HEWLETT PACKARD

Distribution of Video to the Home: Cell or Circuit?

Riccardo Gusella, John O. Limb
Media Technology Laboratory
HPL-94-07
March, 1994

video on demand, cable-TV, interactive multimedia services, communications, fiber/cable hybrid networks, circuit switching, cell-switching, ATM, SONET

Efficient distribution of video to the home is central to many new services such as video-on-demand, home shopping, and educational services. Implementors are now debating whether to deploy Asynchronous Transfer Mode (ATM) distribution right to the set-top in the home, or whether a circuit approach (Synchronous Transfer Mode) might be more appropriate. In this paper, we explore both approaches by proposing an implementation for each and then comparing them in detail. The criteria we use for the comparison are cost, transmission efficiency, maintainability, scalability, flexibility, and interoperability. We assume that the distribution plant is typical of a fiber-rich, analog cable-TV installation.

We conclude that the circuit approach is somewhat simpler to implement and more bandwidth efficient. The two approaches are likely to have similar maintainability and scalability properties. For applications that require changing mixes of voice, video, and data the cell approach gives more flexibility. Since ATM connectivity to the wider world will be mandatory, both cell and circuit must provide it. For the cell the connection point would be at the set-top, while for the circuit it would be at a more central point.

1 Introduction

Communication to the home in the past has been rather limited. It has consisted of little more than telephone service, broadcast and cable TV, and in some cases data modems for home business, telecommuting, and, increasingly, connection to data information services. In the past, there have been many trials to determine whether new types of information service to the home can be offered profitably. Examples are movies-on-demand, home shopping and customized news, information and reservation services. The results invariably suggest that, except possibly for entertainment services, customers do not find the services attractive enough to want to pay an amount necessary to make the services profitable.

What is changing to warrant the enormous current interest in providing a broad range of new services to the home? Besides the regulatory and political, there are a number of changes that warrant it. One change is that cost of processing power is falling very dramatically. This enables customized information processing, storage and control. This, in turn, means that image and video can be used to enhance services and improve user interfaces. Thus, a video-on-demand service can provide users with exactly what they want, when they want it, and with the quality they want. A customized news or weather service can filter information for users, providing only the information they really seek.

Another change is that communications costs are falling dramatically resulting in very high speed broadband infrastructures being deployed in many countries. The ability to move large amounts of data very cheaply over high-speed backbones is fast becoming a reality using SONET [Bellcore89, Omidyar93] and ATM technology [DePrycker89]. Distribution within the local area, the so-called community network, is more problematic. Costs can be amortized over fewer users the closer the service gets to the home. Most "new" services provide a faster and more convenient way to get information and entertainment, not fundamentally new service. Users have to be enticed to adopt the new services by their convenience and low cost. Further complicating matters is that if penetration of the new service is not high, the cost per user will be.

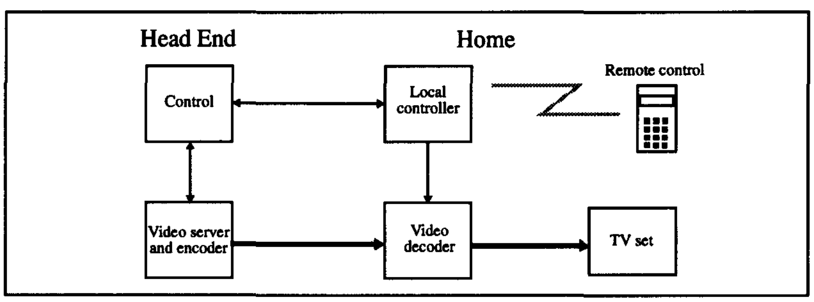

Figure 1. Video-on-demand system

As represented in Figure 1, Video-on-Demand (VOD) communication within the community network consists primarily of a downstream video channel from the head-end or central office (HE/CO) to the home and a bidirectional data and control channel between the HE/CO and home [Hodge93]. One could argue that the downstream data and control could be multiplexed with the downstream video channel. Since, however, video streams account for a very large portion of the downstream traffic, we believe that there is cost and performance advantage in optimizing this traffic.

