Experimental Performance Evaluation of ATP in a Wireless Mesh Network

Xingang Zhang, Randy Buck, Daniel Zappala
Computer Science Department
Brigham Young University
Provo, UT 84602, U.S.A.
zhang.xingang@byu.edu, sutekistudent@gmail.com, zappala@cs.byu.edu

Abstract—It is well known that TCP performs poorly in wireless mesh networks. There has been intensive research in this area, but most work uses simulation as the only evaluation method; however, it is not clear whether the performance gains seen with simulation will translate into benefits on real networks. To explore this issue, we have implemented ATP, a transport protocol designed specifically for wireless ad hoc networks. We choose ATP because it uses a radically different design from TCP and because reported results claim significant improvement over TCP. We show how ATP must be modified in order to be implemented in existing open-source wireless drivers, and perform a comprehensive performance evaluation on a mesh testbed under different operating conditions. Our results show that the performance of ATP is highly sensitive to protocol parameters, especially the epoch timeout value. To improve its performance we design an adaptive version that utilizes a self-adjustable feedback mechanism instead of a fixed parameter. We illustrate its advantages by using a measurement study to compare it with the original ATP and to a standard TCP Tahoe implementation.

I. INTRODUCTION

Wireless mesh networks provide an economical, yet flexible solution to the edge Internet connectivity for both urban areas and developing countries, where the cost of laying fiber may be too expensive. However, TCP performs poorly, in terms of both throughput and fairness, in multi-hop wireless networks [1], [2]. This is primarily due to the unique characteristics of wireless networks, including spatial reuse and interference constraints, which can be further exaggerated by the IEEE 802.11 MAC layer protocol that was initially intended for single-hop wireless communication. Furthermore, packet loss caused by signal fading, route changes, or interference can be misinterpreted as a sign of congestion by TCP, and the subsequent rate reduction may lead to under-utilization of the wireless network.

Due to the importance of this problem, substantial effort has been made to improve the efficiency and fairness of the 802.11 MAC [3]–[5], provide better scheduling of flows [6], [7], to improve TCP performance [8]–[11], and to create new transport protocols [12]–[15]. Much of the earlier work in this area relies on mathematical modeling or packet-level simulation to validate the improvements made, with substantially less work validated with implementation and experiments. This can be attributed to the difficulty in modifying the network stack, which is normally built into the kernel of modern operating systems. However, despite this difficulty, experimental evaluations are indispensable since it's difficult to accurately model radio wave propagation with simulations [16], and simulation results based on simplified assumptions may differ significantly from experimental results [16], [17]. As a result, recent work has often recognized the need for implementation and experimental results [6], [7], [15], [17].

In this paper, we add to the body of experimental work on wireless transport protocols by implementing and conducting a thorough performance evaluation of ATP (Ad-hoc Transport Protocol). ATP is a clean-slate design of a transport protocol for wireless networks, using rate-based congestion control and cross-layer feedback on MAC-layer packet delays [12]. We chose this protocol because relatively little work has been done to evaluate clean-slate and cross-layer designs in an experimental setting, and the original work on ATP included only packet-level simulations. Our approach differs from one previous implementation of ATP [6] in that we have implemented the entire protocol, rather than just the congestion-control algorithm. This enables us to study additional features, such as quick-start rate probing and epoch-based feedback, that have not previously been evaluated with experiments. Our work focuses on the *details* that go into developing a new transport protocol, whereas prior experimental work has typically examined overall performance or fairness [3], [7], [13].

Our implementation of ATP consists of three parts: driver-level delay averaging, per-hop delay collection, and an end-to-end user-level transport protocol. ATP's congestion control algorithm relies on measurement of the average transmission and queueing delay experienced by packets along the path used by a connection. We have modified an open-source wireless driver to collect this data. ATP's delay measurements must then be collected and inserted into packets as they are forwarded along a path. We have implemented a daemon that intercepts packets at each hop, reads the current delay measurement from the driver, and inserts this information into an ATP header. Finally, we have implemented the transport protocol itself on top of UDP using Python.

Our experimental results are obtained by running ATP in an indoor wireless mesh testbed located at BYU and comparing its performance to a TCP Tahoe implementation, also written in Python. By examining packet-level traces of ATP, we are

able to identify several issues that cause it to perform poorly. First, the quick-start probe used to calculate an initial sending rate is extremely inaccurate, often resulting in an initial rate that is much too slow or much too fast. Then, two other default ATP parameters, the epoch timeout (1 second) and the rate increase factor (0.2), result in ATP adjusting its rate too slowly. Second, ATP is highly sensitive to the operating environment – the MAC layer transmission rate and link quality; this may explain why our results show worse performance than the original simulations. Making ATP more aggressive improves its performance somewhat, but it is difficult to find one set of parameters that works well in all conditions. We design an adaptive scheme for ATP that provides more prompt feedback as needed and also includes a better initial rate estimation. Our results show that these improvements to ATP allow it to provide better performance and fairness than TCP over paths of varying lengths.

II. ATP

ATP is a clean-slate design of a transport protocol for wireless networks that uses rate-based congestion control. The design can be divided into three separate functions, performed by intermediate nodes, the receiver, and the sender.

A. Intermediate Nodes

Intermediate nodes maintain an exponentially weighted moving average (EWMA) of the queuing delay (Q_t) and the transmission delay (T_t) for each packet:

$$Q_t = \alpha * Q_t + (1 - \alpha) * Q_{sample},$$
$$T_t = \alpha * T_t + (1 - \alpha) * T_{sample},$$

where $\alpha = 0.75$. Each data packet carries the maximum delay it has encountered so far, D_{max} in an ATP header. A node calculates $D = Q_t + T_t$ and compares it to D_{max}, replacing the packet's value if $D > D_{max}$. Whenever the node observes an idle channel, then $D = \eta * (Q_T + T_t)$, where $\eta = 3$. This multiplier may increase to 5 for a path length of 5 or more hops [12].

B. ATP Receiver

For every received packet belonging to a flow, the receiver maintains an EWMA of the received D_{max} along the path:

$$D_{avg} = \beta * D_{avg} + (1 - \beta) * D_{max}$$

where $\beta = 0.85$. Whenever an epoch timer expires ($E = 1\,s$), the receiver sends a feedback packet to the sender with this D_{avg} and up to 20 SACK blocks, which helps the sender know which packets are missing.

