Research Contributions of The UCSD Multimedia Laboratory

P. Venkat Rangan

Multimedia Laboratory
Department of Computer Science and Engineering
University of California, San Diego
La Jolla, CA 92093

January 25, 1994

The Multimedia Laboratory at UCSD was founded and established in 1990 by *Professor P. Venkat Rangan* as a center for research in the following areas arising out of the synergies between computing and media communication:

1. MULTIMEDIA ON-DEMAND SERVICES (integrated access to non-live media sources)
   (a) *Storage*
      - **Storage models and structures for digital continuous multimedia**

   (b) *Retrieval*
      - **Admission control** algorithms for servicing multiple users simultaneously
      - Techniques and protocols for **synchronization** between **multiple media** during retrieval over heterogeneous networks

   (c) *Personalized Multimedia Services*
      - Architectures for supporting **personalized multimedia dial-up services**
      - **Information caching** algorithms for exploiting storage-network trade-offs

2. MULTIMEDIA COLLABORATION SERVICES (integrated access to live sources)
   (a) *Collaboration Management*
      - **Multi-level conferencing** paradigm for collaborations
      - Architectures for integration of **synchronous and asynchronous collaborations**

   (b) *Communication Management*
      - Network architectures and algorithms for **media mixing**
      - **Frame-Induced Packet Discarding** for efficient video communication over ATM networks

The main research contributions in each of the above areas are outlined below:

Storage Models and Structures for Multimedia

GOAL: To develop multimedia-on-demand information servers capable of providing services similar to those of a neighborhood videotape rental store, over metropolitan area networks.

1. **Constrained Allocation Model**

   - *Previous Approaches*: Use either unconstrained or contiguous allocation of media blocks on disks. Unconstrained allocation cannot guarantee continuous retrieval; contiguous allocation of media blocks is fraught with inherent problems of fragmentation, and can entail enormous copying overheads during editing. Thus, contiguous allocation of media blocks, although suitable for read-only storage systems (e.g., compact discs, CLVs, etc.), is not viable for interactive, read-write storage systems.
   - *Our Idea*: Constrained allocation maintains the access time of media blocks within the real-time playback requirements of strands (which refer to continuously recorded sequence of media samples) by bounding the separation between successive media blocks on disk. We have developed a model that relates disk and device characteristics (such as, disk read/write latency, and video capture/display times) to the recording rate, and derives the storage **granularity** (the size of each block), and the lower and the upper bounds on **scattering** (i.e., the separation between successive media blocks) that result in continuous retrieval. Whereas the lower bound on scattering determined using the tolerable copying overheads (within any user-specifyable limit) during editing, the upper bound is determined frm the requirement for continuous retrieval.
   - *Publications*:

     (a) P. Venkat Rangan and Harrick M. Vin
     "Designing File Systems for Digital Video and Audio"
     In Proceedings of the 13th *ACM Symposium on Operating Systems Principles (SOSP '91)*,
     Operating Systems Review, Vol. 25, No. 5, Pages 81-94, October 1991

2. **Merging**: In order to efficiently utilize the disk space, a new incoming media strand is stored on disk such that its storage pattern (i.e., the media block and the gap sizes, as defined by the granularity and scattering parameters) is merged into those of media strands already existing on the disk, by storing the new strand's media blocks in the gaps of existing strands.

   - *Previous Approaches*: Algorithms have been proposed for on-line merging of strands with fixed storage patterns. In these algorithms, merging is possible only if the exact pattern of an incoming media strand can be maintained even after its media blocks are placed in gaps of existing strands, making the algorithms highly restrictive in applicability. Relevant mostly for storage of homogeneous media strands.
   - *Our Idea*: Algorithms for merging an incoming strand with any storage pattern into any set of strands already existing on disk, by allowing finite perturbations/alterations in the incoming strand pattern so as to make it fit into existing strand patterns. However, the deleterious effects of such perturbations on continuity of retrieval of the incoming strand are smoothed out by read-ahead and buffering, both of which are precisely computed so as to exactly cancel those

effects. We have developed an **off-line merging algorithm** that can be applied a priori to the storage of a set of media streams before any of them have been placed on disk, and an **on-line algorithm** suitable for merging a new media stream into a set of already stored media streams, both of which have been shown to yield **optimal storage space utilization** in their respective categories.

