Investment Function: Enhanced Fairness and Performance in Multi-hop Wireless Networks

Pradeepkumar Mani, David W Petr [1]

Information & Telecommunication Technology Center, University of Kansas

mpradeep@ittc.ku.edu dwp@ittc.ku.edu

Abstract

In multi-hop wireless networks, longer-hop UDP and especially TCP flows receive a smaller fraction of the available bandwidth than single-hop flows in the absence of specific resource allocation techniques. Also, inter-node interference effects in multi-hop networks lead to incredibly poor bandwidth utilization. This paper begins to address these and other related issues through the novel concept of an investment function which captures the investments (network resources consumed, price paid for service quality, etc.) that are made in packet flows. The investment value of packets or packet flows can then be used to control their treatment in the network, with the goal of providing fair and efficient utilization of network resources through minimizing wasted investment and/or maximizing investment throughput. We show that a simple investment function can significantly improve flow fairness by smoothing out the disparities in distribution of bandwidth among flows with different hop length while simultaneously improving bandwidth utilization efficiency.

1. Introduction

Multi-hop wireless networks find application in a number of environments; they can provide an alternative to a last-mile wireline infrastructure in geographically infeasible areas, can serve as limited-lifetime networks that provide communications infrastructure for disaster response (e.g., hurricane Katrina) or for security for infrequent, large-audience events (e.g., NASCAR races), and can provide communications for networks of sensors. Besides the relatively well-understood "wireless link" problems of limited bandwidth, high error rates, and time-varying effects such as fading, there are networking differences that arise from mobility, limited node processing power, and lack of centralized control. The focus of our research is another fundamental problem in multi-hop wireless networks: extremely biased and inefficient bandwidth utilization due to omni-directional broadcast transmission over a shared medium. This interference problem is exacerbated in the multi-hop case due to contention from packets belonging to the same flow but at different hops. For the particular case of 802.11 DCF, **[1]** showed that the throughput of a single flow traversing a chain of four or more wireless hops is upper-bounded by 0.25 of the throughput attainable if the flow traverses only one hop, with actual reduction factors closer to 0.14. When multiple flows are present, the situation deteriorates even further, with strong dependencies on traffic patterns.

The goal of our research is to explore various means to enhance multi-hop fairness and efficient utilization of the scarce bandwidth in multi-hop wireless networks. We introduce a new concept in networking, the ***investment function***, to achieve a two-pronged objective: significant increases in network bandwidth utilization, while allocating and distributing the bandwidth among flows to promote service quality and ensure fairness among flows. Although we believe the investment function potentially has very broad networking applications, we will concentrate on its applicability in a multi-hop wireless network context, where we believe its contributions can be substantial.

2. The Investment Function

In this section we introduce the concept of investment function and illustrate its flexibility. We begin by identifying four dimensions of investment that can be made in network traffic, a list that illustrates the flexibility of the investment function concept, but is by no means exhaustive. The first is that larger packets require larger investment (in both bandwidth and buffers) than smaller packets. The second is the increasing investment of network resources with each successful packet transmission (hop). The third dimension is that of congestion: more has been invested in a packet that has been transmitted over a congested link or by a congested node than in one transmitted in a relatively congestion-free environment. Finally, the customer or user will have invested monetarily in the traffic, with greater relative investment tied to greater service quality expectations.

[1] ***Corresponding Author*** *Ph: (785) 864-7762*

.

The following example scenarios illustrate the need for an investment function that combines these forms of investment. Assume three descending levels of service quality expectation (and hence relative user investment) for flows: Blue, Red and Yellow. Suppose at a given node we need to drop a packet from a given set of packets because the node buffers are full. Which of the following packets would be the right ones to drop?

1. A Blue packet that has traversed 1 hop or a yellow one that has traversed 5 hops?
2. A Blue packet that has traversed 2 hops under relatively congestion-free conditions or a Red packet that has traversed 2 hops under severely congested conditions, with numerous re-transmissions?
3. A 128-byte Blue packet that has traversed 4 hops or a 1024-byte Red packet that has traversed 6 hops?

From an efficient network utilization perspective, it seems like the Blue packet should be dropped in case 1, but the difference in user investment between Blue service and Yellow service might override that conclusion. Similarly, in the other cases the answer depends on the relative values placed on the different investment factors.

This diversity of investments in network traffic can be unified by means of an *investment function*, which can be considered to represent the global value of the packet or packet flow. Here we introduce one possible investment function, with other examples introduced later in this paper. Each packet carries in its header a *Beginning Investment (I_B)* value that is based on the packet size and the relative user investment, and a *Network Investment* (i_N) factor that reflects number of hops already traversed and network conditions at the upstream nodes. The Current Investment *(I_C)* value of a packet arriving at a node is computed as the product

$$I_C = I_B \cdot i_N$$

and the Current Investment *(I_C)* is used to make decisions about packet handling, for example, which packets are to be discarded (if necessary) at this node. Details of the investment function computations follow.