It seems clear that a video-based information service must interoperate with the wider video world; much of the video material to be delivered to the home will come via very high-speed wide-area networks and, increasingly, these networks will be ATM based. At one extreme, all video information entering the HE/CO could be decoded into synchronous streams and then be delivered to the home. At the other extreme, the video information could be delivered to the home (and even distributed within the home) on totally standard, full-function ATM connections, in which the video streams are transported by 53-byte cells. (A compromise would be to adapt the ATM protocol to fit the more limited requirements of video distribution.) Our purpose is to present a comparative assessment of these two technologies for video information services in spite of the current euphoria associated with ATM.

Since it is impossible in a single paper to consider all types of servers and distribution networks, we have selected those designs that we feel are most likely to be deployed in the short and medium term [Brown93, Robbins93]. We assume that the distribution network consists of fiber trunks connected to a cable plant. Furthermore, we assume that the existing 6 MHz cable channelization will be retained as a cheap and efficient technique for both analog and digital distribution to the set-top.

Our yardstick of performance includes the cost of the communications subsystem, the efficiency with which the transmission medium is used, the ability to grow the system efficiently from a small number of users to perhaps tens of thousand (scalability), the ability to change and add new services (flexibility), and the ability to connect to standards such as Synchronous Transfer Mode (STM) and ATM. Issues of security, privacy, etc. are outside the scope of this comparison, but we suspect that the solutions would be analogous in the two environments.

2 Reference configuration

In this section, we describe a generic video-on-demand system at a level of abstraction such that it will constitute the basis for both a circuit and a cell implementation. In the following section, we then specialize our reference configuration adding those details that are necessary for comparing the two technologies.

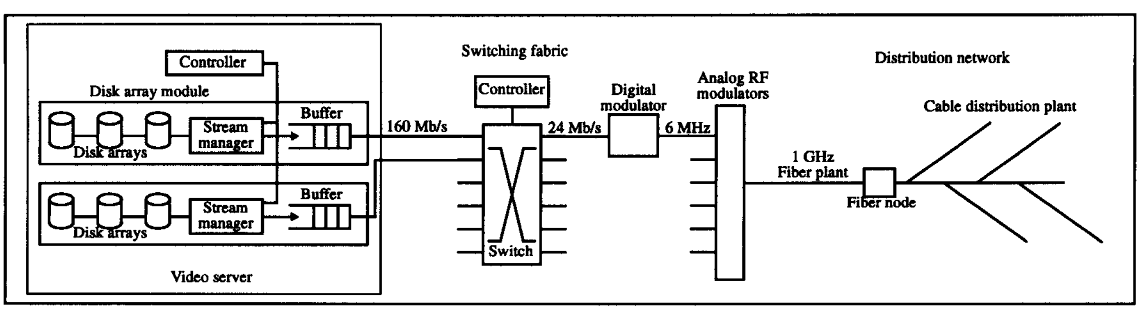

Figure 2. Architecture of a generic video-on-demand system

Figure 2 shows a plausible configuration of a VOD system consisting of a video server, a switching fabric, and a distribution network that delivers video to TV set-top devices in people's homes. A key component of the design, not shown in this figure and not addressed in the paper, is the signaling channel used by set top devices to establish, modify, and terminate service.

The video server consists of several *disk array modules*, each controlled by a *stream manager*, which is responsible for fetching the data from the disk array and streaming them toward the users. The video information is stored in the disk arrays in compressed format according to the MPEG-2 standard. We call each sequence of data directed to a user's set-top a *stream*. The server's *controller* terminates the signaling channel and routes calls for service to the appropriate module, which then offers the service independently of the controller. Disk blocks transferred on behalf of a stream manager are queued into a *buffer* organized as a series of FIFO queues, one for each of the streams being viewed by VOD viewers. The main purpose of the FIFO buffer is to remove the bursty patterns of disk data access.

