The Case for Reexamining Multimedia File System Design

Position Statement

Prashant Shenoy
Department of Computer Science,
University of Massachusetts,
Amherst, MA 01003.
shenoy@cs.umass.edu

Research in multimedia file systems—file systems that support different types of data and applications with diverse performance requirements—has stagnated since the commercial failure of video-on-demand systems in the 1990s. However, the past two years have seen an acceleration in the development and deployment of several new techniques aimed at web-based multimedia applications. These technologies have opened up a fertile research ground for new file system optimizations and rekindled interest in multimedia file systems. In this paper, we describe some of these technology trends, examine their implications on file system design, and present a sampling of the new research problems that arise due to these technological changes. Based on these observations, we argue for a renewed effort in designing next-generation multimedia file systems.

Consider the following technology trends that are fueling the growth of web-based multimedia applications.

- *Proliferation of streaming media content:* Although envisaged in the early 1990s in the context of video-on-demand servers, it is only in the past two years that streaming media has gained widespread use (see Table 1). Today, an increasing fraction of web sites provide content in Real Audio/Video, Quicktime, or Windows Media formats. Audio files in MP3 format are immensely popular (although illegal) on most university campuses. This trend is only likely to accelerate with the advent of inexpensive multimedia PCs and widely available, free software decoders.

- *Broadband to the home:* Concurrent to the growing use of streaming media, technologies such as cable modems and digital subscriber lines (DSL) have emerged, which promise high-speed Internet access to home users. In contrast to the slow 56 Kbps analog modem speeds, these technologies provide either a 10 Mbps shared connection (in case of cable modems) or a 384 Kbps dedicated connection (in case of DSL modems). Table 2 depicts the projected growth of users with high-speed Internet access. The broad availability of higher bandwidth connections is resulting in increasing use of bandwidth-starved applications such as 3D multi-player games, online virtual worlds, and high-quality continuous media.

- *Proliferation of networked devices:* The past few years have seen a sharp increase in the use of networked personal digital assistants (PDAs). Popular PDAs, such as Palm VII, come equipped with a low bandwidth wireless connection that enable them to instantly access information on the Internet. Moreover, many cellular

Table 1: Streaming Media Statistics (from www.internetnews.com)

No. of users per month with streaming media software	50 million
No. of live Radio webcasters	1997: 36, 1999: 1652, 2000: 2615
No. of live TV webcasters	58 stations
No. of on-demand TV webcasters	34 stations
No. of web pages with streaming media	400,000

Table 2: Projected growth in number of users with high-speed Internet access (PC Computing, April 2000)

Year	Cable Modem	DSL	Wireless broadband
2000	1.6 million	0.78 million	52,000
2001	3.1 million	1.7 million	160,000
2002	5.3 million	3.1 million	337,000
2003	8.3 million	5.0 million	445,000

phone providers now offer wireless data services that enable a user to access the Internet via a mobile computer or a custom browser built into the phone. A key characteristic of these networked devices is that they are *resource-scarce*—they have significantly smaller processing and display capacities as compared to a typical desktop as well as a lower bandwidth network connection (see Table 3)

- *Emergence of storage area networks:* Traditionally, storage servers have employed the server-attached disk architecture, in which disks are locally attached to servers and clients access data on disks via the server (see Figure 1(a)). Recently, a new storage architecture has evolved that envisions a separation of storage devices from servers [1, 2, 7, 10]. This architecture consists of a *storage area network (SAN)* to which storage devices such as disks are attached; servers access these devices via the storage area network (see Figure 1(b)). Such a network-attached disk architecture is markedly different from the traditional server-attached disk architecture: (i) it allows clients to directly access data from disks without the server in the data path; (ii) it is fault tolerant, since a server failure can be handled by employing another server to manage the set of disks.

We believe that these technology trends will profoundly impact how file systems will be designed and used in the next decade. In what follows, we examine the impact of these trends on file system design.

The first two technology trends—the growing use of continuous media files and the availability of high-speed Internet access—argue for file systems that can efficiently handle demanding continuous media applications. Much of the research on video-on-demand servers in the past decade envisaged the use of a separate special-purpose server for continuous media files and a traditional file system for handling textual files. In contrast, current users are unwilling to bear the inconvenience of storing and managing different types of files on different servers. Moreover, system administrators are looking to simplify administration tasks by reducing the number of disparate systems they need to maintain [3]. This argues for the design of a single file system that integrates the storage of different types of data and efficiently services different types of applications (e.g., streaming applications, traditional client-pull applications). Such file systems are referred to as *integrated multimedia file systems*. Recently several such integrated file

Table 3: Characteristics of PDAs and typical desktops.

Characteristic	PDA (Palm Pilot)	Desktop
Processor	16.6 MHz Motorola 68328	700MHz Intel Pentium III
Memory	8MB	128MB
Display	150x150, 1 bit color	1280x1024, 16 bit color
Network	9.6 Kbps wireless	100Mb/s Ethernet

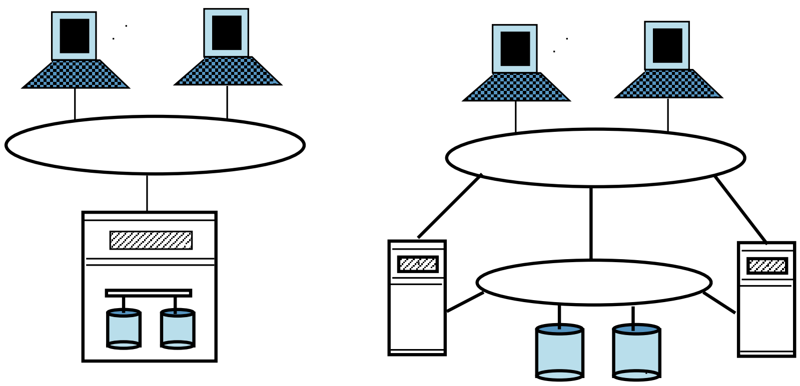

Figure 1: The server-attached disk architecture and network-attached disk architecture for file system design.

systems have been designed [4, 5, 6, 9]. Performance studies have shown that, in addition to increased convenience to users, sharing of resources in such file systems yields significant performance benefits over a partitioned architecture employing separate servers [8]. Based on these trends, we believe that an increasing number of conventional file systems will begin providing native file system support for data types such as streaming media, thereby migrating towards an integrated architecture.

Most existing file systems, whether integrated and partitioned, are based on the traditional server-attached disk architecture. The emergence of storage area networks will require us to rethink the policies and mechanisms employed by existing file systems. In a storage area network, disks are separate from servers, resulting in an inherently distributed architecture. Moreover, network-attached disks have embedded microprocessors, enabling the file system to offload simple file system tasks to disks. Doing so results in a partitioning of file system tasks, some of which are handled by the server, others by disks. To illustrate, a recent study has shown that a file system can offload the placement policy to the on-disk microprocessor and improve performance—the on-disk microprocessor can determine the current disk head location and place data close to this location so as to reduce head movement and reduce the latency of write requests [11]. A further advantage of storage area networks is that they enable users to directly communicate with disks and eliminate the server from the data path (much as direct memory access (DMA) devices enable an I/O device to transfer data directly to memory and eliminate the CPU from the data path). This key architectural change has fundamental implications on the policies for placement, retrieval and caching employed by a file system. For instance, network striping of data on disks in a storage area network will involve different tradeoffs as compared to striping files in a centralized file system. Similarly, the efficacy of a buffer cache at a server will need to be reexamined if users bypass the server and communicate directly with storage devices. Further, file system interfaces will need to be appropriately extended to enable users to communicate with storage devices (in addition to communicating with the server). These examples illustrate the need for redesigning integrated file systems so that they can take advantage of the features offered by a network-attached disk architecture.

Clearly, an integrated file system that employs the network-attached disk architecture can efficiently meet the needs of a large number of users and demanding multimedia applications. However, with the proliferation of networked devices (e.g., PDAs), these file systems will not only need to scale up to demanding applications, but also *scale down* to accommodate resource-scare end-hosts. Moreover, the file system protocols employed by such file systems will need to be redesigned to address this heterogeneity in end-host capabilities. On the other hand, file system protocols (e.g., NFS) will need to be extended to handle streaming media. On the other hand, these protocols will need to be more lightweight to handle resource-scare end hosts (protocols such as NFS that require the entire state information to be sent with every request may impose excessive communication overheads on such hosts). Employing a proxy to act as an intermediary between a server and an end-host is one technique for handling this mismatch. However, we believe that native file system support is required to handle this heterogeneity in the long term. This is because the number of resource-scarce networked devices is likely to far exceed the number of high-end workstations accessing a server.

In conclusion, several technology trends are likely to impact file system design in the next decade. In particular, (i) the growing use of streaming media, (ii) easy availability of high-speed Internet access, (iii) proliferation of networked devices, and (iv) emergence of storage area networks are four trends that will dictate how file systems of the future are built and used. Based on these observations, we argued that next generation file systems should: (i) employ an integrated architecture, (ii) be optimized for the network-attached disk architecture, and (iii) should efficiently support the heterogeneity in end-host requirements. We provided a sampling of research problems that need to be addressed to meet these requirements. We believe that these issues provide a fertile ground for future research on multimedia file system design and are currently designing a file system to address these challenges.

References

[1] G A. Gibson et. al. A Case for Network-Attached Secure Disks. Technical Report CMU-CS-96-142, School of Computer Science, Carnegie Mellon University, 1996.

[2] G A. Gibson et. al. A Cost-Effective, High-Bandwidth Storage Architecture. In *Proceedings of the 8th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-VIII), San Jose, CA*, pages 92–103, Oct 1998.

[3] R. Golding, E. Shriver, T. Sullivan, and J. Wilkes. Attribute-managed Storage. In *Proceedings of the Workshop on Modelling and Specification of I/O (MSIO), San Antonio, TX*, October 1995.

[4] R. Haskin. Tiger Shark–A Scalable File System for Multimedia. *IBM Journal of Research and Development*, 42(2):185–197, March 1998.

[5] M. Holton and R. Das. XFS: A Next Generation Journalled 64-bit File System with Guaranteed Rate I/O. Technical report, Silicon Graphics, Inc, Available online as http://www.sgi.com/Technology/xfs-whitepaper.html, 1996.

[6] C. Martin, P. S. Narayan, B. Ozden, R. Rastogi, and A. Silberschatz. The Fellini Multimedia Storage System. *Journal of Digital Libraries*, 1997.

[7] Storage Networking: The Evolution of Information Management. Technical report, Seagate Technology, Inc., Available online as http://www.seagate.com/corp/vpr/literature/papers/sn.shtml, 1998.

[8] P. Shenoy, P. Goyal, and H M. Vin. Architectural Considerations for Next Generation File Systems. In *Proceedings of the Seventh ACM Multimedia Conference, Orlando, FL*, November 1999.

[9] P J. Shenoy, P. Goyal, S S. Rao, and H M. Vin. Symphony: An Integrated Multimedia File System. In *Proceedings of the SPIE/ACM Conference on Multimedia Computing and Networking (MMCN'98), San Jose, CA*, pages 124–138, January 1998.

[10] The Intelligent Storage Network. Sun Microsystems, Inc., http://www.sun.com/storage/vision, 1998.

[11] R Y. Wang, T E. Anderson, and D A. Patterson. Virtual Log Based File Systems for a Programmable Disk. In *Proceedings of third symposium on Operating systems design and implementation (OSDI), New Orleans, LA*, pages 29–43, Feb 1999.