The network provider assigns a *User Investment* (i_U) factor for each service quality level, in such a manner that the separation between the i_U values for each service quality level reflects the extent of service differentiation desired. The choice of i_U values is very flexible; each network provider is free to implement a structure of their own. At the source node, the *Beginning Investment (I_B)* value is computed based on packet size S:

$$I_B = S \cdot i_U$$

Also at the source, the *Network Investment* (i_N) factor is set to some initial value γ where $0<\gamma<1$ and the *Current Investment (I_C)* at the source node is computed as the product $I_C = I_B \cdot \gamma$.

At each node (including the source node), after computing and storing the Current Investment *(I_C)* value of the packet, the Network Investment (i_N) factor is updated:

$$i_N = i_N + (1- B_A)$$

where B_A is the normalized available bandwidth as seen by this node. The calculation of available bandwidth is a difficult task which will be investigated as part of future research. Finally, embed the Beginning Investment *(I_B)* value and Network Investment (i_N) factor into the data packet before transmission to the next node. Note that the investment function is generally a time-varying quantity, and its computation constitutes a cross-layer exercise.

Along with the investment function, we introduce a new, related network performance measure: *investment throughput*, defined as investment units delivered to all destinations per unit time. For example, if we use a simpler investment function that is simply hop count times packet size, the investment throughput reduces to "network throughput" expressed in units of hops-bits per second. Specifically, the delivery rate of each flow (in bits per second) is multiplied by the number of hops traversed by that flow, and the resulting values are summed over all flows. This is identical to the "one-hop throughput" in **[1]** and similar to the bits-meter per second unit proposed in **[2]**. We also introduce an auxiliary metric: *wasted investment rate*, the investment rate of packets that are dropped before reaching their destination. This is calculated in the same manner as investment throughput, but for packets that are dropped.

The flexibility offered by the investment function can be exploited in various ways. For example, the investment function can be used by packet handling applications to control packet access to node buffers (packet dropping) during congestion, with the goal of minimizing wasted network investment and improving bandwidth parity among flows with different hop counts. Or, the investment function could be used to control access to bandwidth based on investment values. This could be implemented through priority service, control of backoff parameters in wireless network protocols such as 802.11 DCF, etc. Similarly, average flow investment could be used for flow-level control decisions.

3. Results from Example Scenarios

We use simulation to demonstrate the potential utility of the investment function concept and illustrate its characteristics introduced in section 2 by applying it to some example scenarios. Based on the general form described, one can totally customize the investment function to suit one's needs. For our simulation studies, we simplified the investment function as follows: i_N here is simply the number of hops traversed (regardless of congestion), packet size S is the same for all packets in a given simulation, and the user investment factor (i_U) equals 1 for all flows, unless otherwise specified.

We would like to demonstrate the effectiveness of the investment function as a tool to achieve the following sample objectives: (a) improve flow fairness across multiple hops, (b) increase network utilization efficiency by reducing wasted investment, and (c) enforce *controllable* flow priority using user-defined priority. In scenarios where flows of variable hop counts and variable user-priorities compete for resources, such as in a multi-hop wireless network, it seems reasonable that at any given node, service be accorded to a packet based on accumulated investment of a packet *relative* to other packets currently in the node, rather than strictly user-assigned priority to the packet. By considering accumulated investment, not only is *multi-hop* flow fairness maintained but also scarce wireless network resources are efficiently utilized. However, a relatively larger separation between user-priorities can be used to enforce absolute priority, and thus the term *controllable* priority. If user-priorities are specified, then one would normally perform *static priority mapping (SPM)* to map packets to a priority service level. Under *static mapping*, the mapping of the packet is solely based on the user-defined priority. The priority of a packet does not change as it progresses through the network. In contrast, the investment function allows us to perform *dynamic priority mapping (DPM)* at each hop, so that treatment given to the current packet is relative to investment carried by packets currently in the node. One may design a suitable mapping procedure based on one's need.

It has to be noted that these objectives are simply sample objectives. Keeping in mind the above objectives, we conducted each simulation by choosing one or more options from the following: (1) *Queuing disciplines*: simple FIFO or 3-level non-preemptive priority (3LPQ), (2) *Packet dropping*: Tail-Drop packet dropping (TDD) scheme (drop arriving packets for which there is no buffer space) or Investment-Based Dropping (IBD) scheme, in which the packet to be dropped when a queue would overflow is the packet with the smallest investment value, and (3) *User Priority (UP)*: all equal or specified.