A stream is moved through the *switching fabric* from the buffer to the *distribution network*, and then to the *set-top*. The switching fabric consists of a data switch that multiplexes several VOD data streams together and delivers the aggregate flow to *digital modulators* that, in turn, convert the digital data into standard 6-MHz TV channels. (In Section 3.2 we will also discuss a system that employs 12-MHz modulators.) The TV channels are then combined with standard TV channels coming typically from a satellite distribution system and fed to a fiber trunk and from there to the cable distribution plant to the homes. Since the design and properties of the switching fabric will depend greatly on which of the two approaches, cell or circuit, are used, we will deal with it in detail in the following section. The distribution network will invariably produce some jitter that will have to be compensated at the set-top.

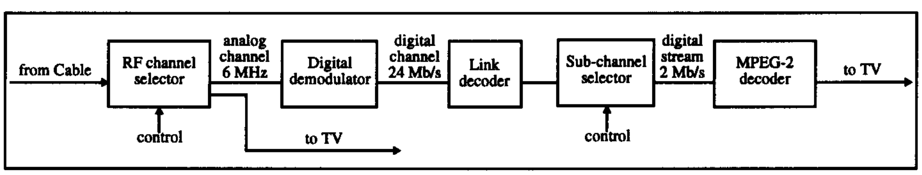

Figure 3. Architecture of a generic set-top device

The block diagram of a prototypical set-top device is shown in Figure 3. An *RF channel selector* tunes the device to one of the many analog channels available on the cable system. The information on a channel dedicated to VOD distribution is fed to a digital demodulator, which extracts from it a bit-stream consisting of several multiplexed VOD channels (*sub-channels*). The *link decoder* deals with transmission-layer frames, which we will address in Section 3. A *sub-channel selector* demultiplexes out of the aggregate stream the byte stream corresponding to the sub-channel of interest to the user and feeds it to the MPEG decoder, whose output, finally, drives the TV receiver. (We note that the MPEG-2 Adaptation Layer specifies a Packetized Elementary Stream that includes a packet cyclic redundancy check (CRC) for error detection [MPEG93]. Because we assume that the CRC is handled by MPEG, there is no need in our system for a transport protocol that detects transmission errors.)

In order to make our comparison of technologies more concrete, we make the following numerical assumptions derived from our analysis of today's technology. The recorded material is stored on the disks compressed at various rates ranging from 2 to 6 Mbit/s. Data are transferred from the disk arrays in blocks of 256 Kbytes. Such a block size corresponds to an interval of 1 second at the rate of 2 Mbit/s. A disk array module's aggregate transfer speed is 20 Mbyte/s. The speed of the input lines to the switching fabric matches that of a server's module; that of the output lines is four times as large, allowing for a total of 320 2 Mbit/s streams. The digital modulators can encode 24 Mb/s into a 6 MHz analog cable channel. The important numbers are shown in Figures 2 and 3.

3 System description

3.1 Circuit Solution

Let us now look at how a circuit connection may be established from the buffer stage of the disk array module to the TV set of the user. We first consider the reference system described in the previous section, and then address how the system may be scaled up.

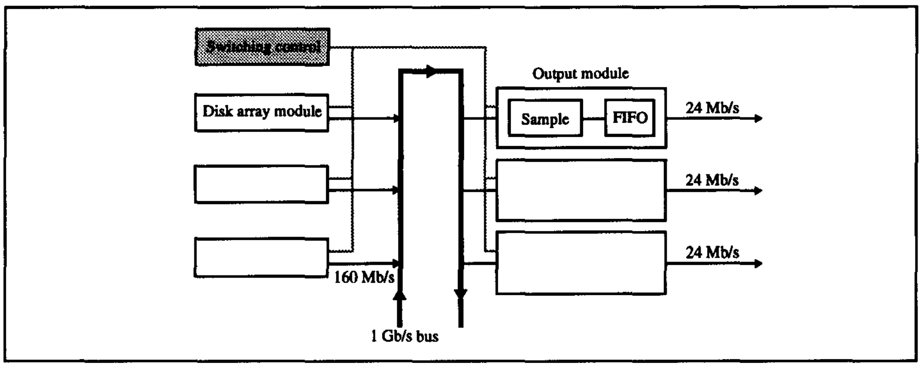

Figure 4. Bus-based time-division switch

The architecture of a flexible switch is shown in Figure 4. The output of each disk array module is gated onto the bus, multiplexed one word at a time. Streams are assembled into channels by gating them from the bus as required. A simple, single FIFO, of maximum size 12 words, is required to assemble the streams. [If the order of the sub-channels within the channel is important, then a time slot interchanger (TSI) is necessary, rather than a FIFO.]

