A Digital Fountain Approach to Reliable Distribution of Bulk Data
This paper was published at SIGCOMM in 1998, SIGCOMM had 226 papers submitted this year and accepted only 27 for an acceptance rate of 12%.

At the time of publication the need for an efficient and effective solution for distribution of bulk data to a large number of clients was on the rise. As far as I can tell, the existing erasure code based solutions did not achieve acceptable performance when it comes to encoding and decoding of data.

This paper suggests a solution called "digital fountain". The name comes from the analogy of drinking water where every drop is equal to the person drinking, as long as they get enough drops they are happy, no matter which drops they get. The main contents of this paper is the description and analysis of what they call "Tornado Codes". This is a erasure code system that has orders of magnitude better performance than the existing dominant solution, Reed-Solomon. I won't go into details as to how Tornado works in this analysis.

After showing the performance benefits of Tornado over regular Reed-Solomon, the paper goes on to compare what is called an interleaved approach with Tornado Codes. Again they show that Tornado has several benefits over the existing solution, some of these benefits are higher efficiency, less complexity (interleaved solutions require careful selection of a parameter), better performance and more scalable. The paper also shows that the performance benefits apply to real world data by comparing Reed-Solomon and Tornado on a publicly available MBone trace. Finally the paper shows a prototype of how Tornado codes can be implemented in a multicast system by implementing it on top of IP Multicast. The results from this implementation seem to be fairly preliminary, the authors admit that the code used is not identical with the proposed solution and they also mention that they have ideas for further testing.

The paper does a good job of describing Tornado codes and show their effectiveness, the authors also show how this can be used in a real multicast system. What they don't talk about is wether or not this is a solution that really fixes a real problem in today's internet, but this is not the goal of the paper either. I don't know wether or not Tornado codes has gained any traction since this paper was published, but even if it has, the contributions made by this paper to today's internet have not been huge since multicast itself has not become very popular. It seems like this is a technically sound paper, but maybe not very important in the long run.

Evaluating and Deploying the Interactive Multimedia Jukebox
This paper was published in IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, VOL. 11, NO. 4, JULY/AUGUST 1999. I was not able to find much information about this journal, but what I can say is that the paper was published by IEEE, which is a very well regarded organization.

This paper looks at the current state of Video-on-Demand systems and suggests a new solution called the Interactive Multimedia Jukebox (IMJ). In 1999 true VoD was not a feasible solution due to cost issues, a service provider would not be able to make a profit on a system where every client needed several gigabytes of storage and bandwidth and only paid a few dollars to rent a movie. The paper starts out with a look at the different types of scheduling paradigms and their associated pros and cons. I found this discussion to be a bit obvious, the graph (Fig. 1) seems especially superfluous. The paper then goes on to describe the IMJ, which basically is a hybrid between true and near VoD.

"The paradigm is based on three properties: 1) A set of channels are multicast to all viewers tuned to the respective channel.
2) Viewers may watch a program playing on any channel or make a request for something of their own choosing. Viewers' requests are scheduled on one of the jukebox's channels using some scheduling criteria. 3) A schedule of currently playing and scheduled programs, updated in real-time, is available to all viewers. Viewers can watch any program, including those scheduled by others, by tuning to the appropriate channel."

Requests are sent via WWW, while the video itself is transferred via MBone. The paper also describes some of the options a service provider has for tuning the IMJ system to fit their needs. The paper goes through some real world data collected from an implementation of the IMJ service. The data shows that the max wait time observed was over 30 minutes, which sounds a bit high to me, I would not be willing to wait that long. Some of the graphs (like Fig. 10) could probably be conveyed just as effectively without the figures.

This paper is very dated, Moore's law has made cheap true VoD a reality in today's internet. There are several solutions that provide this, for example Zune Video on the Xbox 360. The solutions suggested in this paper are more bad-aids/temporary workarounds than new technical progress that forwards the state of the art.