The *DPM* scheme used in our study is described as follows: A running mean of the I_C values seen so far (μ_n) is maintained at each node (suffix n denotes n^{th} packet arrival), along with the standard deviation (σ_n). Both are maintained as running variables, updated with the arrival of the I_C value of the n^{th} packet (I_{C-n}).

$$\mu_n = w.\mu_{n-1} + (1-w).I_{C-n}, 0 \le w \le 1$$

$$\sigma_n = \sqrt{w.\sigma_{n-1}^2 + (1-w).(I_{C-n} - \mu_n)^2}$$

The value for the weight w is chosen as 0.99. Dynamic mapping of an incoming packet (n^{th} packet) with current investment I_{C-n} is performed as follows:

Lowest Priority: $I_{c-n} < (\mu_n - \sigma_n)$

Middle Priority: $(\mu_n - \sigma_n) \le I_{c-n} \le (\mu_n + \sigma_n)$

Highest Priority: $I_{c-n} > (\mu_n + \sigma_n)$

If DPM is desired when user-priorities are specified, the user-defined priorities can be translated into i_U values for purposes of network investment computation.

Based on the above options, we broadly divide the simulations into two categories: (1) *Constant User-investment* simulations *(CU),* and (2) *Variable User investment* or user-defined priority simulations *(VU)*. Under *CU*, I_C is directly proportional to hop count, due to equal packet sizes and equal i_U values across all flows. The baseline case for *CU* experiment was FIFO + TDD (or simply TDD), and the results were compared against FIFO+IBD (or simply, IBD) and 3LPQ+DPM+IBD. The results were quite identical between 3LPQ+DPM+IBD and 3LPQ+DPM+TDD. Under *VU,* the I_C value is directly proportional to the product of hop count and i_U (equal packet sizes). The baseline case in VU is 3LPQ+SPM+TDD, and the results were compared to 3LPQ+DPM+IBD. The idea behind *VU* is to illustrate that the investment function can be used to balance the objective of providing different service types to the users while utilizing resources efficiently and fairly. *Separation between i_U values will determine the extent of service differentiation*. This flexibility, we believe, is one of the more attractive features of the investment function. Although it may appear that the investment function is very similar to the classic priority concept, there are fundamental differences. With investment function, priority is assigned dynamically on a hop-by-hop basis, and the assignment of priority at any given hop depends on the relative priority of the current packet with respect to the packets already present in the node. This is a major distinguishing feature of the investment function, with DPM, when compared to the conventional form of priority service.

All simulations were conducted "grid" topology using the ns-2 simulator, and both TCP and UDP traffic patterns were used. The transmission range of each antenna was approximately 250 meters, while the Carrier Sense (interference) range was approximately 550 meters. We had IEEE 802.11 DCF MAC running on these nodes, with a maximum data rate of 1 Mb/s. AODV was the routing protocol used, while TCP-Tahoe was the flavor of TCP used. To minimize routing overhead, mobility in the nodes was disabled. The total queue size (across all priorities) was fixed to 30 packets in all cases (including FIFO), and hence under conditions of buffer overflow, lower priority packets are dropped to accommodate incoming higher priority packets. The 802.11 RTS threshold is set to 400 bytes. In all of our VU simulations, the i_U values for lower, middle and highest priorities are set as 1, 2 and 3 respectively. The flows are randomly assigned one of the three i_U values.

3.1 Grid Topology

The grid topology we used (Figure 1) consists of 16 nodes arranged in a 4x4 grid. Adjacent nodes are separated by 185m (within transmission range), while diagonally opposite nodes are separated by 265m (not within transmission range). The maximum possible hop count in this topology is 6. For each simulation, we chose source-destination (SD) pairs randomly, while enforcing the requirement that there be exactly 12 1-hop flows, 6 2-hop flows , 4 3-hop flows, 3 4-hop flows, 3 5-hop flows, and 2 6-hop flows (total of 30 flows). This was done to ensure that the total network offered load by flows belonging to various hop counts was equal (approximately so for 5-hop flows). To improve the accuracy of our results, we conducted 40 simulation runs for each experiment, with (different) random SD pairs for each run. The duration of each simulation run was 400 seconds. The performance metrics are *Investment Throughput*, *Wasted Throughput*, *Mean end-end flow delay*, *Mean flow throughput (TCP only)*, and *Mean flow packet delivery ratio (UDP only)*. To assess fairness across multiple hops for the throughput and investment throughput metrics, we make use of *Jain's fairness index (JFI)* **[3]**. The JFI for a metric X with values $\{x_1, x_2, \ldots, x_n\}$ is computed as:

$$JFI\ (x_1, x_2, \ldots, x_n) = \frac{(\sum_{i=1}^{n} x_i)^2}{n \sum x_i^2}$$

. The JFI is bounded in the range [0, 1], with higher values indicating higher degree of fairness. Consider the case where the metric *X* indicates allotted flow bandwidth: if all flows have equal bandwidths allotted (all x_i's equal), then the JFI will have a value of 1. If only k of the n flows receive equal share (while others get none), then the index is k/n. For any other of our other performance metrics X, the variance of X, in conjunction with an absolute performance measure (the mean of X) is taken as the fairness measure. The mean value is required for correct interpretation of the variance. For example, for identical absolute measures, a reduction in variance implies increased fairness.