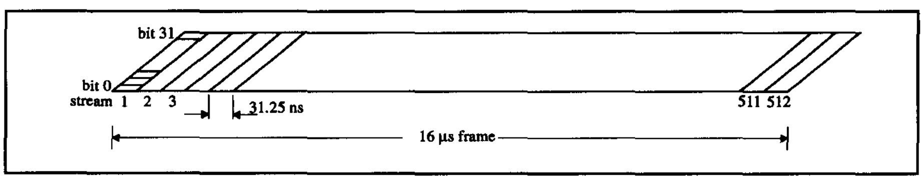

Figure 5. Frame structure for a 512-stream switch

The controlling dimension for such a switch is the bandwidth required by the maximum number of streams that it must switch. If we assume a bus speed of 32 MHz and a bus width of 32 bits, the total bus capacity is 1024 Mb/s. This means that the switch can support 512 simultaneous 2 Mb/s streams. This bus speed is achievable using today's inexpensive technology.

To effect a connection from a disk array module to a channel, the controller selects a slot in the switch and, firstly, instructs the disk module to write into that slot and, secondly, the output module of the required channel to read from that slot. Note that the output module FIFO works at the speed of the channel, a modest 24 Mb/s in this instance.

This architecture is non-blocking, it is capable of multicasting on the output and selecting on the input. The control and implementation are simple since all switching is totally synchronous, and determining a route through the switch requires little more than consulting and maintaining tables of the connections to and from the bus as inferred above.

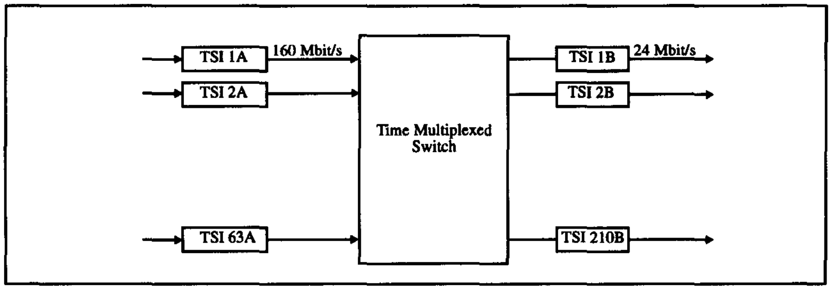

Figure 6. High-capacity circuit switch

This specific architecture is limited by the speed of the bus. The speed in the above example is 1024 Mb/s. Bus-based switches are marketed today with speeds of 2.4 Gb/s (e.g. Fore Systems' ASX Switch). For large switches, perhaps 5000 2 Mb/s streams, a multistage switch would be necessary. A possible structure is a three stage switch consisting of a virtual time slot interchanger followed by a Time Multiplexed Switch (TMS) and finally, a second time slot interchanger. An example for a 5000 stream switch is shown in Figure 6, where we again assume that each server module delivers a 160 Mb/s stream with the output ordered for input to the TMS. The TMS, in turn, may be realized in multiple stages. Such switching architectures have been deployed since the introduction of digital switching and their technology is well understood [Hui90, McDonald83]. However, the functionality and generality of existing switches are not well matched to the specific needs of video distribution; an implementation specific to video distribution would lead to a lower cost solution. Circuit switches could be scaled to very large sizes using well known technology and well known control and routing techniques.

Figure 5 shows the frame structure for a 2 Mb/s, 512-stream switch, using a 16 μs frame. With minimal change to the control, a bit stream can be assembled for any multiple of 2 M/bs, up to 24 Mb/s (limited by the digital capacity of a channel). For example, one could package two high quality 6 Mb/s channels together with two medium quality 4 Mb/s channels and two low quality 2 Mb/s channels.

If 2 Mb/s is too coarse a granularity, the frame structure could be lengthened. A frame structure of 64 μs would permit the channel to be allocated in 0.5 Mb/s chunks. This would require little change to the overall system.