- *Publications*:

  (a) P. Venkat Rangan and Harrick M. Vin
  "Efficient Storage Techniques for Digital Continuous Multimedia"
  To appear in the *IEEE Transactions on Knowledge and Data Engineering*, August 1993

Admission Control

GOAL: To service a large number of multimedia storage/retrieval requests simultaneously without violating the continuity requirements of any of the requests.

- *Previous Approaches*: Existing real-time scheduling techniques become intractable on all but tiny media strands.

- *Our Ideas*:

  1. Using the property of periodicity in both the display and (constrained) placement of media strands, we first developed a simple **round-robin servicing algorithm** that proceeds in rounds, retrieving a fixed number of blocks of each strand in each service round. The number of blocks retrieved per round, same for all strands, is computed to be the minimum necessary for satisfying the continuity of display of the strand with highest playback rate.
  2. We then developed an **optimal** servicing algorithm called **QPMS: Quality-Proportional Multi-Subscriber servicing**, that is shown to satisfy the maximum number of simultaneous requests. In each service round, the QPMS algorithm tries to retrieve, for each media strand, media blocks proportional in number to the playback rate of the strand. However, more often than not, this number may turn out to be a fraction, in which case, the fractional block retrieved cannot be displayed, thereby causing discontinuities. To address this problem, we have developed a **technique of staggered toggling** (among the various media strands) between the ceiling and the floor of the fraction, so that optimality and continuity properties are both guaranteed. The algorithm is powerful enough to accommodate **bounded availability of display buffers**, and permits **dynamic additions and deletions of requests** in a transparent manner (i.e., without causing discontinuity in the retrieval of any of the existing requests). Initial performance estimations have demonstrated that QPMS is two orders of magnitude more scalable compared to straightforward multiplexing techniques such as servicing one request per disk head.

  Publications:

  1. P. Venkat Rangan, Harrick M. Vin, and Srinivas Ramanathan
  "Designing an On-Demand Multimedia Service"
  IEEE Communications Magazine, Vol. 30, No. 7, Pages 56-64, July 1992

2. Harrick M. Vin and P. Venkat Rangan
"Designing a Multi-User HDTV Storage Server"
IEEE Journal on Selected Areas in Communications - Special Issue on High Definition Television and Digital Video Communications, Vol. 11, No. 1, Pages 153-164, January 1993

Inter-Media Synchronization

GOAL: Multimedia objects may be composed of multiple media streams such as audio and video. Display of such objects must proceed so as to not only maintain continuity of each of the constituent media streams, but also preserve the temporal relationships among them, even when each of the streams are played back at a different display site on the network. We have addressed the problems of providing synchronous access to distributed multimedia on-demand services over heterogeneous networks.

- *Previous Approaches*: Earlier work has focussed mainly on media synchronization in the presence of global clocks or in single-site multimedia workstations.
- *Our Ideas*:
  1. **Feedback Mechanisms**: We have developed feedback techniques and protocols for inter-media synchronization in the presence of network jitter and non-deterministic mismatches in playback rates of media capture/display sites. We have devised algorithms by which, in order to facilitate synchronization, a multimedia server creates a relative time system in which, media units (from different media capture sites) whose recording times are within a bounded time window are assigned the same relative time stamp while they are being stored. The rate of advance of the relative time system is driven by a *master* media source that can be chosen on the basis of application. During retrieval, the media display sites periodically transmit light-weight messages called feedback units back to the multimedia server. Feedback units are transmitted concurrently with the playback initiation of media units, and hence, the multimedia server uses estimates of transmission times of feedback units for detecting impending asynchronies at display sites. The server (which also serves as the synchronizer) then preventively readjusts media transmission so as to steer the media display sites back to synchrony. We have shown that media synchronization is not equivalent to clock synchronization, and that, clock synchronization is neither necessary nor feasible in the playback of stored multimedia objects whose media components may have been recorded possibly at different sites and different times.
  2. **Aggressive, Conservative, and Probabilistic** policies for resynchronization: We have investigated various resynchronization policies ranging from aggressive (which reacts as soon as there is even a slight chance that playback may become asynchronous) to conservative (which reacts only when playback is guaranteed to become asynchronous). Between these two extremes are probabilistic policies making use of statistical distributions of network delays and playback rate variations. Aggressive policies exhibit high reactivity and oscillatory behaviour at lower asynchronies but show a dramatic improvement in performance at higher asynchronies. Conservative policies perform best at lower levels of asynchrony, but degrade badly at higher

asynchronies. In sharp contrast to both of the above deterministic policies, probabilistic policies perform uniformly well at all levels of asynchrony.

3. **Minimum Feedback Rate** computation: Given the maximum tolerable asynchrony, we have presented methods to determine the minimum rate at which feedback units must be transmitted to enable a multimedia server to readjust transmission so as never to let asynchrony go beyond the tolerable limit.
4. **Adaptive Control of Feedback Transmission**: We have devised strategies by which a synchronizer can dynamically predict how soon playback may be expected to go out of synchrony at a display site, and furthermore, based on its most recent prediction, adaptively control the latest instant at which feedback units can be transmitted by a display site, so as to prevent asynchrony from exceeding the user-specified tolerable limit ever.

Publications:

1. Srinivas Ramanathan and P. Venkat Rangan
   "Adaptive Feedback Techniques for Synchronized Multimedia Retrieval over Integrated Networks"
   IEEE/ACM Transactions on Networking, Vol. 1, No. 2, Pages 246-260, April 1993
2. P. Venkat Rangan and Srinivas Ramanathan
   "Feedback Techniques for Continuity and Synchronization in Multimedia Information Retrieval"
   To appear in *ACM Transactions on Information Systems*, 1993
3. P. Venkat Rangan, Srinivas Ramanathan, and Thomas Kaeppner
   "Performance of Inter-Media Synchronization in Distributed and Heterogeneous Multimedia Systems"
   To appear in *Computer Networks and ISDN Systems*, 1993

Personalized Multimedia Dial-Up Services

GOAL: Multimedia services are witnessing a radical shift from the conventional broadcast mode of service to an on-demand mode of service. We are exploring how on-demand services can further evolve into *personalized* multimedia services that are automatically customized to suit the individual needs and preferences of users.

- *Previous Approaches*: Earlier efforts are directly mainly towards increasing the number of multimedia entertainment channels available to the end user.

- *Our Ideas*:

  1. We have developed a system architecture in which personalized multimedia services are provided by a symbiosis of *Storage providers*, who manage information storage at multimedia servers (a role akin to that of video rental stores and libraries in today's context), *Network providers*, who are responsible for media transport over integrated networks (a role similar to that of the telephone and cable companies of today), and *Service providers* such as publishing houses, news

distributors, entertainment houses, radio, and television stations. In this architecture, Personal Service Agents (PSAs) play a central role in tailoring the fabric of multimedia services to suit the individual needs and preferences of users. They employ content analysis techniques to select programs that closely match users' interests, and then schedule the selected programs for users' preferred viewing times. Since the network bandwidths necessary for supporting simultaneous, independent program selections in such personalized services are prohibitively large, PSAs employ intelligent caching strategies that judiciously store media information at strategic locations (such as neighborhood servers) in the network, thereby trading storage for network retransmissions. We have devised a simple, yet near-optimal greedy caching strategy, using which PSAs can determine *when*, *where*, and *for how long* media information must be cached in a metropolitan-area network, so as to optimize the overall service costs borne by users. Together with the media storage and communication architectures that we have developed, these caching strategies will be instrumental in economically offering a variety of services in various spheres: from digital libraries and distance learning in education to video on-demand in entertainment.