3.1.1 UDP Simulations

The packet sizes are fixed at 128 bytes (hence the 802.11 RTS/CTS is disabled), and the mean packet inter-arrival time is set as 0.15 seconds to generate a total network load of 204.8 kb/s. Table 1 shows the results for packet delivery ratio (*pdr*) and mean flow delay, respectively, under the various schemes in CU (no user-defined priorities). Please note that there is no SPM under 3LPQ in CU; only DPM is performed, which is totally transparent to the user. The 3PLQ+DPM scheme was included to illustrate the flexibility provided by the investment function to network providers for packet handling applications, in a manner that is totally transparent to the user. When compared to the TDD scheme, the IBD and 3LPQ schemes significantly improve *pdr* and delay fairness performance across various flow hop counts. Flows with higher hop counts are the biggest beneficiaries in both IBD and 3LPQ schemes, while flows with lower hop counts suffer a mild penalty when compared to TDD scheme. As can be seen from Table 1, the variance of the *pdr* for schemes using investment function (IDB, 3LPQ) is much smaller than for the conventional TDD scheme. Moreover, one can also see a slight improvement in the mean *pdr* values as well, when using the investment function. Likewise, for almost identical mean delay values, the variance of delay is reduced by more almost one order of magnitude when using the investment function. Hence, one can safely conclude that the investment function indeed improves multi-hop fairness for the UDP-CU scenario.

Table 2 shows the investment throughput and wasted investment across the three schemes. It can be seen that both IBD and 3LPQ marginally improve investment throughput, while simultaneously decreasing wasted investment. From the high pdr values, one can conclude that the network has been loaded just beyond saturation. We achieved higher improvements in network investment throughput at even higher loads, but do not present results here due to space constraints.

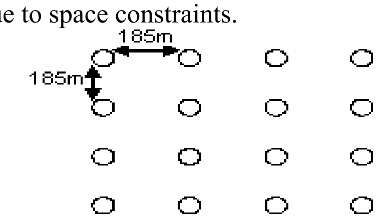

Figure 1. Grid Topology

For the VU case (user-defined priorities), we compare SPM+TDD (no investment function) with DPM+IBD (using the investment function). Figures 2 and 3 compare *pdr* and mean delay performance across flows with different hop counts. Table 3 shows the investment throughput and wasted investment across the two schemes for various priorities. The increase in investment throughput is around 5%, while the decrease in wastage is around 50%. Figure 4 compares mean delay performance of the two schemes under VU across priorities. From Table 3A, it can be seen that using the investment function significantly improves fairness across hop counts, even with user-defined priorities. As with UDP-CU, one can see from Table 3A that the investment function (DPM) is very effective in providing significant improvement in delay and *pdr* **multi-hop fairness**, when compared to SPM. The variances of *pdr* and delay values when using DPM are much lower when compared to the variances when using SPM, while offering better mean delay and pdr performance across various hop counts. This implies that the DPM offers better overall performance, while at the same time offering significant improvements in multi-hop flow fairness.

From Table 3A, it can also be seen that for comparable mean values, the variances with DPM are much lesser than the variances with SPM, when compared across the three priority levels. From figure 4 and Table 3, it is clear that the SPM+TDD scheme clearly favors the higher priority flows at the expense of the lower priority ones. However, the distinction between various priorities is blurred when using the investment function. Lower priority flows, especially ones that traverse a larger number of hops, are not starved for bandwidth. It can also be seen that though the investment function (DPM+IBD) "softens" the distinctions between priorities when compared to SPM, it does not eliminate them, as required by the user. If larger separation is used for user-defined priorities, we expect to see higher variance values for comparable mean values (lower fairness). It has to be noted that increasing the separation between user-investment values (user-priorities) does not impact the SPM scheme at all, because priority service will be provided in descending order of user priority values in SPM, irrespective of the actual user-priority values.

Table 1. UDP CU: Pdr and Delay Results

Flow Hop Count	Flow Packet Delivery Ratio			Mean End-End Delay (sec)		
	TDD	IDB	3LPQ	TDD	IDB	3LPQ
1	0.946	0.925	0.902	0.1440	0.2840	0.2940
2	0.901	0.893	0.868	0.2750	0.4460	0.4530
3	0.864	0.870	0.857	0.4120	0.4900	0.4980
4	0.823	0.856	0.848	0.5390	0.5380	0.5260
5	0.784	0.811	0.821	0.8090	0.6390	0.6410
6	0.729	0.772	0.810	1.0800	0.8030	0.7230
σ^2	0.0063	0.0031	0.0011	0.1215	0.0311	0.0224
mean	0.8412	0.8545	0.8510	0.5432	0.5333	0.5225