One or two sub-channels could be allocated for random traffic for control, file transfer, downloading, etc. These sub-channels would be kept separate from the sub-channels allocated for synchronous traffic, but they could be placed arbitrarily in any channel. They are perhaps best regarded separately from the video streams as shown in Figure 1.

3.1.1 Circuit set-top device

Since there are a greater number of set-tops than channels it is even more important to have an economical "set-top". There will not be just one type of in-the-home box. Some customers will want just their existing analog services, others will want movies on demand, others will want to play sophisticated, interactive games.

Figure 3 shows how a digital video stream would be recovered. The digital demodulator produces a continuous digital stream at 24 Mb/s. In the example considered here, we have 12 2Mb/s sub-channels, and the framer (link decoder) detects the start of the frame. Since each slot of the sub-channel contains 32 bits, a frame is 12 x 32 = 384 bits in length. Any simple framing structure that preserved a clear 384 bit frame would be appropriate.

The acquisition time for frame synchronization could be quite low, a few milliseconds. An eight-bit framing word every four frames might be adequate, resulting in a 0.5 percent overhead. If faster acquisition were required, a longer frame word with 1 to 2 percent overhead might be necessary. Selecting the required stream is then a matter of counting down to the selected sub channel and gating out the 32-bit word. The resulting stream can be fed directly to the video decoder. Other streams on the same channel can also be selected with no more than an additional counter and gate. These additional channels could transport additional video streams or separate data streams.

3.2 Cell Solution

In the ATM solution to video data delivery, we transmit MPEG data from the video server to the set-top device using the ATM layer. Because we assume that the MPEG Adaptation Layer will allow the receiver to detect transmission errors (see Section 2), we need not employ an ATM Adaptation Layer protocol (such as, for instance, AAL5 [Lyon93]) that would duplicate that function. The disk-array module will then segment the video data into 48-byte ATM-cell payloads, and add the ATM cell header. This processing is likely to be done in hardware, by one of the many ATM segmentation chips under development (for an available chip-set see [National93]).

At the heart of the cell implementation is an ATM switch. For possible switch architectures, we refer the reader to [DePrycker89] and [Tobagi90]. As far as the physical transmission layer is concerned, among the various choices, SONET appears a natural one for the distribution network because of SONET's network management features. We note that here we have two physical layers: SONET, and the encoding supported by the modulator. However, while the modulator provides bit timing and line coding, SONET is used as a transmission convergence sublayer, according to B-ISDN nomenclature, to multiplex and delineate ATM cells. Note also that the output of the baseband digital modulator is further encoded by an analog RF modulator for transmission over the cable network.

More debatable is instead the physical layer to be used between the disk subsystem and the switch. Many contend that SONET should be considered, even over links that will span a few meters at most, because of the anticipated high availability, high volume, and consequent low cost, of integrated circuits implementing it. We believe that a custom physical layer would be more advantageous in order to reduce the transmission overhead and the system costs. However, to make the following discussion more concrete in various numerical examples, we will assume that the system employs SONET.

