling each node to determine its own proactivity radius based on local information. Specifically, the proactivity radius in TAF is a function of the amount of data traffic destined to that node and the mobility rate. This function is determined locally by each node, and updates to the radius are disseminated through its proactive zone by piggy-backing them on periodic messages.

Changing the proactivity radius in TAF entails little overhead. Expanding the radius from r to s is done by broadcasting a control (CTL) packet that advertises the new radius with a time-to-live field of s . The proactive zone is maintained implicitly by piggybacking the current value of the radius onto the periodic packets exchanged by the proactive protocol. Nodes receiving this packet participate in the proactive protocol. Shrinking the radius from s to r is done by broadcasting a different CTL packet with a time-to-live field of s , and new radius field of r . In response, nodes in the proactive zone at a distance greater than r terminate their proactive activity for this destination. Note that this scheme exhibits graceful degradation without need for costly reliable multicast services or distributed consensus protocols. If the control packet is lost in the network, the nodes within r hops can maintain their participation in the proactive protocol, while nodes between r and s hops will time out and drop out of the proactive zone.

This mechanism based on proactivity radius provides a virtual 'slider' by which TAF can efficiently control the trade-off between the proactive and the reactive routing protocol at fine granularity. The choice of the precise setting for the proactivity radius depends on the goals of the system. By varying the radius selection strategy, a TAF node can try to optimize for different network metrics. In the rest of this section, we describe how we applied the general TAF framework to create three

different hybrid protocols optimized for overhead, latency and loss rate. In all three of these protocols, the radii are determined independently by destinations based on a common analytical model of the network. This model captures the inherent trade-offs between overhead, latency and loss rate and is used by TAF to determine the optimal setting for the proactivity radius.

3.1 Model

In this section, we outline the analytical model that forms the foundation of the TAF framework and enables an informed tradeoff between proactive and reactive routing protocols.

Proactive routing relies on periodic transmission of route updates. Consequently, the cost of proactive routing at each node is independent of the communication patterns in the network. Let the notation N_r^A represent the number of nodes in proactive zone of radius r around node A. If the network topology has a uniform density, there would be approximately same number of nodes in the proactive zone throughout. Let the proactive routing protocol send periodic packets with a frequency f at each node. Then the cost, in number of packets, of setting a proactive zone of radius r around A is $f \cdot N_r^A$ pkts/sec. A proactive routing protocol running for T seconds would incur an overhead of the order $T \cdot f \cdot N_r^A$ packets. This cost is independent of the number of data sources with A as the destination.

Reactive routing protocols incur an overhead at the time of route discovery. The overhead for node B to discover a route of length h can be estimated to be N_h^B packets, where N_h^B is the number of nodes at distance at most h from B. This overhead is incurred by the broadcast of route request packets. Most reactive protocols use optimizations to restrict the route discovery to a few hops beyond the actual distance. In a static network, this would be the only overhead for reactive protocol. However, mobility in the network causes routes to break, requiring extra overhead to discover alter-

native routes. Consequently, the overhead of a reactive routing protocol depends on the number of link failures in the network as well as the route lengths.

Let the parameter λ define the average lifetime of a link in the network. If the link breaks occur independently and the link lifetime follows an exponential distribution, the mean lifetime of a route of length h hops is defined by $\frac{\lambda}{h}$. In practice, link failures do not occur independently, making this quantity an approximation. In Section 5, we show that this formula approximates the observed values of average route lifetime quite closely. Thus, running the reactive routing protocol for T seconds generates approximately $T \cdot \frac{h}{\lambda}$ route-breaks for each route if h is the average number of hops of the routes found in this time.

The total overhead faced by a reactive routing protocol can be estimated to be $T \cdot \frac{h}{\lambda} \cdot N_h$, where h is the average length of the route and N_h is the average number of nodes at that distance from B. This expression gives the cost for route discovery and maintenance of a single route. If there are S sources routing packets to the same destination, then the overall cost can be expressed as $T \cdot S \cdot \frac{h}{\lambda} \cdot N_h$. This can be compared to the cost $T \cdot f \cdot N_r$ of maintaining the routes using a proactive routing protocol.

The foregoing discussion provides the intuition behind the commonly held belief that reactive routing protocols have low overhead when mobility is low and connections are sparse, while proactive routing protocols are more efficient when mobility and rate of route reuse are high. For equal values of r and h , the cost of reactive routing increases with the number of sources as well as the mobility rate. Thus, the fixed cost of proactive routing can be amortized across the multiple sources that are sending packets to the same destination, enabling it to outperform reactive routing. Similarly, when the mobility in the network increases, the average link lifetime λ decreases in proportion, forcing reactive routing protocols incur higher aggregate costs for route discovery.