Table 2. UDP CU: Network Investment

Investment (Hops-kb/s)	FIFO + TDD	FIFO + IBD	3LPQ+ DPM
Throughput	380.9	387.1	390.7
Wasted	26.6	15.4	12.7

Table 4 shows the mean flow throughput and mean flow delay across various flow hop counts under CU. While the TDD and IBD schemes perform quite similarly, the 3LPQ scheme performs exceedingly well by reducing the mean delay of the higher hop counts. However, delay values beyond 3 hops are very high for all schemes. This

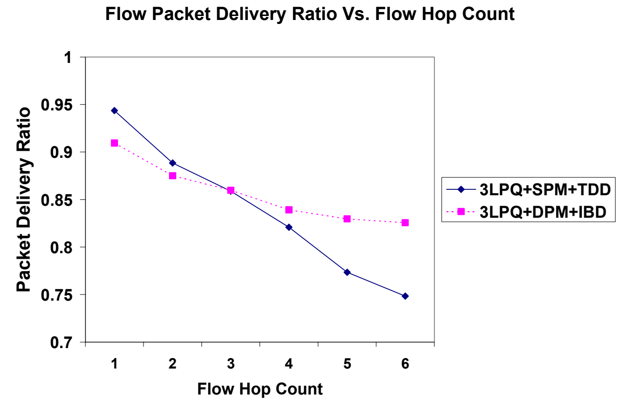

Figure 2. UDP VU: PDR vs. Hop Count

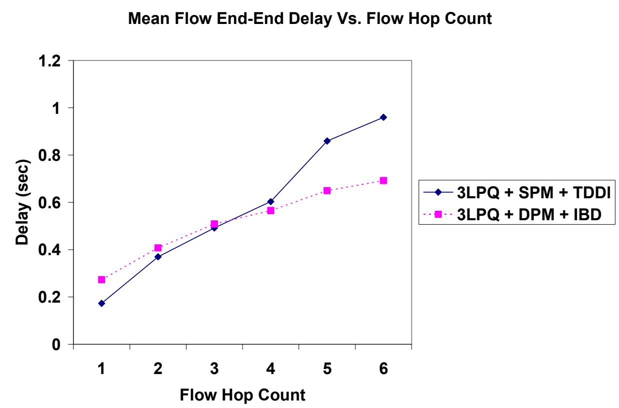

Figure 3. UDP VU: Mean Delay vs. Hop Count

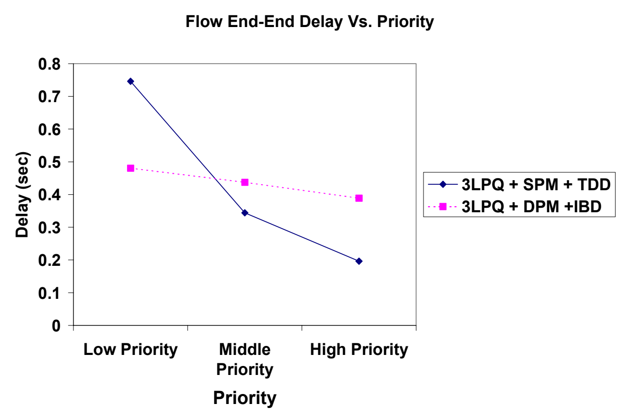

Figure 4. UDP VU: Mean Delay vs. Flow Priority

is because the flows that traverse more than 3 hops are starved for bandwidth. This asymmetric nature of wireless links and its undesirable interaction with TCP has been studied by A.Rao **[4]**. Our investigation of this problem revealed that lack of transmission opportunity for the Ack packets (reverse direction) at the MAC layer is the primary reason for poor throughput performance for higher hop count TCP flows. In other words, if the node sending the Ack is just outside the *Carrier Sense range* of any node transmitting multiple TCP flows, then the transmission opportunities to the node sending the Ack will be very limited, because the node sending TCP packets is likely to occupy the channel for the most part.

Table 3. UDP VU: Network Investment Results

Priority	Investment Throughput		Wasted Investment	
	SPM	DPM	SPM	DPM
Low	10.37	12.93	1.76	0.49
Middle	12.80	12.69	0.54	0.37
High	13.98	13.36	0.19	0.38
Overall	374.34	391.16	25.09	12.24