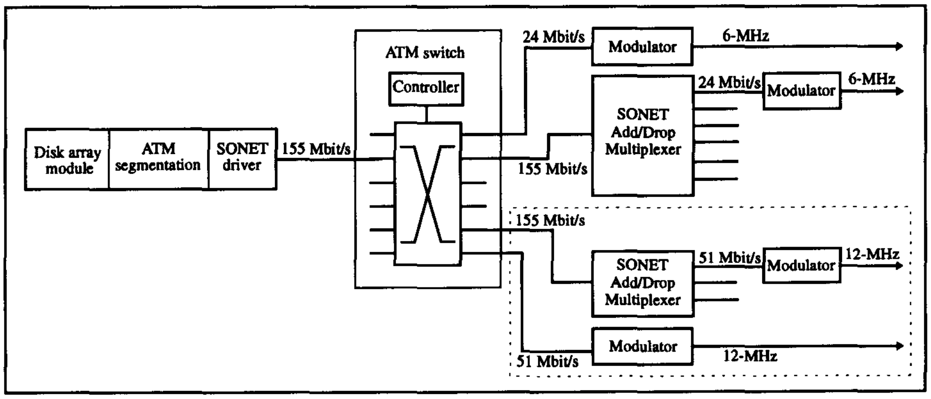

Figure 7. The ATM solution

In Figure 7 we show a diagram of the ATM/SONET solution. A disk module is served by an STS-3 interface at 155.52 Mb/s. Because of SONET overhead and ATM cell-mapping overhead, the bit rate available to ATM cells is 149.76 Mb/s. From this, we subtract the 5-byte ATM-cell-header overhead resulting in a maximum effective MPEG-layer bit rate of 135.6 Mb/s. Thus, out of the theoretical 80 2-Mbit/s channels that the disk array could support, an STS-3 SONET is able to transmit at most 67 of them.

For the distribution network we show two cases that differ for the bandwidth of the digital modulators used. In the case of standard 6-MHz channels, the modulators encode one half of the bit rate of STS-1, or 25.92 Mb/s. (Bellcore has recently put forward a proposal for scaling the STS-1 51.84 Mb/s convergence sublayer rate to lower bit rates [Robe93].) We note that modulators could be built to encode at exactly that bit rate. In this Section, we assume that whenever we use the value of 24 Mb/s, we indeed refer to one half of the STS-1 rate. We show a switch interface line that produces a 24-Mbit/s SONET stream that is fed directly to the modulators. In addition, we also show a standard STS-1 switch line that drives a SONET add/drop multiplexer (ADM) that extracts six 24 Mb/s substreams.

In the second case (the circuit is shown within a dotted box), we have a distribution network with modulators that occupy 12 MHz of the spectrum available on the cable system. This arrangement allows us to connect an STS-1 ATM-switch port to a single modulator, or to connect a 155 Mb/s port to a SONET ADM that drops three STS-1 lines, which are followed by the modulators. The key advantage in this solution with 12-MHz modulators is that we are able to employ standard ATM/SONET components. Note also that we could use DS3 lines (DS3 is one of the ATM physical layers adopted by the ATM Forum) in place of STS-1 lines, with just a small reduction in the effective bit rate.

3.2.1 ATM Set-top device

The ATM/SONET set-top device will include a potentially expensive SONET decoder. However, we show that ATM processing can be performed entirely in software. In the calculations below, we assume that no special tricks are used and, therefore, come up with an upper bound on the instruction counts. The calculations are for the case in which 6-MHz modulator encode 24 Mb/s streams.

The set-top digital demodulator (Figure 3) produces a SONET bit stream, from which the framer—in this case a SONET decoder—extracts a sequence of ATM cells. The ATM header, a 5-byte field, includes a virtual-circuit identifier (VCI) and a 1-byte ATM header checksum. (We notice that since the SONET ATM cell delineation algorithm involves computing the ATM header checksums, depending on the SONET decoder implementation, we could assume that received cells do not have header errors.) Header processing involves computing the header checksum, comparing it with the value in the cell, and demultiplexing on the VCI in order to decide whether the cell's payload should be queued into the MPEG input FIFO, or should instead be ignored for not being addressed to the set-top.

We assume that ATM cells are queued on a short FIFO and that the set-top CPU gets interrupted when the FIFO is half full. We estimate that ATM header processing (including header checksum) can be performed using 20 32-bit instructions on a standard RISC processor; a total of less than 1.2 MIPS for the 25.92 Mb/s SONET stream. In addition to header processing, the data belonging to the desired 2 Mb/s stream should be copied from the ATM FIFO to the MPEG-2 input FIFO. Assuming three instructions for every four bytes copied, the copy operation requires 0.75 MIPS to be performed. Thus, a 2-Mbit/s stream requires no more than 2 MIPS—a fraction of the capacity of the CPUs likely to be utilized in future set-top devices. Should an ATM reassembly chip be developed for set-top devices, it would be a low-complexity, low-speed part.

4 Comparison

4.1 Cost

Of the components illustrated in Figure 2, the transmission subsystem is the same in the circuit and the cell solutions. Therefore, we compare the costs of the two implementations by looking at the switching fabric, at the set-top device, and at the disk array module, of which we will consider different configurations. The cost-per-port-per-user will be higher in ATM not only because of the inherent higher implementation costs, but also because the lower transmission efficiency will reduce the number of users who will share a particular port.

