"An Alternative Paradigm for Scalable On-Demand Applications: Evaluating and Deploying the Interactive Multimedia Jukebox," IEEE Transactions on Knowledge and Data Engineering Special Issue on Web Technologies, vol. 11, num. 4, pp 658-672, July/August 1999.

This paper suggests a new protocol to implement scalable on-demand applications – the Interactive Multimedia Jukebox (IMJ). The paper begins by giving the reader a little bit of background as to why such a protocol is useful. It then describes some different methods for scheduling followed by a description of the different aspects included in IMJ. Throughout these descriptions, the authors provide insight about the performance of IMJ through the use of graphs and tables which display IMJ's characteristics for various metrics. As was explained in the introduction, this paper set out to "describe the jukebox paradigm and relate it to other on-demand paradigms", "describe some of the challenges in providing an on-demand program service, "describe our efforts to prototype a jukebox-based service", and "discuss some of the issues related to building the IMJ and our ongoing efforts in improving the jukebox paradigm."

I think the paper does a good job convincing the reader why a protocol such as IMJ would be useful. In my opinion, for a protocol to be considered a contribution it has to be able to provide a new technology that is not currently existent or not currently implemented and do so at a reasonable cost. This paper accomplished both of these requirements. Considering the paper was published in 1999, it is safe to assume that video on demand (VoD) was not as prevalent as it is today. This paper provides a new technology which would utilize the existing multicast backbone on the internet and allow the end user to have more interaction with their television programming. By using multicast and different scheduling mechanisms, IMJ allows the service providers to continue to broadcast (in the television sense not network sense) their usual programming while servicing the requests of many users without having too much load increase at their servers. This also gives the service providers potential for more profit since they would theoretically be charging some sort of subscription fee to use the IMJ service and could implement it without significantly increasing their expenses. IMJ also provides end users with a very nice capability which is to request certain programs without having to pay too much more for service, only the required subscription fee. The paper did not touch on the financial specifics of putting IMJ into production but I believe it gave a nice, brief overview of the potential there.

In order to have a decent comparison to other existing VoD services, the paper discusses two factors:

1. "The number of viewers who can watch a particular program stream."
2. "A combination of how much viewer input is considered in program scheduling and whether the program schedule is developed in real-time or prearranged."

I think these two factors cover a good sized spectrum of comparison when it comes to VoD services and I believe that the inclusion of figure 2 is helpful to see where the IMJ spectrum fits compared the other VoD services. Clearly the most important as far the service providers are concerned is factor 1. Factor 2 I do not see as being so important. Being able to compare the technology as it is today, ten years after the publication of this paper, we still do not have real time scheduling and I do not see this as a problem. I think the ability to have what the paper considers "true VoD" (i.e. the ability to pause, rewind, and fast forward), is more important than choosing what shows will come on next. The paper had the right idea when it talked about having a few channels dedicated to the viewers' programming choices. This is almost comparable to the on-demand channels some cable packages come with today however within the increase capacity of memory and speed of buffering at both endpoints (provider and client), the issue of waiting until a large group requests the same program seems almost irrelevant.

The paper then goes on and reports about various measurements that were performed in the IMJ test implementation. I think figure 5 provides interesting feedback. For all three experiments performed (IMJ WWW page hits, program requests, and join requests for two IMJ channels), it seems like the trend lines stay relatively consistent. Granted there is a spike at the beginning which signifies the initial launch of IMJ, I think the consistency shows that a service like IMJ would be heavily used should it be deployed worldwide, another detail that helps show the contribution this paper provides. I also appreciated the discussion of hot programs versus cold programs in section 4. This takes into account a real world situation that is the providers trying to make as many of their customers happy and exploits the problem with a voting scheduling policy.

One thing I thought the paper could have done a slightly better job on was in the analysis of figure 9 when the authors discuss the skew value. It is not clear based on the text what the significance of this value is and I think giving a more detailed description of its implication would have strengthened the impact of the results. Overall I thought the paper was well-written. I think the contribution of the paper at the time was significant but at the same time am not surprised that the protocol was not widely deployed since with the increase of bandwidth due to high speed internet connections and the increased memory capacity of cable boxes, a service like IMJ becomes unnecessary.