C. ATP Sender

ATP senders implement a rate-based congestion control algorithm that operates in three phases: increase, decrease, and maintain. If the feedback rate ($1/D_{avg}$) from the receiver is (ϕ) times greater than the current rate ($\phi = 1.1$), the sender increases its sending rate linearly by $1/5$ ($\kappa = 0.2$) times the difference between the current rate and the feedback rate

Parameter	Meaning	Default Value
E	epoch timeout	1 s
α	intermediate nodes averaging	0.75
β	receiver averaging	0.85
η	idle multiplier	3
ϕ	rate increase threshold	1.1
κ	rate increase factor	0.2

TABLE I
LIST OF IMPORTANT ATP PARAMETERS

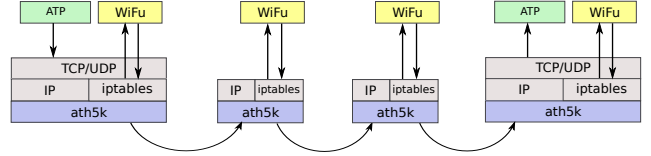Fig. 1.   ATP Implementation Architecture

[12]. If the feedback rate is smaller than the current rate, the sender immediately lowers it sending rate to the feedback rate. Otherwise, the sender simply maintains its current sending rate.

To determine the initial sending rate, the ATP sender performs *quick-start* by sending a short probe packet along its path. The receiver returns the probe immediately, without any averaging applied to the maximum measured delay along the path. The ATP sender repeats quick-start whenever it misses three consecutive feedback packets from the receiver, effectively restarting the congestion control algorithm.

For clarity, we list the relevant ATP parameters, their meaning, and default values in Table I.

The only part of ATP that we do not implement is link failure notification from intermediate nodes, similar to ELFN [8]. We omit this because we use static routing for our experiments, without testing any link failure scenarios.

III. IMPLEMENTATION

The architecture of our ATP implementation consists of three major components: 1) *delay averaging*, 2) *delay collection*, and 3) *transport protocol*. Figure 1 illustrates how these components interact. Delay averaging is done by modifying the open source `ath5k` driver [18]. Delay collection is performed by writing an application for the WiFu toolkit, which is software that has recently been developed at BYU. The transport protocol, shown as ATP in the figure, is implemented on top of UDP in Python.

A. Delay Averaging

ATP requires each intermediate node to measure the real-time queuing and transmission delay for every packet transmitted and to maintain a moving average of each of these values. However, these measurements are not readily available in the `ath5k` driver we use in our experiments, because packets are dequeued and transmitted in hardware. This means that even if we can tell exactly when a packet is put into the transmit queue, it's not possible for us to know when it's dequeued and

the actual transmission starts. Thus to approximate ATP's delay measurements, we modified the `ath5k` driver to measure the sum of $(Q_{sample} + T_{sample})$. This is the total time from when a packet enters the transmit queue until an interrupt is received, signalling the packet was successfully transmitted. The driver maintains a moving average of this sum:

$$D = Q_t + T_t = \alpha * (Q_t + T_t) + (1 - \alpha) * (Q_{sample} + T_{sample})$$

Mathematically, this is equivalent to keeping two separate averages of the delays and then summing them. The driver writes D to the */procfs* file system every time it changes.

One complication we encountered involves interrupt handling for transmitted packets in the `ath5k` driver. In the driver initialization phase, an interrupt mask is configured to determine which interrupts will be handled by the driver. Although this can be enabled to generate an interrupt for every single successful packet transmission (`TXOK`), this may cause a performance degradation due to high interrupt load. Thus, by default the driver is configured to only enable two interrupts: one for the end-of-line (`AR5K_INT_TXEOL`), which indicates this data frame is the last one in the transmit queue; and one for the transmit descriptor (`AR5k_INT_TXDESC`), which indicates that a group of frames were transmitted from the transmit queue. This default interrupt handling configuration will create an inaccuracy for our measurement of the total delay for each packet under heavy load. When the wireless network card is not busy transmitting data frame, the EOL interrupt can be used as a good approximation for transmission finish time for a packet because it may be the only packet in the transmit queue. However, if the wireless card is busy sending data frames, the descriptor interrupt only reflects the last transmitted packet's finish time, and no interrupt is sent for the other packets in the batch. Therefore, the trade-off between a more accurate measurement and better performance under load is inevitable. We choose the latter since this is the default configuration for `ath5k` driver and we want to evaluate ATP performance in a more realistic environment. We will discuss the impact of this trade-off in our experimental results.

B. Delay Collection

We implement delay collection using the WiFu toolkit, which allows user-space applications to intercept and process IP packets as they are forwarded by the kernel. With WiFu, the application specifies a set of `iptables` rules to indicate which packets to intercept, and then reads these packets using the `netfilter` interface. Thus our delay collection daemon reads packets from the `netfilter` queue, and then compares D_{max} from the packet with the current delay measurement stored in */procfs*. If the local delay is larger, then the daemon replaces D_{max} with this node's current delay measurement. All packets are then sent back down to the kernel to continue the forwarding process.

C. Transport Protocol

We implement ATP's transport protocol functionality using Python, running on top of UDP. This includes connection establishment and teardown, reliability, and congestion control.
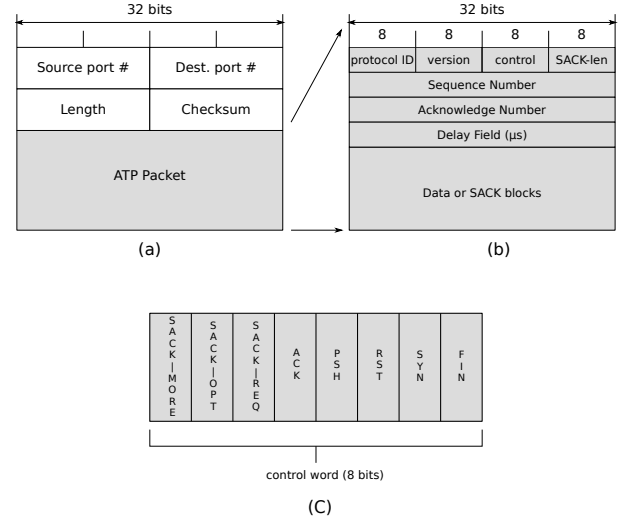

Fig. 2. ATP Header Format

During connection establishment, ATP uses quick-start to find the initial rate, and it begins transmitting packets at this speed. The TCP receiver averages delay measurements and sends a periodic ACK to the sender, once per epoch. After each received ACK, the sender will calculate a new rate and may choose to retransmit some missing packets.

Our use of Python has some drawbacks. It is well established that interpreted languages are usually slower than compiled languages. When building a transport protocol in Python, this means that fine-grained timers may be less accurate than in a C implementation, and that per-packet network I/O may be slower. In our experiments, these drawbacks become more obvious with multiple flows when multiple senders reside on the same nodes and compete for the system resources. However, for most cases our code is efficient, and we are able to make comparisons between ATP in Python and TCP in Python. We are currently rewriting our code in C to make it more efficient, so that we can compare ATP to a kernel implementation of TCP.