Thus, the model enables TAF to quantify the tradeoff between different routing regimes in terms of overhead.

A similar tradeoff exists for reliability. Proactive protocols maintain routes constantly. Consequently, they incur low loss rates as they can quickly find alternative routes in response to link failures. In contrast, reactive protocols detect route breaks by attempting to send packets and hence suffer from packet loss whenever routes are broken. Since the frequency of route breaks is given by $\frac{h}{\lambda}$, the loss rate of a reactive routing protocol can be expressed as $\frac{h}{p \cdot \lambda}$, where p is the rate at which packets are sent by the source. Thus, when the mobility in the network is high, reactive protocols might suffer much higher packet loss compared to proactive protocols. Other factors, such as congestion, also affect the loss rates, but in a mobile environment, the impact of link breaks often surpasses other factors.

While the model presented here provides a quantifiable metric that can guide how to modify the proactive radius, it is an approximation. The values it computes may diverge from the actual behavior of the deployed routing protocols. Optimizations such as expanding-ring search, route caches, local route repair, multiple routes would impact the actual cost observed in the network. However, we show in the evaluation section that the model captures the overheads of routing protocols with sufficient accuracy and leads to the construction of adaptive hybrid protocols that outperform purely proactive and reactive routing regimes.

In the next section, we discuss the application of this framework to the construction of specialized routing protocols. Since each node makes independent decisions, further discussions in the paper only describe adaptation at a single node. However, these protocols apply equally well to multiple nodes, as the adaptation does not require any consensus or communication between participating nodes.

3.2 Minimizing Packet Overhead

Routing overhead is a critical consideration when choosing routing protocols. In mobile environments, nodes are typically limited by battery power. Routing algorithms that require excessive communication will experience greatly diminished system longevity.

We propose a protocol for minimizing the per-packet overhead of routing algorithms based on the TAF framework. Called TAF-PO, this protocol performs a dynamic adaptation between fixed, high cost proactive routing protocols versus the varying costs of reactive protocols in order to minimize routing overhead. The cost of proactive routing shows little variation with mobility and traffic patterns, and instead depends mostly on the number of nodes in the proactive zone. However, the cost of reactive routing protocol varies with the number of sources communicating with a given destination, as well as the mobility in the network. Depending on the instantaneous values of these parameters, there is an opportunity for optimization by choosing one routing regime over another.

The goal of the TAF-PO protocol is to dynamically find the values for proactive radii that optimize the total cost. Using the model introduced in Section 3, the expression $\frac{h}{\lambda} \cdot N_h$ describes the cost of reactive component for each source, where h is the number of hops along the route that uses reactive routing protocol to forward packets and λ is the mean lifetime of a link. By increasing the proactive radius, we can reduce the value of h and decrease the cost of the reactive component. The cost of the proactive component is given by the expression $f \cdot N_r$, where N_r is the number of nodes in the proactive zone. By keeping track of the values of h , λ , N_h , N_r , the destination can predict whether an increase or decrease in the proactivity radius would lead to an improvement in routing overhead.

Keeping track of the metrics required for TAF-PO is straightforward. The value of h , route length, can be obtained from the time to live (TTL)

value in the IP header of the data packet. The value of λ , mean link lifetime, is tracked at each node within the proactive zone by measuring the average lifetime of each of the links. Each node appends the measured value of λ and the number of its upstream nodes to the periodic beacon packet it uses to send updates. This information is aggregated by the proactive nodes and the cumulative results are passed on to the destination. Thus the destination can obtain the values of h , λ , and N_r for nodes and links in its proactive zone. It then approximates the value for N_h based on the value of N_r assuming that the node density is approximately the same around both regions. Estimation of N_r poses a restriction that the proactive radius under TAF must be greater than or equal to one, but the impact of this restriction on the overall cost of TAF-PO is small and conservative.

Under TAF-PO, the destination estimates the cost benefits of increasing or decreasing the current radius based on the measured parameters and the analytical model. It then decides to increment or decrement the radius if the estimated benefit is beyond a threshold. We pick a threshold of 1.2 for expanding the proactive zone and a threshold of 1.5 for shrinking it. We picked these numbers based on a set of simulations performed for different values of thresholds. A more rigorous establishment of the threshold values is being explored. A higher threshold is used for decrementing the radius because a decrease in radius could invoke link breaks and hence increase the overhead of the reactive component.

3.3 Target Loss Rate

Loss rate is a critical parameter for a network-layer routing protocol. Higher layer protocols such as TCP are quite sensitive to the loss in the underlying layers. A routing protocol that results in a high loss rate will experience greatly diminished TCP throughput [4].

We used the TAF framework to construct a protocol, named TAF-TLR, for achieving a target loss

rate. The core operation of the protocol is to adjust the proactive zone in response to perceived loss at the destination such that the protocol does not experience loss greater than the targeted rate. A secondary goal of this protocol is to achieve the targeted loss rate with the lowest possible routing cost. Clearly, in the absence of such a restriction, expanding the proactive radii to encompass the network would trivially propagate routes to all nodes. However, this approach is suboptimal due to the excessive packet overhead and consequent power consumption it would require. TAF-TLR uses the TAF framework to pick the minimal sufficient proactive radii to guarantee a targeted loss rate without incurring excessive overhead.

TAF-TLR uses the perceived loss at each node as the primary driving metric for adaptation. In essence, high perceived loss will drive the protocol to expand a proactive zone, while low loss rates will enable it to shrink the zone size. There are many direct and indirect techniques for measuring the loss rate at a node. For instance, it is often trivial to extract this information from TCP sequence numbers without any extra space or time overhead. For simplicity, and in order to support any protocol on top of IP, we follow a more straightforward and conservative approach for measuring the loss rate that requires slightly more space in each packet. TAF-TLR attaches an IP option header to each packet with the number of packets generated in the last few seconds. The destination node records the number of packets it received in an interval of the same length and uses the ratio to estimate the current loss rate for the routing protocol. While a production implementation would use implicit data collection from higher layer protocols; we note that the scheme represented here is general and biases TAF-TLR performance towards the conservative side.

Once a perceived loss rate metric is calculated, TAF-TLR manipulates the proactivity radius to achieve the target loss rate without excessive overhead. TAF-TLR operates in epochs, each of which consists of a measurement phase followed by an

adjustment to the radius of proactivity. If the exponentially decaying average of loss rates measured in the last measurement phases is higher than the target rate, TAF-TLR increments the radius of proactivity by one. If the perceived loss rate is well below the target loss rate, the radius of proactivity is lowered to reduce excessive routing overhead. TAF-TLR thus hunts for the appropriate zone radius setting in a similar manner to the TCP congestion control mechanism[5].

3.4 Latency Optimization

The third protocol we constructed based on the TAF framework is TAF-LO, a hybrid protocol for minimizing network latencies while reducing routing overhead. Reactive routing protocols may entail long perceived latencies, on the order of several seconds, since they perform costly route discovery operations on-demand. This route discovery operation is repeated from scratch whenever broken routes are detected. Thus the latency of reactive routing protocols increase with greater mobility in the network. In contrast, the latency of proactive routing protocol typically depends only on the distance between the source and the destination. Recovery from lost packets by transmitting through an alternate route might increase the overhead slightly when the network is highly mobile. These differences between proactive and reactive routing protocols make it possible to devise a TAF-based adaptation for finding a combination to minimize latency versus packet overhead.

TAF-LO manipulates the proactivity radii in order to achieve low latency with minimal routing overhead. Like TAF-TLR, TAF-LO requires a metric that captures the observed latency. Again, such measurements may be performed implicitly from information embodied in the transport layer protocols. However, we pick a simple approach and measure it directly from data embedded in packets. This implementation decision is separable from the rest of the protocol. TAF-LO attaches a packet origination time into each packet in an

IP option header. The latency is then estimated at the destination. Since we are interested only in the increase and decrease in latencies rather than the actual values, the sender and receiver need not be synchronized and their clocks may be skewed by any arbitrary amount. We do, however, assume that the clock drift between sender and receiver is negligible compared to the round trip time.

TAF-LO operates in a manner analogous to TAF-TLR, but with latency as the metric for optimization. In each epoch, the destination increments its proactivity radius by one. In the next epoch, it observes any changes in the latency and continues to increase the radius if the latency decreases beyond a threshold factor. If the latency increases beyond a certain threshold, the destination shrinks the proactive zone, but waits for two epochs before incrementing the radius again. This exponential backoff stabilizes the sizes of proactive zones and avoids frequent changes as the TAF-LO protocol searches near the optimal value. In our implementation, we use a threshold of 1.2 for increments, and 1.5 for decrements.

4 Implementation

In this subsection, we describe the details of the adaptive routing protocols we built based on the TAF framework.

TAF uses TORA and AODV as off-the-shelf components of the hybridization framework. TORA, Temporally Ordered Routing Algorithm [10], is the proactive routing component in TAF. TORA operates by maintaining a destination rooted directed acyclic graph independently for each destination node. The DAG is defined by a five-tuple height computed for each node. TORA performs routing by forwarding packets from high nodes to lower nodes that are closer to the ultimate destination. This height-based approach enables TORA to have many alternative paths and thus avoid excessive communication, as update messages need only be sent when a broken link is the

last down stream edge to a destination. For the reactive routing component in TAF, we use AODV, Ad-hoc On-demand Distance Vector routing protocol [13]. An accompanying internet draft [12] describes the detailed operation of AODV and how to set its timeouts and parameters.

