Multimedia Systems Research: The First Twenty Years and Lessons for the Next Twenty

Prashant Shenoy, University of Massachusetts Amherst

This retrospective paper examines the past two decades of multimedia systems research through the lens of three research topics that were in vogue in the early days of the field and offers perspectives on the evolution of these research topics. We discuss the eventual impact of each line of research and offer lessons for future research in the field.

Categories and Subject Descriptors: H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems—*Multimedia streaming*

General Terms: Design, Performance

Additional Key Words and Phrases: Multimedia systems

1. INTRODUCTION

Multimedia systems research as a field has evolved significantly since its early days in the 1990s. A look back at the research topics that were in vogue in those early days offers numerous broader lessons for future researchers.

Research as a discipline is always forward-looking, and as such, it is inherently difficult to predict whether a research idea will eventually pan out in terms of impacting practice and making its way into products. I present three examples of "hot" research topics from the 1990s and offer my perspective on whether these have "failed" or "succeeded." There are many such examples, and the choice of these three examples is based on personal experience in conducting research in these topics and having watched their evolution over the past two decades. These three topics have yielded different outcomes in terms of impacting practice: one was widely believed to be a failure in the early days but is now a thriving success, another saw successful early adoption by industry but eventually failed to get market traction, and the third ended up influencing a different use-case from the one for which it originally envisioned. Collectively, these examples corroborate the unpredictable nature of research but also offer some useful lessons to today's researchers.

2. A RETROSPECTIVE LOOK AT THREE MULTIMEDIA SYSTEMS RESEARCH PROBLEMS

Video on Demand: Video on Demand (VOD) and streaming servers were one of the early hot research topics in the field in the 1990s [Sitaram and Dan 1999]. Dozens of papers appeared on this topic in

Author's address: P. Shenoy, School of Computer Science, 140 Governors Drive, University of Massachusetts, Amherst, MA 01003; email: shenoy@cs.umass.edu

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

multimedia conferences [Bolsoky and et. al. 1996; Martin et al. 1996; Mayer-Patel and Rowe 1997; Shenoy et al. 1998; Vernick et al. 1996] and high-profile VOD trials were launched by large companies such as Time Warner in the mid 1990s. However the trials were soon deemed to be a failure and the conventional wisdom at the time was that the field of VOD had reached a dead end. The primary reasons for the failure of VOD trials was the lack of broadband networks to deliver video content to end-users and the large capital investment needed to deploy such a network was a major deterrent. The presumed failure of VOD was disheartening to researchers at the time, and there was even a bias among conference review committees to "move on" to other topics. Fast-forward to today and we see that on-demand streaming services are now a thriving success. Youtube, Netflix, TiVo, and many other companies have built successful businesses around video streaming; cable companies routinely offer on-demand streaming services. Internet streaming is, in fact, threatening traditional models of watching television, cable, and movie content. One may ask how a field that was assumed to be doomed by conventional wisdom ended up as a multi-billion dollar industry. The primary reason is the relentless advances in technology—network infrastructure advancements eventually led to widespread deployment of home broadband Internet services, which in turn, made it feasible to deploy video streaming services of various sorts. User expectations was another factor—users found utility in viewing small video clips and did not expect Internet-streamed content to be of the same quality and reliability as, say, broadcast television, which made it possible to deploy streaming services at a lower cost and complexity (e.g., using existing IP-based networks).

UDP Streaming Protocols: The design of network streaming protocols was another hot research topic of the 1990s [Amir and McCanne 1995; Yavatkar and Manoj 1994; Gong and Parulkar 1994]. The conventional wisdom at the time held that TCP was not suitable for video streaming since its congestion control and packet retransmission mechanisms were "incompatible" with the timeliness constraints imposed by video playback (late packets were no better than lost packets was the general premise). UDP-based streaming was deemed to be "correct" approach and many adaptive streaming protocols and mechanisms based on UDP were proposed [Padhye et al. 1999; Schmitt et al. 2002; Talley and Jeffay 1996]. UDP-based protocols such as RTP, RTSP [Schulzrinne et al. 1998] and others were standardized by the IETF. These ideas were quickly embraced by industry and companies such as Real Networks and Microsoft designed their own versions that were implemented into server products and clients players such as RealPlayer and Windows Media Player. Research on UDP streaming was held to be a success due to its quick adoption by industry. However, despite this adoption and incorporation into commercial products, there was a lack of market adoption, and today much of video streaming occurs over TCP rather than UDP. TCP streaming over HTTP has become the de-facto approach for delivering video content, despite the drawbacks of TCP's congestion control and retransmission mechanisms for streaming. Again, one may ask how a field that was deemed to be success ended up with a lack of broad market adoption. The answers may be found in technological advances and the convenience of using the broadly-deployed HTTP protocol. Advances in broadband network infrastructure meant that the bandwidth needed to stream typical video content were commonly available to end-users. Memory capacities on clients devices have also grown in accordance with Moore's law, and these devices can easily support larger buffers needed to mask the jitter due to TCP's rate control mechanisms. Finally the convenience of delivering video using the widely-deployed HTTP protocol and through existing web browsers (rather than stand-alone players) accelerated the adoption of TCP streaming. Today dynamic adaptive streaming over HTTP (DASH) streaming continues to be both an active area of research and is widely used in practice.

Quality of Service in Operating Systems: Quality of Service (QoS) support in operating systems was yet another active area of research in the 1990s. Numerous QoS-aware OS schedulers were designed to deliver better resource management support and service quality to multimedia applications [Steinmetz and Nahrstedt 1995; Jeffay et al. 1992]. The main premise of this line of work was that multimedia applications such as client players and streaming servers required guaranteed resources and needed to be isolated from other non-multimedia applications. Fair-share schedulers are an example of such QoS-aware mechanisms for allocating CPU, network interface and disk resources to competing applications in a fair, predictable fashion, rather than in a best-effort manner [Goyal et al. 1996; Keshav 1995; Lin et al. 1998]. While some of these schedulers were implemented into open-source operating system kernels such as Linux [Sundaram et al. 2000] and FreeBSD [Bruno et al. 1998] (and a few commercial ones such as SGI's IRIX [Sgi 1999]), they did not see broad adoption by general-purpose operating systems such as Windows and Mac OS X. Once again, advances in technology made the need for resource isolation and quality of service guarantees less important. For instance, there was sufficient CPU capacity and network bandwidth on modern multi-core hardware to easily fulfill the needs of both multimedia and non-multimedia applications. Lack of resource contention on resource-rich hardware meant that QoS-aware mechanisms were unnecessary for most users. However this has not meant that this line of research has been a failure. Rather these mechanisms have been adopted and are in use for a different scenario than the one envisioned by multimedia researchers in the 1990s. Today, resource-isolation mechanisms such as fair-share schedulers are used by virtual machine platforms such as VMWare and Xen; such platforms allow a certain amount of resources to be allocated to resident each virtual machine, and QoS-aware mechanisms provide a natural way to enforce these allocations. Thus, these OS mechanisms have found use in a different context, and they have been adopted by virtual machine hypervisors, which act as the "kernels" on virtualized hardware.

3. LESSONS LEARNED

These examples offer several lessons for the future; these lessons are not specific to the multimedia field and apply to other research disciplines within Computer Science as well.

First, there is typically a ten to fifteen year time lag for a research idea to mature and get incorporated into products. We observe a similar time lag to maturity in the above examples. While this time-frame may be shrinking as technological advances accelerate, this "idea to product" time-frame is still long, which means one must not rush to judgment on the (eventual) impact of a research idea. As always, many research ideas will fail, some will succeed, and still others will be adopted in new ways not envisioned by their original designers. The key take-away lesson is that the long time-frame to maturity means it is nearly impossible to predict which ones will fall into which category. As a corollary, it is not always wise to "follow the crowd"—one must not abandon a line of research simply because the crowd has moved on, since eventual impact (or lack thereof) will be only known many years later.

Second, research problems that get harder due to technological advances offer more potential for long-term impact than ones that get simplified due to technological advances. For instance, the need for OS-level QoS support was a short-term problem where the application needs overwhelmed the resources available on typical machines at that time; with advances in hardware capabilities, these resource constraints were mitigated, and the mechanisms became less useful. In contrast, research problems that are independent of technological advances or those that get harder with advances in technology (e.g., building larger, more scalable or more complex systems, or testing complex software systems for security bugs) are more likely to remain relevant or useful in the long-run.

In summary, there have been stunning advances in the first two decades of the field of multimedia systems. Use of multimedia by everyday users has become common, whether is producing content using a plethora of devices or consuming them using myriad services. The future is equally bright; new trends such as mobile and social are changing how users create, share and view content; with these advances come associated challenges ranging from mobile streaming, privacy, and new modes of interaction. Our community will be well served by drawing upon the lessons from the past as we embark on future research in these exciting areas.

Acknowledgments

Ideas presented in this paper were influenced by discussions with Harrick Vin and Dan Swinehart over the years.

REFERENCES

AMIR, E. AND MCCANNE, S. 1995. An application level video gateway. In *Proceedings of ACM Multimedia'95, San Francisco, CA*.

BOLSOKY, W. AND ET. AL. 1996. The Tiger Video Fileserver. *Proceedings of 6th Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, 97–104.

BRUNO, J., GABBER, E., OZDEN, B., AND SILBERSCHATZ, A. 1998. The eclipse operating system: Providing quality of service via reservation domains. In *Proceedings of the USENIX Technical Conference*. 235–246.

GONG, F. AND PARULKAR, G. 1994. A two-level flow control scheme for high speed networks. *J. High Speed Networks 3,* 3, 261–284.

GOYAL, P., VIN, H., AND CHENG, H. 1996. Start-time fair queuing: A scheduling algorithm for integrated services packet switching networks. In *Proceedings of ACM SIGCOMM (SIGCOMM'96), Palo Alto, CA*. 157–168.

JEFFAY, K., STONE, D. L., AND SMITH, F. D. 1992. Kernel support for live digital audio and video. *Computer Communications 15*, 388–395.

KESHAV, S. 1995. On efficient implementation of fair queuing. *Journal of Internetworking Research 2*, 157–173.

LIN, C., CHU, H., AND NAHRSTEDT, K. 1998. A soft real-time scheduling server on the windows nt. In *Proceedings of the 2nd USENIX Windows NT Symposium, Seattle, WA*.

MARTIN, C., NARAYAN, P. S., OZDEN, B., RASTOGI, R., AND SILBERSCHATZ, A. 1996. The fellini multimedia storage server. *Multimedia Information Storage and Management , Editor S. M. Chung, Kluwer Academic Publishers*.

MAYER-PATEL, K. AND ROWE, L. 1997. Design and Performance of the Berkeley Continuous Media Toolkit. In *Proceedings of SPIE/ACM Conference on Multimedia Computing and Networking (MMCN), San Jose, CA, USA*. 194–206.

PADHYE, J., KUROSE, J., TOWSLEY, D., AND KOODLI, R. 1999. A model based tcp-friendly rate control protocol. In *Proceedings of IEEE NOSSDAV'99, Basking Ridge, NJ*.

SCHMITT, J., ZINK, M., THEISS, S., AND STEINMETZ, R. 2002. Improving the Start-Up Behaviour of TCP-Friendly Media Transmissions. In *Proceedings of the International Network Conference (INC'02), Plymouth, UK*. 73–80.

SCHULZRINNE, H., RAO, A., AND LANPHIER, R. 1998. RFC 2326 - Real Time Streaming Protocol (RTSP). Standards Track RFC.

Sgi 1999. React: Irix real-time extensions. Silicon Graphics, Inc., http://www.sgi.com/software/react.

SHENOY, P. J., GOYAL, P., RAO, S. S., AND VIN, H. M. 1998. Symphony: An integrated multimedia file system. In *Proceedings of the SPIE/ACM Conference on Multimedia Computing and Networking (MMCN'98), San Jose, CA*. 124–138.

SITARAM, D. AND DAN, A. 1999. *Multimedia Servers: Applications, Environments and Design*. M. Kaufman.

STEINMETZ, R. AND NAHRSTEDT, K. 1995. *Multimedia Computing, Communications and Applications*. Prentice Hall.

SUNDARAM, V., CHANDRA, A., GOYAL, P., SHENOY, P., SAHNI, J., AND VIN, H. 2000. Application Performance in the QLinux Multimedia Operating System. In *Proceedings of the Eighth ACM Conference on Multimedia, Los Angeles, CA*.

TALLEY, T. AND JEFFAY, K. 1996. A general framework for continuous media transmission control. In *Proceedings of the 21st IEEE Conference on Local Computer Networks, Minneapolis, MN*. 374–383.

VERNICK, M., VENKATRAMINI, C., AND CHIUEH, T. 1996. Adventures in building the stony brook video server. In *Proceedings of ACM Multimedia'96*.

YAVATKAR, R. AND MANOJ, L. 1994. End-to-end approach to large-scale multimedia dissemination. *Computer Communications 17,* 3, 205–217.