Publications:

(a) Srinivas Ramanathan and P. Venkat Rangan
"System Architectures for Personalized Multimedia Services"
Invited paper, To appear in *IEEE Multimedia Magazine*, 1994

Multimedia Collaboration Management

GOAL: Video and audio differ fundamentally from text in their semantics and usage, and hence require a radically different approach to designing services and service access paradigms. The goal of this research is to develop paradigms, software architectures, and specification languages for supporting a wide range of collaborative multimedia applications, structured interaction support, and interactive multimedia environments.

- *Previous Work*: Earlier work in multimedia collaborations has mainly been in the area of simple person-to-person tele-conferencing.
- *Our Ideas*:

  1. We have proposed a comprehensive **taxonomy of multimedia collaborations** that includes static and dynamic conferences, nested conferencing (in which a conference as a whole can participate in another conference), transient and persistent conferences, homogeneous and heterogeneous conferences, and symmetric and asymmetric conferences.
     Publications:

     (a) P. Venkat Rangan
     "Video Conferencing, File Storage, and Management in Multimedia Computer Systems",
     Computer Networks and ISDN Systems, Vol. 25, Pages 901-919, March 1993

(b) P. Venkat Rangan and Harrick M. Vin
"Multimedia Conferencing as A Universal Paradigm for Collaboration"
In *Multimedia - Principles, Systems, and Applications*, Lars Kjelldahl (ed.), Chapter 14, Springer-Verlag, April 1991

(c) Srinivas Ramanathan, P. Venkat Rangan, and Harrick M. Vin
"Integrating Virtual Reality, Tele-Conferencing, and Entertainment into Multimedia Home Computers",
IEEE Transactions on Consumer Electronics, Vol. 38, No. 2, Pages 70-76, May 1992

2. We have developed a **model that can capture the requirements of diverse types of structured collaborations**. The model combines efficiency and power are combined via a hierarchy of three paradigms, at the lowest level of which are *streams* for media communication modulated by access rights of participants within collaborations. The higher two levels of paradigms are *sessions*, which represent collections of semantically related media streams, and *conferences*, which represent temporally related sequences of sessions. These paradigms support unification of both interactive and non-interactive collaborations, sophisticated access control, and intra-group and inter-group collaborations, thereby yielding a rich environment for carrying out multimedia collaborations.

   Publications:

   (a) Harrick M. Vin, P. Venkat Rangan, and M. S. Chen
   "System Support for Computer Mediated Multimedia Collaborations"
   In Proceedings of *ACM Conference on Computer Supported Cooperative Work (CSCW '92)*, Toronto, Canada, November 1-4, 1992

3. We have been involved in the Etherphone project at Xerox Palo Alto Research Center to **integrate video storage and transmission services** in the Etherphone environment. The multimedia extensions to the Etherphone environment permits a user to participate in multiple multimedia conferences simutaneously using multiple channels of audio and video, as well as access stored audio and video documents within a multimedia conference.

   Publications:

   (a) P. Venkat Rangan and D. C. Swinehart
   "Software Architecture for Integration of Video Services in the Etherphone Environment"
   IEEE Journal on Selected Areas in Communication, Vol 9, No. 9, Pages 1395-1404, December 1991

   (b) Harrick M. Vin, P. T. Zellweger, D. C. Swinehart, and P. Venkat Rangan
   "Multimedia Conferencing in the Etherphone Environment"
   IEEE Computer, Vol. 24, No. 10, Pages 69-79, October 1991

Media Mixing

GOAL: Once a collaboration is established, media communication can take place among its participants. The most demanding applications for multimedia communication are those in which video and audio streams are being continuously transmitted by various sources (such as users and storage services), and the individual video images and audio streams have to be continuously mixed together to obtain a single composite image and audio stream, and displayed synchronously (e.g., a tele-orchestra). The goal of this research is to develop algorithms, protocols, and communication architectures for supporting media mixing in such high-end collaborative multimedia applications.