The following example will best illustrate the situation: Consider a chain topology of 4 nodes, arranged as shown in figure 7, with inter-node distance being 200m and two TCP flows A (node 1 to 2) and B (node 1 to 4). Even under normal circumstances, the rate of traffic out of node 1 is expected to be much higher than the rate of traffic out of node 4, partly due to the fact that node 1 is a source to two TCP flows, while node 4 is a destination to only one flow (only transmits Ack to source), and partly due to the fact that TCP packets are usually much larger in size when compared to Ack packets. Hence we can expect the channel occupancy probability of node 1 to be much higher than that of node 2. However, that is not the main problem. As we can see, nodes 2 and 3 are within carrier sense range of node 1, and hence they will not indulge in any transmissions when node 1 is transmitting (to node 2). Node 3 cannot indulge in any receptions either. But, node 4 is outside the carrier sense range of node 1, and will attempt sending Acks to node 3, as it senses the channel to be free even when node 1 is transmitting. Node 3 does not respond (to node 4) to prevent corruption of the received bits at node 2 (from node 1). Node 4 times-out and backs-off multiple times before it can actually send out the Ack to node 3. In the meanwhile, node 1 will experience a timeout for flow B due to non-receipt of Ack and will throttle its sending rate, causing immense damage to the overall throughput of flow B. We actually noticed significant improvements in flow B's throughput when we moved node 4 within carrier sense range of node 1. For higher hop counts, the probability that at least one node is outside the carrier sense range of a transmitting node is quite high, and hence we can expect poor TCP performance at higher hop counts. The overlay MAC layer solution proposed by **[4]** is expected to alleviate this problem to a large extent. Their solution proposes the use of a distributed weighted fair queuing (WFQ) scheme for fair bandwidth allocation to nodes. However, they use static weights for the WFQ. It is quite difficult to assign weights to flows in the WFQ a-priori, especially in a mobile multi-hop wireless network. One of our future studies will include using dynamic weights in the WFQ scheme, made possible by the investment function, to improve the parity in bandwidth allocation to multi-hop flows.

Table 3A. UDP VU: Fairness Measure

Metric		SPM	DPM
Pdr vs. Hop count	σ^2	0.0057	0.0009
	mean	0.8333	0.8575
E2e delay vs. Hop count	σ^2	0.0930	0.0236
	mean	0.5700	0.4733
E2e delay vs. Priority	σ^2	0.0882	0.0016
	mean	0.4167	0.4367
Inv. Throughput vs. Priority (JFI)		0.9855	0.9995

The dynamic mapping scheme alleviates the starvation problem to some extent, as evident from flow throughput values in Table 4. Though the throughput gains for each flow may not be quite substantial in an absolute sense, they are still quite substantial in a relative sense (improvement factors: around 2 for 3-hop flows, and around 17 for 4-hop flows). From Table 4A, it can be seen that both flow throughputs and investment throughput in schemes using the investment function experience moderate (IBD) to significant (3LPQ) increases in fairness (JFI) when compared to the TDD scheme. The IBD scheme is expected to only improve network utilization efficiency, and hence we do not see significant improvements in fairness relative to TDD. However, the delay variance and mean delay values are significantly reduced (almost 70% reduction in variance and 45% reduction in mean delay) in the 3LPQ scheme when compared to the TDD scheme. A more sophisticated investment function that captures dynamics of TCP may lead to a higher degree of fairness, and a much better throughput performance.

Table 4. TCP CU: Throughput and Delay

Hop Count	Flow Throughput (kb/s)			Mean End-End Delay (seconds)		
	TDD	IDB	3LPQ	TDD	IDB	3LPQ
1	65.2	62.75	53.7	0.41	0.4	0.45
2	7.43	9.81	13.75	1.11	1.18	0.83
3	4.15	2.93	6.25	1.38	1.71	0.96
4	0.18	1.58	3.07	2.12	2.06	1.32
5	0.08	0.67	1.04	3.29	3.27	2.05
6	0.69	0.08	0.14	4.52	4.29	2.68

For VU simulations, as with UDP, we compared performance of 3LPQ+SPM+TDD (no investment function) and 3LPQ+DPM+IBD (investment function). The total investment throughput values (in hops-kb/s) for the SPM and DPM schemes are 775.6 , and 780.2 respectively, while the corresponding wasted investment values are 7.91 and 6.01 respectively. Again, these gains are modest at best. Table 5 shows the throughput and delay performance under both schemes. Both schemes show poor performance for hop counts greater than 3, but as seen from Table 5A, there is a modest improvement in throughput fairness among flows with DPM compared to SPM, with a more significant improvement in delay fairness. This, we believe, is because DPM alleviates the problem of bandwidth starvation to higher hop flows to some extent. As seen from Figures 5 and 6 and Table 5A, DPM shows significant improvement in throughput and delay fairness across priorities, while at the same time maintaining some degree of service differentiation as specified by the user. For comparable mean delay values (across priorities), the delay variance with DPM is much smaller than the delay variance with SPM. As with UDP, the degree of service differentiation under DPM decreases when compared to SPM, which again can be improved by appropriate choice of i_U values.

4. Related Work

The problem of poor transport layer performance in wireless networks has been attributed to various factors such as mobility, erroneous congestion control, contention of TCP packets with ACK packets, link-layer contention (lack of bandwidth), etc.**[4-7]**. Our primary focus is on the problem of inadequate bandwidth in multi-hop wireless networks. To improve end-to-end throughput, numerous localized solutions have been proposed such as tweaking of TCP parameters, modifying 802.11 DCF, modified link-layer schemes, drastic changes in TCP architecture, etc **[4-6][8-10]**. Gupta and Kumar in their seminal work **[2],** derive bounds for the capacity of wireless networks. Our goal is to efficiently utilize network capacity, so that we inch as close as possible towards the achievable theoretic bounds. We strongly believe that end-to-end delivered throughput (sometimes called application goodput) can be substantially improved by maximizing network utilization efficiency. To our knowledge, we are the first to explore this avenue. The unifying tool that we use to achieve this objective (and others) is the *investment function*. Numerous fair channel access schemes like **[11][12][13]** have been proposed for wireless networks. Most of these are not relevant to multi-hop wireless networks, which present unique challenges that are absent in single-hop networks. These schemes either involve a centralized scheduler, which is not feasible in ad hoc networks, or do not deal with fairness in a fine-grained sense. For instance, in **[11]**, fairness is discussed in regards to preventing starvation of low-priority flows. Distributed schedulers, like the one proposed in **[12],** do not discuss fair bandwidth allocation when the network is congested. They do not provide any insight into assigning globally recognized weights, which we believe is the key to fair scheduling. The flow weights in fair queuing disciplines, such as WFQ, WCFQ **[14],** and distributed schemes such as the ones described in **[13]** can be made as functions of the mean investment value. This ensures that the weights are dynamic and provide fair treatment to all flows despite varying wireless network dynamics, while at the same time maintaining maximal network utilization efficiency. The investment function that we describe has a small degree of overlap with the price-based approach discussed in **[15][16],** but there are fundamental differences in terms of applicability, objective function and computation of price or investment .

Table 4A. TCP CU: Fairness Measure

Metric		TDD	IDB	3LPQ
Flow throughput vs. Hops (JFI)		0.2329	0.2495	0.3243
Inv. Throughput vs. Hops (JFI)		0.3430	0.3843	0.5591
Mean delay vs. Hops	σ^2	2.3237	2.0087	0.6903
	mean	2.1383	2.1517	1.3817

5. Conclusions and Future Work

The major contribution of this paper is the introduction of a new concept in networking called the *investment function*, the development of which will, we believe, address the problem of multi-hop fairness and poor bandwidth usage in multi-hop wireless networks, in addition to improving network utilization efficiency. Through our simulations, we demonstrated two of the many ways that the investment function can be used for better network performance. Our TCP and UDP simulation results with a grid topology (and earlier results with a “chain” topology, not shown here) indicate that the investment function can provide substantial improvement in throughput and delay fairness properties across multiple hops and priorities, in addition to substantial reduction in network wastage for UDP flows.

Given the flexibility of the network investment function, we plan to research the various forms of the investment function and its application to different scenarios such as provision of QoS, distributed fair bandwidth allocation, etc,. We are also currently studying a distributed scheme that allocates flow-level and node-level bandwidth in a multi-hop wireless network. Another avenue for future work is to study the suitable form of investment function to reduce TCP bias towards 1-hop flows. We also plan to study the fairness-priority tradeoff, by having a larger separation between the user investment factors. In addition to simulation, we plan to make use of probabilistic and optimization tools to more fully characterize the scarce bandwidth and hop-count disparity problems and begin to understand their root causes.

Table 5. TCP VU: Throughput and Delay Results

Hop Count	Flow Throughput (kb/s)		Mean End-End Delay (seconds)	
	SPM	DPM	SPM	DPM
1	44.99	37.66	0.46	0.54
2	11.83	16.04	0.63	0.67
3	9.87	8.73	0.88	0.82
4	0.98	2.01	3.05	1.01
5	0.24	2.19	3.79	1.65
6	0.06	0.36	4.52	2.22

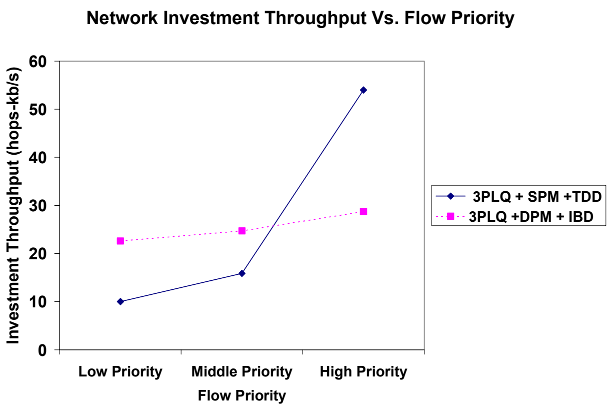

Figure 5. TCP VU: Investment vs. Flow Priority

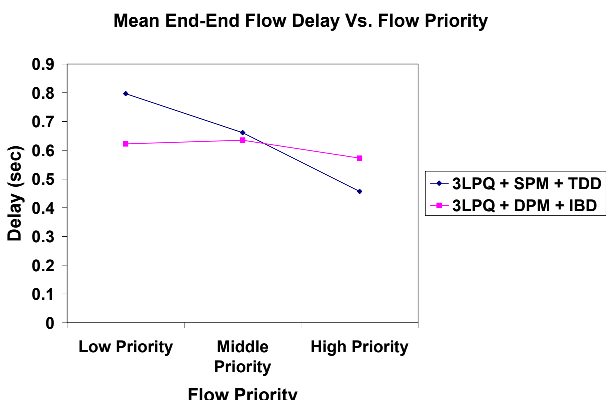

Figure 6. TCP VU: Mean Delay vs. Flow Priority

Table 5A. TCP VU: Fairness Measure

Metric		SPM	DPM
Flow throughput vs. Hops (JFI)		0.3403	0.4248
Inv. Throughput vs. Hops (JFI)		0.5158	0.6876
Mean end-end delay vs. Hops	σ^2	3.1730	0.4252
	mean	2.2217	1.1517
Inv. Throughput vs. Priority (JFI)		0.5880	0.9863
Mean end-end delay vs. Priority	σ^2	0.0336	0.0013
	mean	0.6410	0.6002

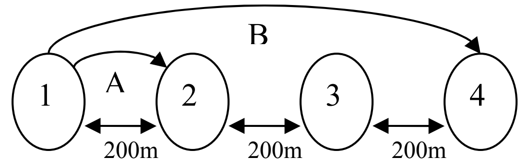

Figure 7. 4-node Chain Topology

6. References

[1] J. Li, C. Blake, Douglas S. J. De Couto, Hu Imm Lee and R. Morris, "Capacity of Ad Hoc Wireless Networks", In Proceedings of MOBICOM 2001, Rome, Italy

[2] P. Gupta, P. Kumar, "The Capacity of Wireless Networks". *IEEE Trans. on Information Theory*, 46(2):388–404, Mar. 2000

[3] R. Jain, D. Chiu, and W. Hawe, "A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems", DEC Research Report TR-301, Sept. 1984.

[4] Ananth Rao and Ion Stoica. "An Overlay MAC layer for 802.11 networks**",** In Proceedings of *Mobisys 2005*, Seattle.

[5] M. Gerla, K. Tang, and R. Bagrodia, "TCP Performance in Wireless Multi-hop Networks," *IEEE WMCSA* 1999

[6] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, M. Gerla, "The impact of multi-hop wireless channel on TCP throughput and loss", in Proceedings of INFOCOM 2003, San Francisco, CA

[7] G. Xylomenos, G C. Polyzos, P.Mahonen and M. Saaranen, "TCP Performance Issues over Wireless Links", *IEEE Communications Magazine*, vol. 39, no. 4, 2001, pp. 52-58

[8] T.Nandagopal, T.Kim, X.Gao, V.Bharghavan, "Achieving MAC Layer Fairness in Wireless Packet Networks.", In *Proc. ACM/IEEE MobiCom*, pages 87–98, August 2000.

[9] S. Lu, V. Bharghavan, and R. Srikant, "Fair Scheduling in Wireless Packet networks," IEEE/ACM Trans. Networking, vol. 7, no. 4, pp. 473-489, Aug. 1999

[10] K. Sundaresan, V. Anantharaman, H. Hsieh, and R. Sivakumar, "ATP: a reliable transport protocol for ad-hoc networks," in Proc. ACM MobiHoc '03, Annapolis, ML, 2003

[11] Mohammad Malli, Qiang Ni, Thierry Turletti, and Chadi Barakat. "Adaptive Fair Channel Allocation for QoS Enhancement in IEEE 802.11 Wireless LANs". Accepted to appear in IEEE ICC 2004 (International Conference on Communications), Paris, France, June 2004

[12] Y. Liu, S. Gruhl, and E. Knightly, ``WCFQ: an Opportunistic Wireless Scheduler with Statistical Fairness Bounds,", *IEEE Transactions on Wireless Communication,* 2(5):1017-1028, September 2003

[13] Vaidya, N.H., Bahl, P., Gupta, S, "Distributed fair scheduling in a wireless LAN", In: Proc. of ACM International Conference on Mobile Computing and Networking (MOBICOM). (2000)

[14] Y. Xue, B. Li and K. Nahrstedt, "Optimal Resource Allocation in Wireless Ad Hoc Networks: A Price-based Approach," in Tech. Report, UIUC, June 2004.

[15] Y. Qiu and P. Marbach, "Bandwidth Allocation in Ad-Hoc Networks: A Price-Based Approach," in *Proc. INFOCOM*, 2003