We also note that the cost-per-port of the switching fabric relative to the cost of the disk array module will decrease as the number of disks in the disk array module increases. (The number of disks in a disk module is related to the number and size of programs that can be stored on that module, where copies of the same movie count as separate programs, and is independent of the number of module streams, which depend only on the bit rate of the module port and the speed of the streams that are requested.) We conclude that for disk modules with a smaller number of disks, the switching fabric will contribute a larger proportion of the total per-user cost. Therefore, a server with a smaller number of titles will benefit more from using a circuit-switching fabric.

The cost difference in the set-top is the difference in cost between a TDM physical layer and the SONET physical layer. Since this difference could be significant, and the set-top is the most cost-sensitive component in the entire VOD system, the cost of the ATM/SONET solution should be weighed against the general advantages of ATM. In the system we are considering, SONET becomes a cost liability. However, ATM does not necessarily require SONET to the set-top, and we could look at simpler and cheaper ATM framing protocols.

4.2 Transmission Efficiency

The transmission efficiency of the ATM solution is much lower than that of the circuit solution. Aside from various rates mismatches that arise from using standard components, SONET introduces a 3.3 percent physical-layer overhead and ATM an 11 percent header overhead, as discussed in Section 3. In contrast, the circuit solution would require an overhead of perhaps 0.5 percent.

4.3 Flexibility

By system flexibility we mean the ability to handle integrated data and voice traffic with video streams of different rates. We have, a priori, assumed that there will be a separate bidirectional data and control path between HE/CO and the home. While bursty traffic could be combined with video streams the cross-section of each channel is not large so that combining the two traffic types in one stream creates a difficult scheduling problem and probably would not save significant capacity. Combining bursty traffic from many users in a channel of larger cross-section might yield greater efficiency while preserving flexibility. However, a more complex scheduling algorithm would be required, policing of each virtual circuit would be necessary, and the set-top would have to operate on a higher-speed stream.

Consider the scenario where video is coded at different rates, say 1.2 Mb/s, 1.55 Mb/s, 2.0 Mb/s, and 2.15 Mb/s. With a cell system these rates could be allocated to a stream in a simple manner until it is full—there is virtually no wastage. A circuit approach adds capacity with a given granularity. By using a long frame the granularity can be made fine at a negligible increase in complexity. Nevertheless, there will be some wastage perhaps ranging from a few percent to 10 percent.

ATM/SONET provides flexibility in the way video may be packed into a stream. In other ways, however, because of standards, and the generality of the system it is not flexible. To illustrate this consider the following situation. In our cell example we have a good match between one half STS-1 rate and the 24Mb/s rate of the digital channel. If a more efficient digital modulation were deployed that operated at 36 Mb/s over a 6 MHz channel, there would be no simple way to use the capacity increase using standard ATM/SONET elements. With the circuit approach the multiplex ratios could be modified. It would require that the switch be initially designed to permit the counters that set the multiplex ratios to be under program control. On the other hand note that, ATM, without SONET, could use a nonstandard physical-layer protocol that was more efficient and flexible.

4.4 Maintainability

SONET and ATM are full featured systems with maintenance functions imbedded in the systems. They are designed to operate reliably in large common carrier systems. Cost-reduced systems could presumably be designed to match the particular needs of a cable system. Similarly, circuit solutions could deploy techniques used by existing common carrier systems, or be customized to the needs of a cable system.

4.5 Interoperability

Interoperability must be provided between the public network (an ATM network) and a cell or circuit solution. The difference lies in where the conversion to a digital stream is made. Both cell and circuit approaches require that an incoming ATM stream from the network be reassembled. In the case of the cell approach the data may be packed in 48 byte lumps and reassembled at the set-top, whereas in the circuit approach because the data is handled as a continuous stream it would have to be reassemembled at the HE/CO.

Cell systems could exist side-by-side with circuit systems if they were interconnected by ATM/SONET, since SONET is a synchronous protocol.

4.6 Scalability

Ideally one would like a switch design that scales smoothly with an increase in capacity of the system. