S. Banerjee and B. Bhattacharjee, "A Comparative Study of Application Layer Multicast Protocols"

This paper started off very promising. After seeing the title of the paper and reading through the "abstract" section, the reader gets a sense that this paper will provide a very good description and comparison of various existing application layer multicast protocols. Unfortunately, after going through the paper in its entirety, I found the information in the "abstract" section to be somewhat misleading.

The paper starts off by giving the reader a decent background of multicast and the problems it solves. It then continues by describing why we need to implement multicast at the application level rather than the network level describing reasons such as how adding multicast to the IP layer "increases the overheads and complexities at the routers" or how there is a "dearth of experience with additional mechanisms like reliability and congestion control on top of IP Multicast, which makes the ISPs wary of enabling multicasting at the network layer." The paper then describes the goal of implementing multicast at the application layer, "to construct and maintain an efficient overlay for data transmission", followed by some metrics that can be used to measure the "goodness" of the different protocols.

In the ensuing sections, the paper proceeds by describing three different kinds of application layer multicast approaches which the authors classify as "mesh-first", "tree-first", and "implicit". For each of these three, various protocols are named and described to demonstrate the difference between the classifications. For each classification, they choose one of the protocols and provide a diagram of an example topology for it along with an example of what adding a new node to the topology looks like. I do not think the authors did a good job with this. For example, when discussing the Narada protocol, the authors mention that in figure 2, panel 3, the link between node A and node C is dropped because it is "non-useful". Upon examining the figure, it is not at all clear to me why this link was dropped and the paper does not go into a good enough explanation to back up the claim that it is not useful. Furthermore, I felt like the description of figure 4 was completely useless and did not help the reader better understand the Nice protocol.

As I kept reading the paper, I found more and more areas that I thought were criticizable. For instance, in the description of the Yoid protocol it is mentioned that in the case that no "potential parent is found in this list, Yoid proposes different heuristics that can be used to find other potential parents on the tree." I thought this statement merited some sort of discussion whether it be a thorough explanation of the type of heuristic that is used or just a simple example. Also, in the description of the Scribe protocol, it seems as if the authors go on a complete tangent when they spend a few paragraphs discussing Pastry. I understand that an understanding of Pastry is necessary to understand Scribe but I feel as if some sort of subsection was needed here or at least a better introduction to Pastry and what it is used for.

Lastly, I found the "comparative study" section to be completely useless. The authors spend 7 pages giving a somewhat detailed description of a handful of different application level multicast protocols and have almost nothing to show for it. The table they provide gives the reader almost no information as to how the different protocols compare to one another in terms of performance (perhaps the metric people care about the most). I think the authors should have spent more time coming up with more insightful experiments to compare the protocols such as comparing the speed at which nodes receive content being transmitted to the different nodes or performing some kind of efficiency measurement for new nodes joining the network, what happens to the performance of the mesh network when a node leaves, etc. Having such experiments would have significantly improved this paper's contribution to the field whereas they are close to none as is since I could have read about the

different protocols individually and probably in more detail. To make matters worse, I did come across some grammatical errors and think the paper could have been more consistent in its formatting (i.e. have the same subsections for all of the protocols that were discussed). When I first began reading this paper I did not realize that it had not been published anywhere but after going through it in detail, I really am not surprised.