An Experimental Study of Multimedia Traffic Performance in Mesh Networks

In the abstract the authors make the statement that testbed experiments facilitate a better understanding of network and application characteristics. I would expect them to make the case as to why the UCSB MeshNet testbed is better than other testbeds and perhaps spend a little time on the ways in which testbeds experiments in general can provide insights not available in simulation or emulation. I will be looking for the promised video and voice traffic evaluations. I will be expecting further explanation of the network card interface configuration experiments, however, this does sound interesting. The authors promise some suggestions for improved performance, presumably this is from the NC configuration parameters and will be demonstrated by some experiments. The authors also make the statement that this study is beneficial for capacity planning and protocol design. I will be looking for this tie in. Overall a decent abstract. I think the NC configuration parameters makes the paper interesting.

Edit: Page 1, Column 1, Introduction, paragraph 2: Word spacing – effi cient --> efficient, simplifi ed --> simplified, traffi c --> traffic ***There are many of these. I am going to stop listing them.

Edit: Page 1, Column 2, Introduction, paragraph 3: Reference needs to be indexed -- [?] --> [4]

The authors make the case that rapid deployment of wireless technology and infrastructure is enabling new applications that require real time data delivery including video, VOIP, and online gaming. They also make the case that better evaluation and analysis is increasingly critical for the facilitation of robust protocol design. This is a little weak. I would like a few statements describing why robust protocol design cannot be done with current tools, i.e., why do we need the MeshNet testbed?

The authors make the case that the MAC layer in most simulators are simplified and that the PHY layer is inaccurately modeled and that evaluation in simulation may not reflect performance in real world deployments. This is true. I would also like to have seen some discussion of the fact that evaluation in testbed may not reflect real world performance either. However, space is limited and the trend of thought in 2005 was undoubtedly different than today.

Okay, here is the discussion of the limitations of the testbed environment. Difficult to configure, results affected by specific configurations, the large number of parameters makes it difficult to find representative values, and the real-world wireless media makes repeatability a problem as well as making it difficult to separate the characteristic under study from the environmental noise.

The authors make the case that multimedia and VOIP has more stringent requirements that regular data delivery. They restate that they will conduct experiments with VOIP and video over the MeshNet and study the capacity of the MeshNet nodes. They restate that they are going to try different NC configurations and that they believe this study is beneficial in both wireless capacity planning and protocol design.

This introduction was okay. It was not excellent. I don't feel like the justifications for the study were strong enough. I realize that this is a 6 page paper and that the authors are space limited but that 2nd to last paragraph in the introduction contained several sentences that were just repeats from the abstract. This space could have been used to better justify the need for this MeshNet study. For instance, the authors say (again) that they believe this study is beneficial for capacity planning and for protocol design. I already know that they believe this from the abstract. What I was looking for is some reasons why this is true. I am still willing to accept the paper, but I feel it could have been much stronger.

The testbed description is detailed enough that a competent researcher should be able to duplicate the MeshNet. The link reliability tool is interesting. I presume that this is meant to address the repeatability issue. I am not convinced that this would make experiments repeatable. Perhaps it would be effective for short duration experiments, but I think that short duration experiments are more repeatable anyways. Interesting tool, however, I would need to have some justification that this makes the testbed experiments repeatable. I mean what happens if a link becomes asymmetric or completely unusable during an experiment and then returns to symmetry before the end of the test? Okay, I get the time thing and that is the NTP algorithm but I don't understand why the authors didn't just use NTP over the LAN. NTP over the wireless will be bad (milliseconds) but over the LAN it is much better (at least in the range of 100-200 microseconds). Anyways why do we need 10 microsecond accuracy?

The description of the topology is adequate, but I would like to know why the nodes were placed in the configuration described. Was this just for convenience or is there some other reasons? Are the nodes on the 3rd floor further away than the others? Why across the hallway?

The authors use the link reliability tool to find a symmetric four hop path and then manually configure the nodes with static routes then they send video and VOIP data. The authors are going to text ARF, RTS/CTS, and maximum retransmission count on this path. This seems interesting. Could have used a table for default parameters. The metrics seem reasonable: latency, loss, jitter, and fairness.

The authors found that the packet sending rate has more effect on the number of flows that can be sent than the actual data rate. This result is interesting but could have used some more explanation. The paper shows that as the number of hops increases the latency and loss also increases. Also the network capacity is constrained by the number of hops rather than the data rate. This is because of contention and interference. The authors note that the number of flows (with hops > 1) supported with ARF is similar to the that of a fixed (2Mbps) data rate. This is because ARF keeps falling back to the low data rate because of packet loss. Also ARF seems to follow a slow start like behavior because ARF cannot adapt quickly with bursty traffic.

The authors conduct experiments to determine the fairness between video flows with ARF turned on. They find that there is a lot of variability when the network is not saturated as opposed to when the network is saturated. Also the channel capture effect is shown in that earlier flows get more of network resources than later flows when the network is not saturated. In section 4.3, the experiments indicate that additional loading induces jitter which could impact video or VOIP flows. The result indicating that RTS/CTS induces latency and loss is an important result. I believe that this was generally unknown at the time of the writing of this paper. The authors found that MAC retransmissions reduce loss but increase latency. However, these results vary widely according to media conditions and there is no single parameter that will be optimal for all conditions. Finally the authors found that under congested conditions MAC retransmissions no longer increase the delivery rate.

I give this paper an accept due to novelty. I believe that it was one of the first testbed studies in wireless. The experiments are interesting and some of the results quite significant. However, the paper could have benefited from stronger writing. The authors made their intentions clear in the abstract, but did not motivate the reasons for theses intentions in the introduction as convincingly as they could have. The description of the testbed was clear, and complete, but the presentation of results was unclear in some places and generally difficult to follow.