TAF adapts between these two routing regimes by adjusting the proactivity radius as described in the previous section. Since we need to restrict the proactivity to a small zone around the destination, we alter TORA to bound its range of operation. Specifically, TAF adds a new component to the height tuple denoting the distance from a given destination. Only nodes with a distance value less than or equal to the proactivity radius participate in TORA. Participation requires sending periodic beacons with the height of the originating node. TAF uses these periodic packets to detect link breaks. We assume that a link is broken when two consecutive periodic beacons are missed.

Whenever the last downstream link of a node is broken, TORA sets the height of that node higher than all its neighbors using a virtual clock to identify time of occurrence of the link break. Since this operation changes the distance of that node from the destination, we make the node guess its distance to be 1 hop more than the neighbor with smallest distance. This modification restricts the proactivity to continuously remain within the zone. In order to prevent a drift over time, the destination broadcasts a control (CTL) packet within the proactive zone periodically that resets the height of all the nodes.

The CTL packets are also used to assert the current radius of proactivity at the end of each epoch. Each CTL packet carries a sequence number and the value of the new radius. Nodes receiving CTL packets for the first time join the proactive protocol based on their distance. This enables TAF to effect an increase in the zone radius with overhead proportional to the number of nodes being added to the zone. TAF uses a similarly low-overhead mechanism to shrink the proactive zone. When the radius is decreased, the nodes at the edge of

the new proactive zone send update packets to the nodes no longer in the current zone. We rely on timeouts to obviate the need for reliable broadcasts. If these update packets are dropped for any reason, the nodes no longer in the new zone would detect link breaks as other members of the zone stop beaconing to them, and naturally prune themselves out of the proactive zone. Overall, this notification mechanism enables TAF to efficiently manage the zone sizes without need for a reliable multicast protocol.

Nodes use TORA to reach the destination if they reside within its proactive zone. Otherwise, they employ AODV to discover routes. If a cached route is not available, AODV initiates a traditional route request. Whenever an intermediate node in the proactive zone for that destination receives a route request, it replies back to the source without further propagating the route requests. In case of link breaks, TORA transparently calculates alternative routes based on node heights. If no downstream node can be found, TORA drops packets. Whenever a TORA node drops packets, it sends an AODV route error back to the source.

These low-overhead mechanisms for integrating AODV and TORA provide a natural, seamless boundary between the two protocols. In the next section, we evaluate their effectiveness and efficiency.

5 Evaluation

We performed dynamic adaptation between reactive and proactive routing protocols in a simulation environment. We chose the three standard parameters to measure performance routing protocols, routing cost, loss rate and latency as the criteria for adaptation. In this section, we present the results observed while adapting based on these parameters.

5.1 Simulation Setup

We evaluate the three specialized protocols based on TAF using GloMoSim [16], a scalable packet-level simulator.

As described in the previous section, we simulated a routing protocol that uses AODV as the reactive routing protocol and TORA as the proactive routing protocol. The operation of AODV was implemented based on the internet draft [12]. Recently introduced optimizations such as gratuitous RREP and local error recovery were not included in this simulation. We implemented TORA as described in the internet draft [2] and made the changes outlined in the previous section.

There are numerous protocol settings to which the TAF framework is agnostic. We nevertheless cite them for repeatability. The TORA periodic packet interval was set to one second while the CTL packet interval was set to five seconds. The bandwidth of the physical channel was set to be 2 Mbps. The radio-layer employs a two-ray path propagation model to simulate signal propagation. The nominal transmission range of this model was 220m corresponding to the WaveLan radio hardware. We used IEEE 802.11 as the MAC protocol. Since IEEE 802.11 guarantees reliable unicast and notifies packet loss AODV neighbor discovery mechanism is not employed to detect link breaks.

The topology in our simulations consisted of 160 nodes distributed randomly using a uniform distribution in a square field of area 1700×1700 . Each simulation was run for duration of 360 simulated seconds. The mobility in the environment was simulated using a random-waypoint mobility model. According to this model, each node randomly chooses a point in the field and moves towards it at a randomly chosen velocity. The node pauses for a specified period at the destination before continuing the same pattern of motion. In our simulations, velocities range randomly between 0 m/s and 20 m/s, and wait times are 60 seconds. We change the mobility rate by varying the number of mobile nodes in the network. A mobility fraction

of 0 corresponds to all stationary nodes while a mobility fraction of 1 corresponds to all nodes in motion.