Because previous ATP publications do not include a header format [12], [19], we design our own ATP header, which is shown in Figure 2. Figure 2(a) shows a UDP datagram with an ATP packet as the payload. Figure 2(b) shows the ATP header and body, including:

- *protocol ID*: a number specifying ATP or TCP,
- *version*: the ATP version number,
- *control*: control flags,
- *SACK len*: total length in bytes of the SACK blocks,
- *sequence number*: sequence number in bytes, as with TCP,
- *acknowledgement number*: ACK number in bytes, as with TCP, and
- *delay*: maximum delay (D_{max}) seen by intermediate nodes along the path the packet traverses, and
- *data/SACK*: contains data for a packet going forward along a path, contains SACK blocks for an ACK.

Note that for simplicity ACK packets do not carry data.

Figure 2(c) shows the structure of control flags for ATP packet. Most of them have the same meaning as that of TCP except the first three: SACK-MORE, SACK-OPT, and SACK-REQ. Our implementation of ATP relies on SACK blocks to trigger retransmission, which means that the SACK functions essentially like a negative acknowledgement. When the sender is bursty or otherwise finishes sending a burst of data, the receiver cannot detect dropped packets from the end of the burst. Thus the receiver sets the SACK-MORE flag whenever it has not received any data for half of the epoch timeout period. When receiving an ACK with this flag set, the sender will retransmit anything that has not yet been acknowledged. The SACK-OPT flag is set by the receiver to indicate that the packet contains a SACK block rather than data. The SACK-REQ flag is set by the sender during quick-start to ask the receiver to immediately send a feedback packet, rather than waiting for an epoch to expire.

The reliability portion of ATP will resend packets that appear to be missing, based on the received SACK blocks in a feedback packet. Retransmitted packets always have priority over new packets. The sender expects a feedback packet once every $1.1 * E$ seconds; the additional 10% of waiting time is used to allow time for the feedback to travel from the receiver to the sender.

IV. EXPERIMENT DESIGN AND SETUP

Experimental evaluation has several drawbacks. First, it can be hard to generalize experiments from one testbed to another, because the operating conditions (transmit power, placement of nodes, link quality, etc.) may be different in each deployment. Second, it can be difficult to achieve repeatable results, because there are so many factors that can affect performance. We have been careful to design our experiments to avoid these pitfalls as much as possible, by varying the operating conditions in the testbed and by limiting variability from run to run as much as we can.

A. Experiment Design

We made an initial assessment of ATP performance using a set of simple experiments and found that the following variables influence performance and repeatability:

- *environmental variation*: Since our mesh testbed is located in a major campus building, there are significant differences in human activities and wireless connectivity between different periods of day. To avoid this fluctuation in interference from other wireless sources, we use IEEE 802.11a rather than 802.11b, since there is no other traffic using this frequency in our building. In addition, we alternate the test sequence between different protocols or parameters in a random order to ensure that differences seen between versions are not due to an abrupt environment change.

- *MAC rate control*: By default, the `ath5k` driver uses a rate control algorithm called `minstrel` [20] to find the best transmit rate given current conditions. This affects
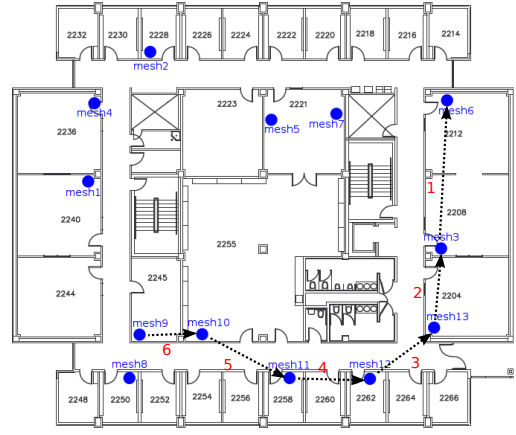

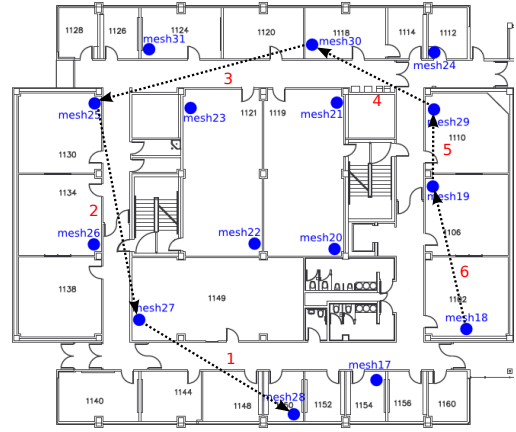Fig. 3. 2nd Floor Mesh Nodes and Path (Testbed A)

Fig. 4. 1st Floor Mesh Nodes and Path (Testbed B)

repeatability when different experiments see different MAC rates. To eliminate this possibility, we turn off minstrel and use a fixed transmit rate for our experiments.

- *routing protocol*: We initially used OLSR [21] to compute routes for our experiments, but found that it would withdraw routes under heavy congestion, causing periods of time without any route between neighboring nodes. As a result, for most experiments we use static routing to ensure the routes are consistent across experiments.

B. Experiment Setup

The primary factor affecting transport protocol performance in a single radio wireless mesh network is path length. For this reason, our experiments primarily use a single path of varying length to evaluate ATP performance. Figure 3 illustrates the portion of our mesh testbed located on the second floor of our building, with a 6-hop chain from mesh9 → mesh6. To provide greater generality, we use a second second 6-hop chain from mesh18 → mesh28 on the first floor of our building, as shown in Figure 4. For brevity, we name these configurations **Testbed A** and **Testbed B**. In our experiments, we activate only the nodes in the chain and turn the rest to a different frequency.

Each mesh node in these two configurations is a Dell desktop running Ubuntu Linux (kernel 2.6.32) with a 3Com 3CRDAG675B wireless card (Atheros AR5413 chip) that supports IEEE 802.11a/b/g. Each wireless card is loaded with our modified `ath5k` driver. In all our experiments, RTS/CTS is disabled and the MAC retransmission retry is set to *zero*. We configure each mesh node to operate on the 802.11a band to minimize uncontrolled wireless interference. In all our experiments, we transfer a 2MB text file from the sender to the receiver and record all packet events at both ends for analysis. We repeat each experiment at least 10 times to obtain a sufficient sample. Finally, we explicitly change the combination of transmit power and transmit bitrate on the wireless card to emulate different link quality and operating conditions. Due to space limitations, we only present some of our results in the following sections.

V. EXPERIMENTS RESULTS AND ANALYSIS

We compare the performance of ATP to our implementation of TCP Tahoe, which we also wrote in Python to provide a fair comparison. We first examine the performance of ATP regulating a single flow of traffic on Testbed A, so that we can examine the details of its performance. Using this experiment, we develop aggressive and adaptive versions of ATP, to overcome some of its shortcomings. We then examine the performance of these versions with multiple flows, and generalize our results to also consider Testbed B. We finish by examining the fairness of ATP in a class stack topology.

In all of our results, the error bars represent the standard deviation of the measure.

A. Single Flow

For our single flow experiments, we vary the length of the path over which ATP and TCP operate, from one to six hops, in Testbed A. We configure the nodes to use a bit rate of $24\ Mbps$ and a power of $10\ dBm$.

1) Default ATP: We begin by running ATP with its default parameters, which are listed in Table I. These are the parameters used by the simulations run by the designers of ATP [12], [19]. Figure 5 shows the average goodput versus path length. Generally, default ATP obtains better goodput than TCP, with significant gains for a single hop and for a six-hop path. ATP also demonstrates more variance in experiments than TCP Tahoe.

To investigate ATP performance in more detail, we plot the goodput and delay evolution over time from one sample 3-hop experiment in Figure 6(a). The upper half of the figure shows the instantaneous goodput of the connection over a 200 ms sliding time frame. The lower half shows a delay trace, with green points representing the D_{max} field of every incoming packet at the receiver, blue points marking the sending delay D_{send} chosen by the sender, and pink points marking the periodic feedback of D_{avg} to the sender every epoch. At the beginning of this trace, the sender transmits with a rate $(1/D_{send})$ that is too fast, which congests the network and causes D_{max} increase by more than 60 ms.
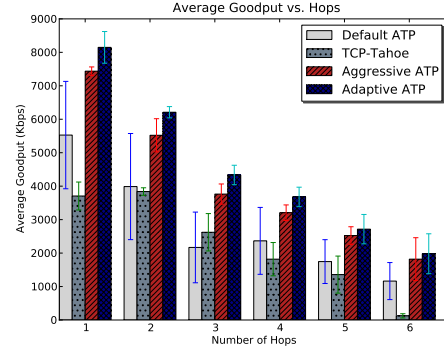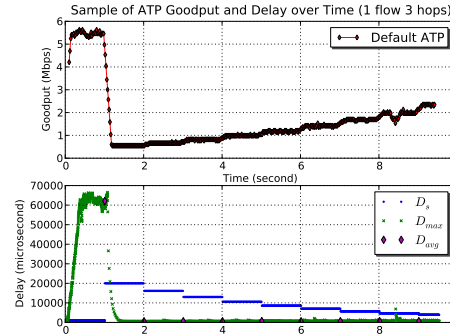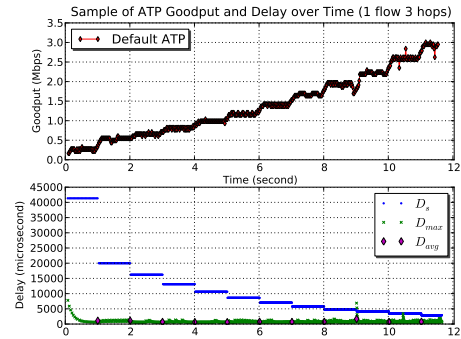

Fig. 5. Goodput, ATP and our variants, single flow, Testbed A

(a) Initial rate too fast

(b) Initial rate too slow

Fig. 6. Goodput and delay traces, default ATP, single flow, Testbed A

Because the receiver only sends feedback at the expiration of the epoch (1 second), the sender won't reduce its rate for a long period. When the sender receives the feedback from receiver, it overreacts by reducing to a sending rate of less than 1 Mbps, and then slowly increases its rate (represented by a steadily decreasing D_{send}) over multiple epochs due to a small rate increase factor (κ). Figure 6(b) shows a different trace where the initial sending rate is too slow and it takes multiple epochs to adjust the sending rate to take advantage of the link speed.

From the above analysis, we can see there are several obvious faults with the recommended ATP implementation:

(a) Aggressive ATP delay trace   (b) Aggressive ATP Zoomed delay trace   (c) Adaptive ATP delay trace
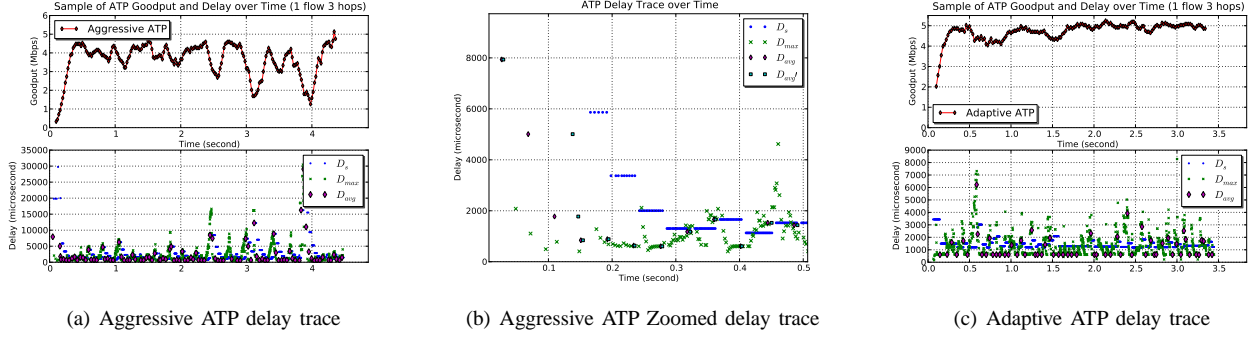
Fig. 7. Goodput and delay traces, aggressive and adaptive ATP, single flow, testbed A

(a) quick-start may not obtain an appropriate sending rate with only one probe, (b) the epoch timeout of 1 second is too large for a high bitrate network, and (c) the rate increase factor κ of 0.2 may be too slow, particularly when combined with a long epoch period.

2) Aggressive ATP: To correct these problems, we explore using a more aggressive set of parameters for ATP by reducing the epoch timeout E to 0.04 second and increasing the rate increase factor κ to 0.5. Referring again to Figure 5, we can see that with more prompt feedback, aggressive ATP performs better than default ATP by an average of 47%. The lower half of Figure 7(a) shows how a faster rate adaption helps ATP adjust to a poor choice of an initial rate. However, we also observe that the instantaneous goodput is highly variable with these new settings. Furthermore, the more aggressive feedback does not necessarily mean more *prompt* feedback, which is best illustrated in Figure 7(b). This figure zooms in on the delay trace and adds an additional set of light blue points to indicate when the feedback from the receiver is actually received by the sender, denoted $D_{avg'}$. There is a significant delay for some of these feedback packets, indicating that they must contend for the channel on the reverse path; note a cluster of feedback packets received at 0.15 seconds in the figure. This clustering occurs because we do not use a priority queue for feedback packets.

We experimented with many different settings for E and κ, and the results shown here are the best we obtained on Testbed A. It is not guaranteed that these parameters will be best in all deployments and along all paths. The designers of ATP realized the potential problems with a fixed epoch timeout of 1 second and mention that it needs to be adapted for different environments [12]. However, no solution has been provided on how to adapt the epoch to fit different wireless paths.

3) Adaptive ATP: To find a better solution for this problem, we introduce an adaptive SACK feedback mechanism that can adjust the feedback delivery time according to different operating situations. Our adaptive version of ATP uses the same settings for E and κ as the aggressive ATP, but augments the receiver to send quicker feedback. For every outgoing packet from the sender, ATP inserts the current sending delay D_s in the ATP header as depicted in Figure 2. It then uses the following rules to send a feedback packet immediately, rather

than waiting for an epoch:

- if the receiver detects two dropped packets,
- if there are 10 *consecutive* incoming packets with $D_{max} > D_s$ (this indicates the current sending rate is too fast and is causing some contention in network), and
- if there are 10 *consecutive* incoming packets with $D_{max} < (1 - \phi) * D_s$ (the sender is going too slowly).

We choose a threshold of 10 consecutive packets to avoid any temporal turbulence in the network that would induce false alarms and oscillation.

In addition, we add a number of hops field N_h to the ATP header. This field simply records the number of hops traversed from sender to receiver, similar to the IP TTL field. The sender can use this total number of hops to better adjust its idle multiplier (η) at the initiation stage of connection. We use $\eta = 3.0 + 0.5 * (h - 1)$, where h is the number of hops. The data collection portion of ATP updates this field as the packet is forwarded.

Referring again to Figure 5, we can see that in this same single flow experiment, adaptive ATP performs 64% better than default ATP. In addition, Figure 7(c) clearly shows a more smooth rate curve for adaptive ATP, which confirms our algorithm's resilience to small turbulence.

B. Multiple Flows

Fairness is another important evaluation criteria for ATP and our variants. We perform experiments with two and five simultaneous flows between the same sender and receiver pair in Testbed A. Figure 8 shows the results for two flows. Figure 8(a) shows that adaptive ATP generally achieves the best average goodput among the variants and TCP. Second, the performance advantage obtained with aggressive ATP over default ATP diminishes with increasing path length. Figure 8 (b) shows a CDF of goodput for 6 hops, verifying performance advantage of ATP over TCP Tahoe in a multihop mesh network. Finally, Figure 8(c) uses the same normalized standard deviation of goodput, which is introduced in [12] as an *unfairness index* to calibrate the fairness between the two flows. Smaller values indicate better fairness, with our modified versions of ATP performing much better than default ATP and TCP.
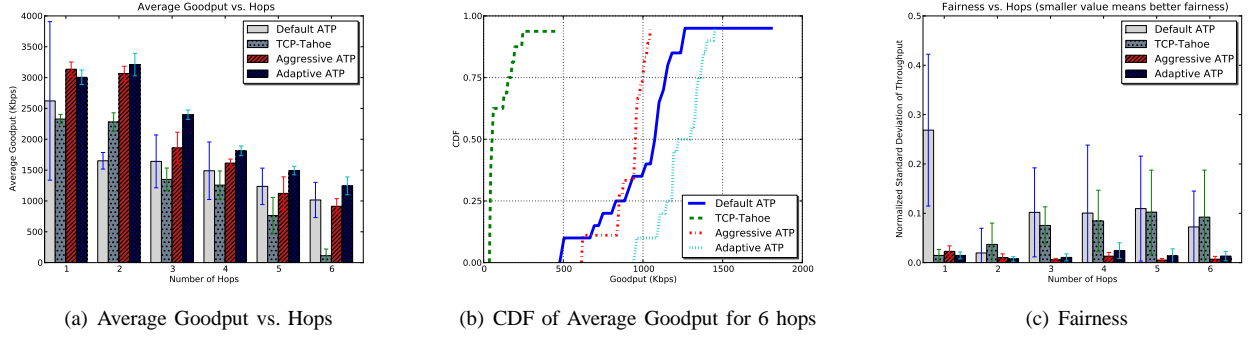
(a) Average Goodput vs. Hops

(b) CDF of Average Goodput for 6 hops

(c) Fairness

Fig. 8.   Two Flows, ATP and our variants, Testbed A

(a) Average Goodput vs. Hops

(b) CDF of Average Goodput for 6 hops

(c) Fairness
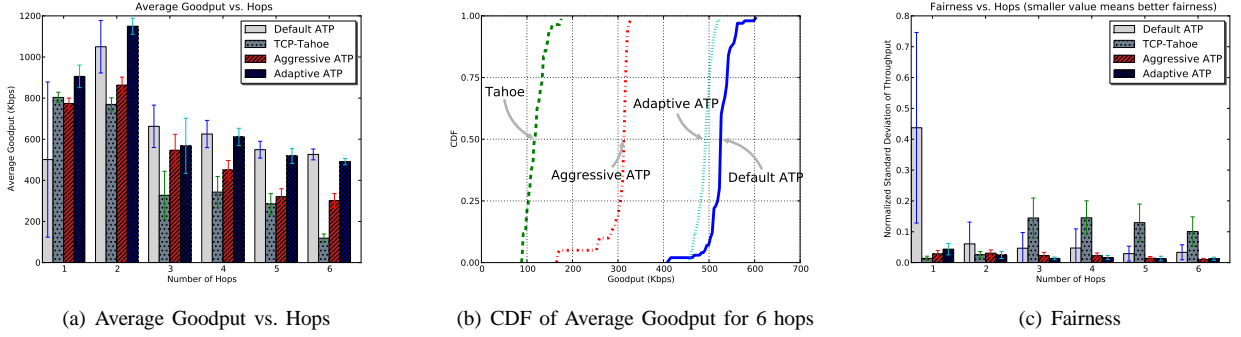
Fig. 9.   Five Flows, ATP and our variants, Testbed A

(a) Single Flow, Multihops

(b) Two Flows, Multihops

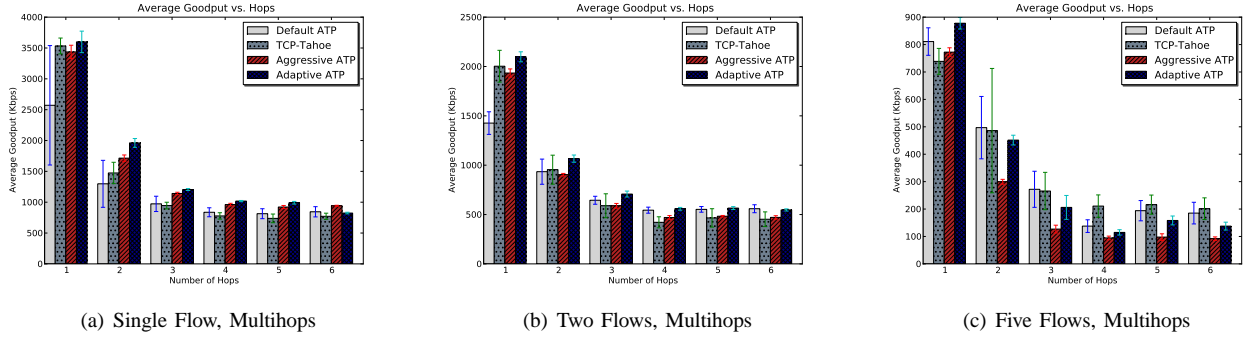(c) Five Flows, Multihops

Fig. 10.   Goodput, ATP and our variants, Testbed B

Figure 9 shows the results for five flows. Default ATP performs surprisingly well in this experiment for flows longer than one hop, and aggressive ATP fares poorly. This makes sense, because conservative rate adjustment should work better when the network is more congested. However, our adaptive version of ATP demonstrates remarkable consistency, and generally performs on a par with default ATP, with better fairness.

C. Generality

To demonstrate the generality of our results, we examine ATP's performance under a significantly different environment using Testbed B from Figure 4. We configure each mesh node to use a bit rate of 6 $Mbps$ and a power of 17 dBm. The increased transmission power and reduced bitrate, in addition
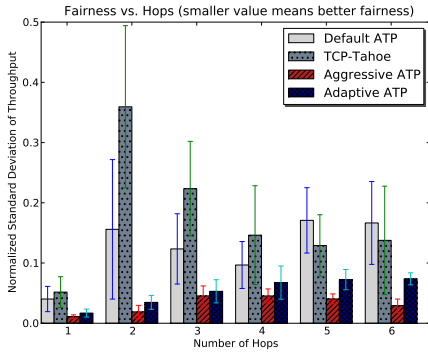
to the sparser distribution of nodes, provides a more reliable, stable mesh with lower bandwidth. For brevity, we show only the goodput and fairness results with respect to the number of flows.

Figure 10 shows the goodput for the one, two, and five flow experiments. Due to the low bandwidth and better link quality, there is not much difference exhibited between different protocols in the one and two flow scenarios, and TCP's performance is much better. For five flows, TCP outperforms ATP and its variants for paths over three hops long. Aggressive ATP generally has the worst goodput among all the tested protocols in this scenario, which further confirms that the aggressive parameters setting may not be beneficial under all circumstances.

Figure 11 shows the fairness results for the same two and
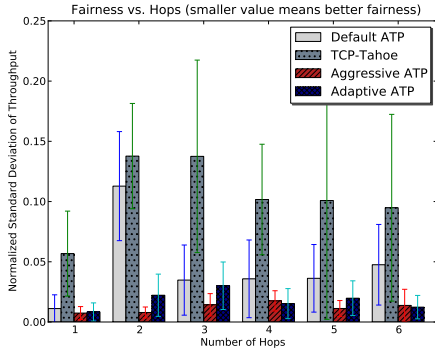
(a) Two Flows, Multihops

(b) Five Flows, Multihops

Fig. 11. Fairness, ATP and our variants, Testbed B

Fig. 12. Stack Topology Constructed on First Floor Mesh

(a) TCP Tahoe

(b) Default ATP

(c) Aggressive ATP

(d) Adaptive ATP
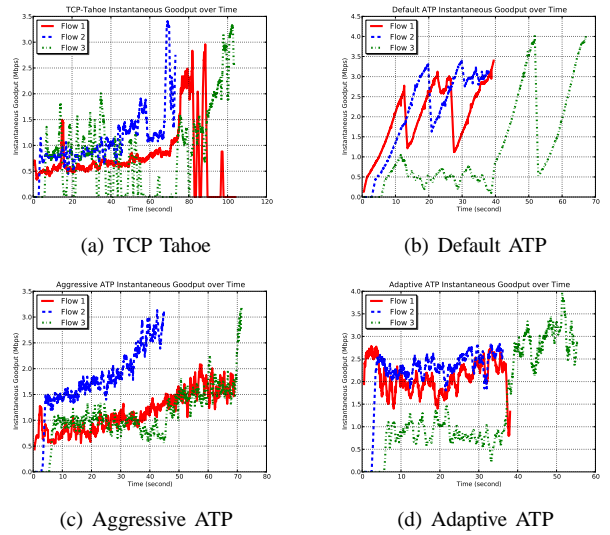
Fig. 13. Goodput trace, TCP and ATP variants, Stack topology

five flow experiments on Testbed B. It is clear that the goodput advantage of both TCP and default ATP shown in the previous figure are achieved at the sacrifice of fairness. TCP generally exhibits the worst fairness among all those tested. Meanwhile, aggressive ATP generally shows the best fairness over all, which suggests that a shorter epoch timeout contributes to better fairness. It is possible that adaptive ATP should also try to adjust the epoch time so that it is more fair in this situation.

D. Fairness with Stack Topology

Our fairness results so far are all for flows that share the same path during experiments. A common stack topology tests fairness when two flows on the edges of a network starve a flow in the middle [22]. The research community has addressed this issue in numerous works [6], [7], [13], [22] with various approaches. To explore ATP's behavior in this situation, we construct a stack topology using our first floor nodes, as shown in Figure 12. We run an experiment that starts the flows sequentially, with a 3 second delay between each flow.

Figure 13 shows the instantaneous goodput of each flow for different protocols. Since our stack topology does not contain the exact spacing that is often used in simulations and mathematical models, we can not estimate the appropriate fair share for each flow. However, the difference among the
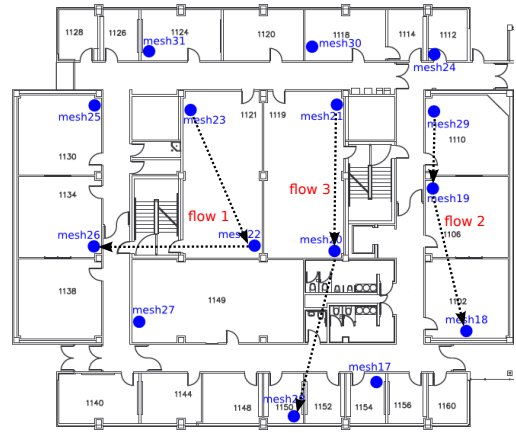
protocols in how they handle this situation are enlightening. Tahoe periodically starves the flow in the middle (flow 3), which validates our topology configuration. Default ATP does not perform as badly, but does provide the middle flow with only 1/5 the throughput of the other two flows. The sawtooth pattern in this case is due to the combination of severe penalty for a missed feedback packet and ATP's conservative rate recovery. Both aggressive and adaptive ATP provide much better goodput for the middle flow, with aggressive ATP taking more bandwidth away from flow 1. Table II lists the average goodput each flow obtains when all three flows are active with transmission. From the channel utilization perspective, both adaptive ATP and default ATP achieve better total goodput during the contention period, and adaptive ATP provides the best utility, as measured by the log of the goodput.

E. Randomized Flows

Finally, we perform a series of experiments that mimic the use of a mesh network. In these experiments, we randomly choose a set of flows in both the first and second floors of

Protocol	Flow 1 (kbps)	Flow 2 (kbps)	Flow 3 (kbps)	Aggreg. (kbps)	Utility
TCP Tahoe	703	1184	480	2368	8.60
Default ATP	2246	2483	546	5275	9.48
Aggressive ATP	966	1992	951	3909	9.26
Adaptive ATP	2085	2410	875	5369	9.64

TABLE II
GOODPUT SHARE OF EACH FLOW IN STACK EXPERIMENT

(a) TCP Tahoe   (b) Default ATP

(c) Aggressive ATP   (d) Adaptive ATP

Fig. 14.   CDF of per-flow goodput, random file transfer

(a) Average aggregate goodput   (b) Average log utility

Fig. 15.   Goodput and fairness, random file transfer

our mesh network, varying the number of simultaneous flows to impose different loads. We examine two scenarios: *file transfer* and *streaming*. For file transfer, we simultaneously initiate a 1 MB file transfer between randomly pairs of nodes and then calculate the overall goodput. For streaming, we simultaneously initiate a backlogged TCP transfer between the randomly chosen pairs and terminate the transmission after 30 seconds.

In these experiments, we use the OSLR [21] routing protocol, rather than static routing, similar to how a mesh network would be operated. To obtain longer routes, we reduce the power to $2\,dBm$ and use a bit rate of $6\,Mbps$. This provides a maximum path length of 4 hops in our mesh. We repeat each experiment 10 times with different random flows.

1) File Transfer: Figure 14 shows the CDF of per-flow measured goodput for different numbers of simultaneous flows. We consider any flow with goodput lower than $1\,Kbps$ as starved. Clearly, TCP Tahoe has a greater proportion of starved flows. With 32 flows, Tahoe starves nearly half of the total flows. All versions of ATP perform very well from the fairness perspective, with almost no starvation for up to 8 simultaneous flows. For 32 flows, ATP also has some starvation, with percentages from 20% to 30%, and adaptive ATP performs best among the variations.

We next consider the tradeoff of aggregate goodput and fairness for this experiment. For fairness, we use the log utility, which is the sum of the log of each flow's goodput, with a minor adjustment to avoid a singularity in the calculation. When taking the log of a flow's goodput, any totally starved flow yields negative infinity. Thus for starved flows, we assign their goodput to one byte.

Figure 15 shows both the aggregate goodput and log utility for the file transfer experiments. Tahoe generally obtains the best aggregate goodput for the multiple flows runs. However, this performance is achieved at the cost of overall utility, which is the lowest of all protocols tested and is zero for 32 flows due to starvation of some flows. Adaptive ATP has both the best aggregate goodput among all ATP versions and also the highest utility.

2) Streaming: The streaming experiments show drastically different goodput and utility results than the file transfer experiments. Figure 16(a) shows the aggregate goodput for the 30 second streaming experiments. Default ATP has much better performance here, and fairly good utility. Adaptive ATP performs competitively with the default settings in most cases, but has much better utility as the number of flows increases.
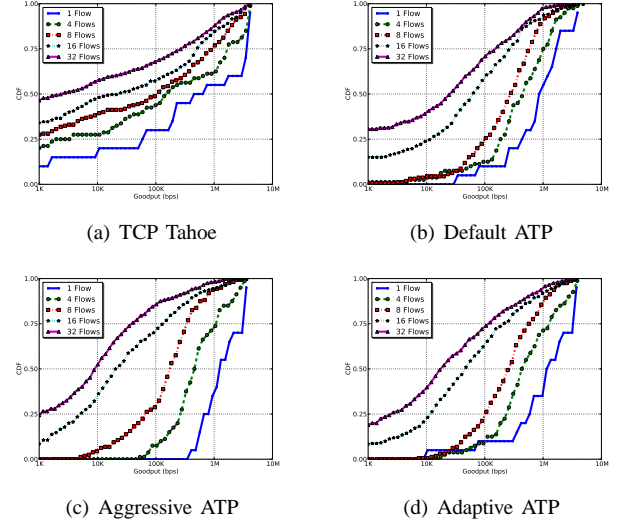
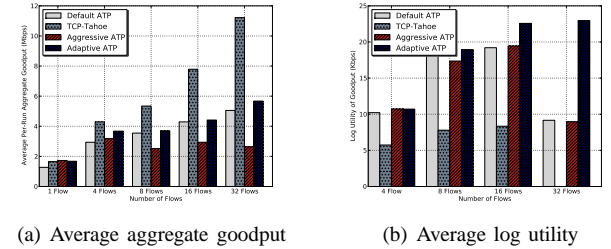The CDF of per-flow goodput is similar to the file transfer case, so we do not show them here.

VI. RELATED WORK

There has been extensive research to improve the performance of transport protocols in multihop wireless networks in the last decade. However, relatively few fully implemented transport protocols are available, and experimental results are still uncommon. Two TCP alternatives that have been implemented are DiffQ [6] and Hop [15]. DiffQ applies the theoretical work of cross-layer optimization and develops a practical backlog-based MAC scheduling algorithm with router-assisted backpressure congestion control. This work

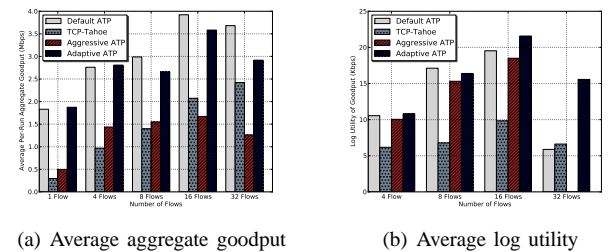(a) Average aggregate goodput   (b) Average log utility

Fig. 16.   Goodput and fairness, random streaming

includes an implementation of ATP, but only the congestion control portion of the protocol. Hop builds hop-by-hop transport protocol with in-network caching. It uses *blocks*, large segments of contiguous data, instead of packets, and uses transport-layer reliability to achieve both higher overall throughput and robustness with lossy wireless links. Their implementation is also done in user space running over UDP. Other recent experimental work focuses on the fairness of TCP in wireless networks, using protocols that allocate rates and work in conjunction with existing TCP protocols [7], [13].

Implementing network protocols at the user-level is not something new to the networking community [23] [24] [25]. The comparative advantages of user-level protocol implementation, such as the ease of prototyping and debugging are well established. Notable transport protocols include Alpine [26], which provides a framework for user-level network protocol development in FreeBSD, and Daytona [27], which implements a user-level TCP stack for Linux.

VII. CONCLUSION

In this paper, we present an extensive performance evaluation of ATP in a wireless mesh testbed. We have examined the performance of ATP as it was designed, with regard to both goodput and fairness, and found that its original design does not perform as well as shown in earlier simulation results. Although it generally outperforms TCP-Tahoe in terms of goodput, especially for the longer paths with lossy links, ATP's quick-start and fixed epoch feedback mechanism need improved designs. This is another piece of evidence that network protocols should be evaluated with both simulations and an implementation in order to properly validate their performance.

Our exploration of a revised design for ATP includes both using more aggressive parameters and creating a more adaptive feedback mechanism. A more aggressive ATP improves performance in some circumstances, but does not perform as well with multiple competing flows. Our adaptive design of ATP demonstrates a good mix of high goodput and fairness.

Due to unfairness at the MAC layer, ATP and our variants alone will not solve all the problems encountered by transport protocols in a wireless mesh network. However, we are encouraged that rate-based congestion control has been shown to be useful in a wireless setting, and that cross-layer feedback of delays encountered at the MAC layer can improve performance.

With the help of the WiFu toolkit, we plan to expand our experimental evaluations to include more transport protocols proposed by other researchers. In particular, those works that haven't been evaluated with implementations and experiments are of special interest to us.

REFERENCES

[1] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, and M. Gerla, "The impact of multihop wireless channel on TCP throughput and loss," in *IEEE INFOCOM*, vol. 3, 2003, pp. 1744–1753 vol.3.

[2] V. Gambiroza, B. Sadeghi, and E. W. Knightly, "End-to-end performance and fairness in multihop wireless backhaul networks," in *ACM MobiCom*. New York, NY, USA: ACM, 2004, pp. 287–301.

[3] A. Aziz, D. Starobinski, P. Thiran, and A. El Fawal, "EZ-Flow: removing turbulence in IEEE 802.11 wireless mesh networks without message passing," in *ACM CoNEXT*. New York, NY, USA: ACM, 2009, pp. 73–84.

[4] M. Heusse, F. Rousseau, R. Guillier, and A. Duda, "Idle sense: an optimal access method for high throughput and fairness in rate diverse wireless LANs," in *ACM SIGCOMM*. New York, NY, USA: ACM, 2005, pp. 121–132.

[5] B. Bensaou and Z. Fang, "A fair MAC protocol for IEEE 802.11-based ad hoc networks: Design and implementation," *Wireless Communications, IEEE Transactions on*, vol. 6, no. 8, pp. 2934–2941, August 2007.

[6] A. Warrier, S. Janakiraman, S. Ha, and I. Rhee, "Diffq: Practical differential backlog congestion control for wireless networks," in *INFOCOM 2009, IEEE*, April 2009, pp. 262 –270.

[7] K.-Y. Jang, K. Psounis, and R. Govindan, "Simple yet efficient, transparent airtime allocation for TCP in wireless mesh networks," in *ACM CoNEXT*. New York, NY, USA: ACM, 2010, pp. 28:1–28:12.

[8] X. Yu, "Improving TCP performance over mobile ad hoc networks by exploiting cross-layer information awareness," in *MobiCom*. New York, NY, USA: ACM Press, 2004, pp. 231–244.

[9] R. de Oliveira and T. Braun, "A dynamic adaptive acknowledgment strategy for TCP over multihop wireless networks," in *IEEE INFOCOM*, vol. 3, March 2005, pp. 1863–1874 vol. 3.

[10] S. M. Elrakabawy, A. Klemm, and C. Lindemann, "TCP with adaptive pacing for multihop wireless networks," in *MobiHoc*. New York, NY, USA: ACM Press, 2005, pp. 288–299.

[11] K. Nahm, A. Helmy, and J. C. C. Kuo, "TCP over multihop 802.11 networks: issues and performance enhancement," in *MobiHoc*. New York, NY, USA: ACM Press, 2005, pp. 277–287.

[12] K. Sundaresan, V. Anantharaman, H.-Y. Hsieh, and A. R. Sivakumar, "ATP: a reliable transport protocol for ad hoc networks," *Mobile Computing, IEEE Transactions on*, vol. 4, no. 6, pp. 588–603, 2005.

[13] S. Rangwala, A. Jindal, K.-Y. Jang, K. Psounis, and R. Govindan, "Understanding congestion control in multi-hop wireless mesh networks," in *ACM MobiCom*. New York, NY, USA: ACM, 2008, pp. 291–302.

[14] D. Scofield, L. Wang, and D. Zappala, "HxH: A hop-by-hop transport protocol for multi-hop wireless networks," in *The Fourth International Wireless Internet Conference*, Nov. 2008.

[15] M. Li, D. Agrawal, D. Ganesan, and A. Venkataramani, "Block-switched networks: a new paradigm for wireless transport," in *USENIX NSDI*. Berkeley, CA, USA: USENIX Association, 2009, pp. 423–436.

[16] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott, "Experimental evaluation of wireless simulation assumptions," in *MSWiM*. New York, NY, USA: ACM Press, 2004, pp. 78–82.

[17] W. Kiess and M. Mauve, "A survey on real-world implementations of mobile ad-hoc networks," *Ad Hoc Networks*, vol. 5, no. 3, pp. 324–339, April 2007.

[18] "Atheros linux wireless drivers." [Online]. Available: http://linuxwireless.org/en/users/Drivers/ath5k

[19] K. Sundaresan, V. Anantharaman, H. Hsieh, and R. Sivakumar, "ATP: A reliable transport protocol for ad-hoc networks," in *ACM MobiHoc*, 2003.

[20] "Minstrel rate control algorithm." [Online]. Available: http://wireless.kernel.org/en/developers/Documentation/mac80211/RateControl/minstrel

[21] A. Tønnesen, T. Lopatic, H. Gredler, B. Petrovitsch, A. Kaplan, S.-O. Tücke *et al.*, "olsrd - an adhoc wireless mesh routing daemon,." [Online]. Available: http://www.olsr.org

[22] K. Xu, M. Gerla, L. Qi, and Y. Shu, "Enhancing TCP fairness in ad hoc wireless networks using neighborhood red," in *MobiCom*. New York, NY, USA: ACM, 2003, pp. 16–28.

[23] R. A. Thekkath, T. D. Nguyen, E. Moy, E. D. Lazowska, and S. Member, "Implementing network protocols at user level," in *IEEE/ACM Transactions on Networking*, 1993, pp. 64–73.

[24] A. Edwards and S. Muir, "Experiences implementing a high performance TCP in user-space," *SIGCOMM*, vol. 25, no. 4, pp. 196–205, 1995.

[25] E. Kohler and E. Kohler, "The click modular router," *ACM Transactions on Computer Systems*, vol. 18, pp. 263–297, 2000.

[26] D. Ely, S. Savage, and D. Wetherall, "Alpine: a user-level infrastructure for network protocol development," in *USITS*. Berkeley, CA, USA: USENIX Association, 2001, p. 15.

[27] P. Pradhan, S. Kandula, W. Xu, A. Shaikh, and E. Nahum, "Daytona: A user-level TCP stack." [Online]. Available: http://nms.lcs.mit.edu/~kandula/data/daytona.pdf