- *Previous Work*: Most existing systems use either wholly centralized or wholly distributed communication architectures, both of which are not scalable. Furthermore, they assume globally synchronized clocks, and do not handle network delay jitters.

- *Our Ideas*:

  1. We have developed a **mixing algorithm** which minimizes the difference between the generation times of media units that are being mixed together in the absence of globally synchronized clocks and presence of transmission delay jitter. The algorithm is shown to be **complete**: given that there are no other message exchanges except media data between a mixer and the media sources, there cannot be any other algorithm that succeeds when our algorithm fails.
     Publications:
     (a) P. Venkat Rangan, Harrick M. Vin, and Srinivas Ramanathan
     "Communication Architectures and Algorithms for Media Mixing in Multimedia Conferences"
     IEEE/ACM Transactions on Networking, Vol. 1, No. 1, Pages 20-30, February, 1993.

  2. We have proposed **hierarchical mixing architectures** for supporting scalable conferences, and have developed **algorithms for designing optimal hierarchies** that minimize end-to-end delays. Two types of mixing are considered: **monostage**, in which each mixer combines media only after it receives media units from all its sources, and **multistage**, in which the mixer combines media units in a pairwise manner as soon as its receives them from at least two sources. Towards a non-hierarchical network such as the internet, we have developed a **Packet Train** protocol in which the routing nodes themselves of a network can be used as mixers, thereby integrating mixing with routing. The packet train protocol also supports **proportional delay mixing**, i.e. participants receive more recent media units from their immediate neighbors, and less recent media units from distant participants.
     Publications:
     (a) Harrick M. Vin, P. Venkat Rangan, and Srinivas Ramanathan
     "Hierarchical Conferencing Architectures for Inter-Group Multimedia Collaboration"
     In Proceedings of *ACM Conference on Organizational Computer Systems (COCS '91)*, SIGOIS Bulletin, Vol. 12, No. 2-3, Pages 43-55, November 1991
     (b) Srinivas Ramanathan, P. Venkat Rangan, and Harrick M. Vin
     "Designing Communication Architectures for Inter-Organizational Multimedia Collaboration"
     Journal of Organizational Computing, Vol. 2, Nos. 3 & 4, Pages 277-302, 1992

(c) Srinivas Ramanathan, P. Venkat Rangan, Harrick M. Vin, and Thomas Kaepppner
"Optimal Communication Architectures for Multimedia Conferencing in Distributed Systems"
In Proceedings of the *12th International Conference on Distributed Computer Systems (ICDCS '92)*, Yokohoma, Japan, Pages 46-53, June 1992

Efficient Video Communication over ATM Networks

GOAL: In order to support the real-time performance requirements of video applications, networks must offer guarantees of minimum bandwidth, and maximum end-to-end delay and media loss. Assuming bounded buffers, violation of either the bandwidth or the delay constraint induces buffer overruns at switches, resulting in video frame losses. The goal of this research is to develop efficient video communication strategies that can provide frame-level loss bounds to video applications in emerging fast packet-switched networks, such as ATM.

- *Previous Work*: In most existing proposals, loss bounds are both expressed and enforced in terms of network-level transport units, such as packets or cells, rather than application-level video frames. Merely using an application-specified frame loss bound itself as the packet-level loss bound may not be sufficient. For instance, if each video frame is composed of 100 packets, an application-specified 1% frame loss bound cannot be guaranteed if the network were to provide a 1% bound on the packet loss, because in the pathological case, one packet could be lost for every consecutive 100 packets transmitted by the application, thereby resulting in the loss of each and every video frame. Of course, a straightforward way in which the network can guarantee the application-requested 1% bound on the frame loss is by providing a 0.01% bound on the packet loss. However, this constitutes an overly stringent demand on the network, potentially leading to its under-utilization.

- *Our Ideas*:

  1. We have proposed a simple to implement, yet effective, strategy called **Frame-Induced Packet Discarding** (FIPD), in which, upon detection of loss of a threshold number (determined by an application's video encoding scheme) of packets belonging to a video frame, the network attempts to discard all the remaining packets of that frame. Using a discrete time markov chain model of the FIPD strategy, we have devised a method for computing the frame loss probabilities of video applications, which then serves as the criterion for admission control. By isolating frames that experience packet loss from those that do not, FIPD reduces the fraction of frames that are lost during transmission. In addition, selective packet discarding releases network bandwidth and buffer space, which can be used to admit and service a greater number of video channels. Performance simulations of FIPD in a target ATM environment corroborate its predicted practical utility: FIPD offers close to two-fold increase in number of video channels for both JPEG and MPEG.
  Publications:

     (a) Srinivas Ramanathan, P. Venkat Rangan, and Harrick M. Vin
     "Frame-Induced Packet Discarding: An Efficient QoS Management Strategy for Video

Networking"
In Proceedings of Fourth International Workshop on Network and Operating Systems Support for Digital Audio and Video, Lancaster, UK, November 1993.

Prototype Demonstrations

- *Hardware Environment*: The multimedia lab is equipped with four multimedia stations, each consisting of a Sun SPARC workstation, a Compaq PC-AT equipped with full-color motion video analog/digital conversion and compression hardware (made by UVC Corp.), a video camera and a microphone. These multimedia stations, together with a NeXT Dimension, are interconnected by both Ethernet and FDDI networks, and are served by a 486 Compaq System-Pro file server. This entire lab cluster was connected to other metropolitan-area sites within San Diego via 45 Mb/s T3 link at the time of demonstrations.

- *Demonstrations at Educom, San Diego, CA, October 16-18, 1991*:

  1. Multimedia storage and on-demand real-time retrieval over a 45 Mb/s T3 link between UCSD Multimedia lab and San Diego Convention Center.
  2. Multimedia tele-presentation system for use as a remote classroom teaching mechanism.
  3. Media mixing in multimedia conferencing for use in tele-orchestras.

- *Demonstrations at ILP, San Diego, CA, April 26, 1992*:

  – Illustration of the dramatic improvement in media synchronization due to feedback technique in the playback of a two-channel video object on a heterogeneous network environment: one video channel retrieved over FDDI onto a 386 PC, and the other channel onto a 286 PC over Ethernet.

- *Publications*:

  1. P. Venkat Rangan, W. Burkhard, R. Bowdidge, Harrick M. Vin, J. W. Lindwall, K. Chan, I. A. Aaberg, L. M. Yamamoto, and I. G. Harris
     "A Testbed for Managing Digital Video and Audio Storage"
     In Proceedings of *Multimedia - for now and the future, USENIX Summer Conference*, Nashville, Tennessee, June 10-14, 1991, Pages 199-208
  2. P. Venkat Rangan, Harrick M. Vin, K. Chan, and I. Aaberg
     "A Window-based Editor for Digital Video and Audio"
     In Proceedings of the *25th International Conference on System Sciences, Hypermedia and Multimedia mini-track*, Kauai, Hawaii, January 7-10, 1992, Pages 640-648

Deliverables

The outcome to date has been:

- 16 Journal publications, 23 Conference publications, and Four Prototype demonstrations
- One Ph.D. thesis completed: Harrick Vin has just graduated from the UCSD multimedia lab to take up a faculty position at the University of Texas at Austin beginning fall '93
  - Harrick M. Vin
    "Architectures and Algorithms for Designing Digital Multimedia Storage Servers"
    Ph.D. Dissertation, Multimedia Laboratory, University of California, San Diego, July, 1993
- A second Ph.D. thesis is in the final stages of completion. In addition, Two M.S. theses, Two M.S. projects, and Two B.S. projects have been completed.

Funding

The total funding to date is $ 0.8 million, and has been provided by NSF, NSF Young Investigator Award, Powell Foundation, UC Micro, Starlight Networks, UVC Corp., Compaq, IBM, NeXT, Codenoll, Network Peripherals, NCR, and Xerox, via both fellowships and grants.