A constant bit rate (CBR) generator drives the data traffic in our simulation. In each simulation trial, 20 nodes attempt to send packets at a rate of two packets per second to a single destination. The sources and the destination were chosen randomly. Packet sizes were set to 512 bytes. The sources start transmitting from a time randomly chosen between 50 seconds and 100 seconds of the simulation, and terminate data transmission after 250 seconds, sending 500 packets. We repeated each simulation 5 times changing the value of the random seed. The results presented here are the averages of these 5 trials.

5.2 Results

In the next few sections, we examine TAF-PO, TAF-TLR, and TAF-LO, a family of TAF-based protocols for minimizing packet overhead, optimizing for a targeted loss rate, and reducing routing latency, respectively. We compare these protocols to purely reactive AODV and purely proactive TORA, changing the experimental conditions over a wide range that enables both types of routing regimes to excel. We show that the TAF-based hybrid protocols outperform fixed, that is, purely proactive or reactive, routing algorithms. That is, TAF-based hybrid algorithms perform as well as or better than the best of the proactive and reactive routing protocols. The reason for this is that TAF-based protocols adapt quickly and with low overhead to locate the sweet spot that represents the good tradeoff between the two routing regimes.

We present detailed measurements to provide the intuition behind these results and demonstrate the case for hybrid routing. We show that the optimal routing strategy often lies somewhere between purely reactive and purely proactive routing protocols. We demonstrate that TAF-based protocols can operate in this realm between the two regimes. We finally show that the model and approxima-

tions employed by the TAF framework are sufficiently accurate and effective.

5.3 Minimizing Packet Overhead

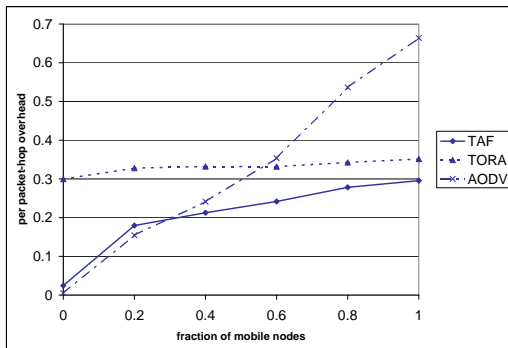


Figure 1: Cost Adaptation: Average Routing Overhead

Figure 1 shows the routing cost of AODV, TORA and TAF-PO, our protocol for minimizing packet routing overhead. The graph shows how much extra overhead the routing protocols extracted from the network on top of the data traffic by plotting the ratio of the total number of control packets to the total number of data packets. A high ratio indicates that the routing protocol extracted a large toll, wasting bandwidth and power, introducing delay, and possibly leading to congestion. As expected, the TORA overhead is independent of the mobility rate, whereas the overhead of AODV increases with increasing mobility and the concomitant reduction in link lifetimes. Our adaptive protocol achieves a lower overhead than both of the pure routing protocols. At very low mobility the hybrid approach shows a slightly higher overhead than AODV because of the restriction to maintain a proactive radius of at least one.

Figure 2 provides the intuition behind why our hybrid approach outperforms the proactive and reactive routing algorithms. It plots the per-packet overhead of different routing protocols as a func-

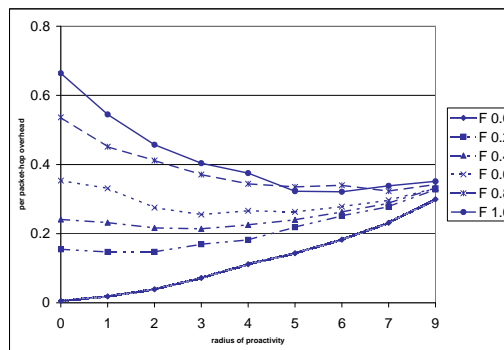


Figure 2: Static Analysis: Average Routing Overhead

tion of the size of the proactive zone. It examines six scenarios in which varying fractions of nodes are mobile, where the 1.0 line corresponds to the case where all the nodes are in motion. The right hand side of the graph corresponds to a purely proactive algorithm (TORA), and shows that its per-packet overhead is high, but also largely independent of the mobility rate. The left hand side corresponds to a purely reactive algorithm (AODV), and shows that the protocol overhead increases with the amount of mobility in the system. The intermediate nodes represent cases where the proactivity radius is statically set to the value shown on the x-axis. This graph clearly demonstrates that no single point on the graph accommodates a wide range of mobility rates. There is no silver bullet; a dynamically adaptive algorithm is necessary to find the optimal tradeoff. It is TAF's temporal adaptation mechanism that allows it to shift the protocol to the appropriate location on the x-axis and realize reductions in packet overhead.

5.4 Achieving a Target Loss Rate

We next examine TAF-TLR, our hybrid protocol for achieving a targeted loss rate with the lowest possible overhead. Figure 3 examines the loss rate characteristics of TORA, AODV and TAF-

TLR as a function of mobility. In this graph, the targeted loss rate is 5%. As expected, the loss rate of AODV increases significantly with mobility. TORA, on the other hand, can achieve a loss rate that does not vary much with the amount of mobility in the network, but this reliability comes at the expense of over-communication. TORA propagates routes constantly and throughout the entire network, and devotes extra bandwidth and power to route maintenance. TAF-TLR, on the other hand, uses just as much proactive routing as is necessary to achieve the target loss rate. This graph shows that TAF-TLR achieves the targeted loss rate. Even in the presence of very high mobility, the adaptive protocol achieves a maximum loss rate of 4.89% .

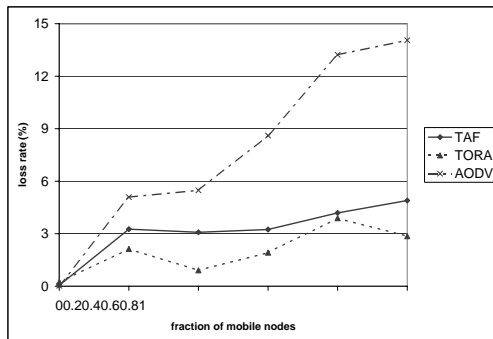


Figure 3: Loss Adaptation: Average Loss Rate

Figure 4 provides the intuition behind TAF-TLR's operation. The graph plots observed loss rate as a function of the proactivity radius and confirms our earlier observation that the loss rate decreases as the amount of proactivity in the network is increased. It also shows that the loss rate of AODV (radius 0) increases with mobility. The loss rate for TORA (radius 9) is quite low, and the curves illustrate the operation of TAF-TLR. In essence, TAF-TLR operates by sliding the hybrid protocol sufficiently right to achieve the targeted loss rate, but not too far right to avoid the excessive overhead required by TORA.

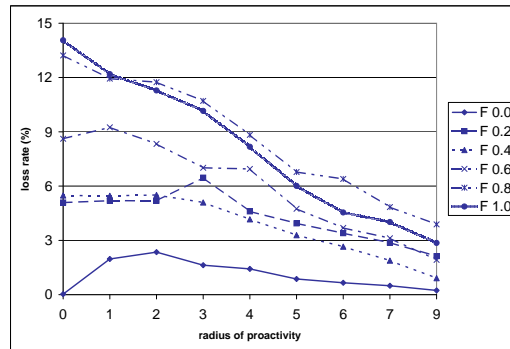


Figure 4: Static Analysis: Average Loss Rate

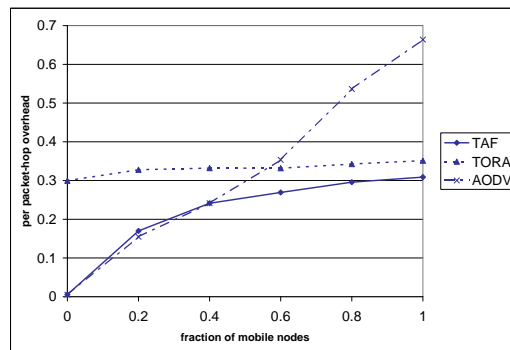


Figure 5: Loss Adaptation: Average Routing Overhead

The hidden benefits of TAF-TLR are shown in Figure 5, which plots the routing overhead as a function of mobility. While TORA achieves low loss rates, it expends excessive energy propagating unnecessary updates throughout the network when mobility rates are low. AODV entails minimal overhead in such static networks due to its on-demand operation. Again, the TAF-based protocol outperforms both AODV and TORA while achieving a given loss rate. This graph demonstrates two related facts. First, no single, static parameter setting is suitable for all scenarios. A beaconing period suited for high mobility rates extracts too much energy in static networks. Long beaconing intervals reduce overhead but increase loss

rate. Second, the TAF framework enables adaptive protocols that can achieve the target loss rate with minimal routing overhead.

5.5 Optimizing Latency

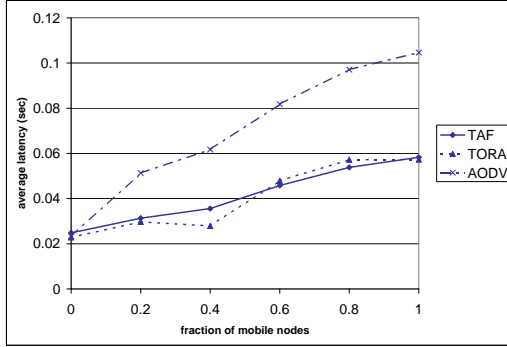


Figure 6: Latency Adaptation: Average Latency

Finally, we examine the latencies of AODV, TORA and TAF-LO. Figure 6 shows how the average latencies observed by these protocols vary with mobility. As expected, AODV latency increases significantly with mobility. The latency of TORA also shows an increase at high mobility as it tries to use alternate routes when a packet drop is reported. TAF-LO achieves performance that is comparable to a purely proactive protocol, showing significant advantages over AODV. In addition, TAF-LO requires a fraction of the overhead that TORA entails.

Figure 7 shows the variation of latency in static simulations with no adaptation. The latency can be seen to drop as the radius of proactivity is increased. At lower speeds, there is no significant change in the latency with increase in radius. At higher speeds there is a small increase in latency when the radius of proactivity is very high. We found that this increase is due to repeated attempts by TORA to find alternate routes as the MAC layer reports packet-loss events. This graph suggests that an adaptive protocol would ideally like to be

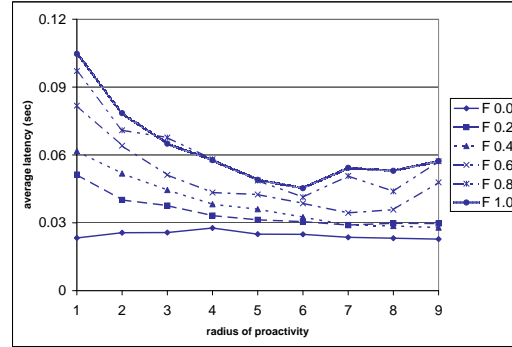


Figure 7: Static Analysis: Average Latency

located at a radius in between the minimum and maximum.

5.6 Analysis

We have shown that adaptive protocols based on the TAF framework can achieve good performance. These adaptive protocols are driven by the model presented in section 3. Hence, it is important to see how well the model is able to match the observed values in the simulations.

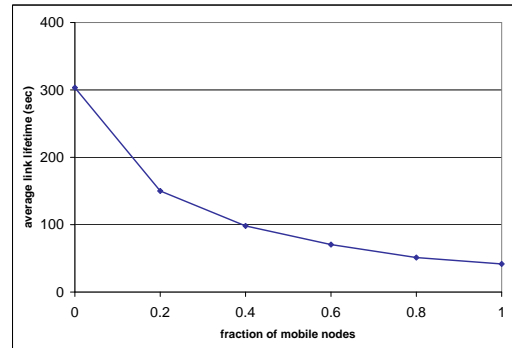


Figure 8: Static Analysis: Average Link Lifetime

Figure 8 shows the variation of average link lifetime observed in the simulations. Figure 9 shows average number of route discoveries performed by AODV. It also shows the expected number of route

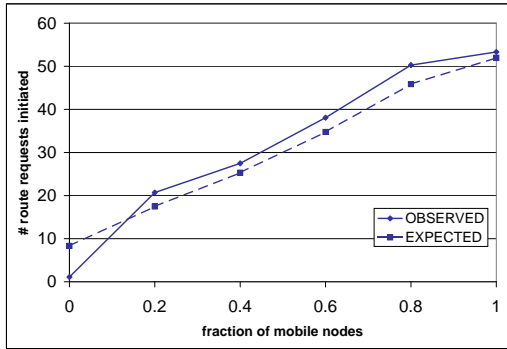


Figure 9: Static Analysis: Average Number of Route Discoveries

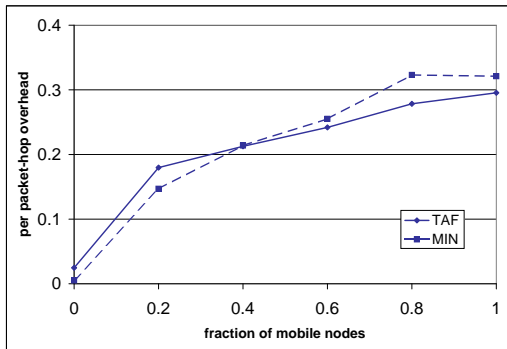


Figure 10: Cost Adaptation: Routing Overhead

discoveries computed from the observed values of λ , average link lifetime. The figure shows that the expected values matches quite closely with the actual observed values except in the case of no mobility. This is because λ values of $360seconds$ is used instead of ∞ .

Figure 10 shows the routing overhead of the cost adaptive routing protocol along with the minimum values observed from the simulations with static values of proactive radius. The adaptive protocol closely follows the pattern of the static values even performing better since it is able to dynamically adapt and hence find a lower minimum. This graph illustrates that performing cost adaptation based on the analytical model is quite effi-

cient.

6 Future Work

In this paper, we treat the adaptation for different destinations independent of each other and do not take advantage of commonality in the network. In the presence of multiple destinations, the overhead of maintaining proactive zones can be shared wherever the zones overlap. Consequently, we could achieve further minimization of the routing overhead by coalescing packets. Further, the periodicity of the proactive protocol could itself be adapted based on the mobility in the network. This could further lower the cost of proactive routing protocols and facilitate greater optimization. The TAF framework provides a foundation for studying such optimizations and adaptation strategies.

7 Conclusions

In this paper, we presented a framework for dynamic adaptation between proactive and reactive protocols. Our quantitative measurements show that there are many combinations of mobility and traffic patterns where the optimal routing strategy lies between purely proactive and purely reactive protocols. Our framework enables the construction of routing algorithms that can operate between these two extremes. Our framework is general, effective and efficient. It enables different nodes on the same network to vary the combination of proactive and reactive routing protocols according to entirely different metrics of their own choice. We outlined the design of three specialized protocols based on this framework and evaluated their performance. In all cases, adaptive protocols based on the TAF framework are as good as or better than pure routing protocols. Overall, there is a large spectrum of design points between proactive and reactive protocols. The TAF framework enables fine grain exploration of the full spectrum.

References

- [1] R. Boppana and S. Konduru. An Adaptive Distance Vector Routing Algorithm for Mobile, Ad Hoc Networks. In *Proceedings of the IEEE Infocom 2001*, March 2001.
- [2] S. Corson and V. Park. Temporally-Ordered Routing Algorithm (TORA) Version 1 Functional Specification, December 1997.
- [3] Zygmunt J. Haas and Marc R. Pearlman. The Performance of Query Control Schemes for the Zone Routing Protocol. *IEEE/ACM Transactions on Networking*, pages 427–438, August 2001.
- [4] G. D. Holland and N. H. Vaidya. Analysis of TCP Performance over Mobile Ad Hoc Networks. In *Proceedings of the Fifth Annual ACM/IEEE Conference on Mobile Computing and Networking (MOBICOM)*, August 1999.
- [5] V. Jacobson. Congestion Avoidance and Control. *Computer Communications Review*, 18(4):314–329, August 1988.
- [6] M. Joa-Ng and I-Tai Lu. Peer-to-peer Zone-based Two-level Link State Routing for Mobile Ad Hoc Networks. *IEEE Journal on Selected Areas in Communications*, 17(8):1415–1425, August 1999.
- [7] Asis Nasipuri, Ryan Burleson, Benjamin Hughes, and Johnny Roberts. Performance of a Hybrid Routing Protocol for Mobile Ad Hoc Networks. In *Proceedings of the IEEE International Conference on Computer Communication and Networks (ICCCN2001)*, Phoenix, AZ, October 2001.
- [8] Navid. Nikaein, Houda. Labiod, and Bonnet. Christian. DDR: Distributed Dynamic Routing Algorithm for Mobile Ad hoc Networks. *First Annual Workshop on Mobile Ad Hoc Networking and Computing, MOBIHOC 2000*, Boston, 2000.
- [9] Navid. Nikaein, Shiyi. Wu, and Bonnet. Christian. HARP: Hybrid Ad hoc Routing Protocol. *International Symposium on Telecommunications, IST 2001*, 2001.
- [10] Vincent D. Park and M. Scott Corson. A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks. In *Proceedings of IEEE INFOCOM*, Kobe, Japan, April 1997.
- [11] Marc R. Pearlman and Zygmunt J. Haas. Determining the Optimal Configuration of the ZRP. *IEEE JSAC*, 17(6), August 1999.
- [12] C. Perkins, E. Royer, and S. Das. Ad-Hoc On Demand Distance Vector (AODV) Routing, March 2001.
- [13] Elizabeth M. Royer and Charles E. Perkins. Ad-hoc On-Demand Distance Vector Routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, February 1999.
- [14] M. Satyanarayanan. Fundamental Challenges in Mobile Computing. In *Fifteenth ACM Symposium on Principles of Distributed Computing*, Philadelphia, PA, May 1996.
- [15] R. Sivakumar, P. Sinha, and V. Bharghavan. CEDAR: A Core-Extraction Distributed Routing Algorithm. *IEEE Journal on Selected Areas in Communications*, 17(8), August 1999.
- [16] Xiang Zeng, Rajive Bagrodia, and Mario Gerla. GloMoSim: A Library for Parallel Simulation of Large-scale Wireless Networks. In *Proceedings of the 12th workshop on Parallel and distributed simulation*, pages 154–161, Banff, Canada, May 1998.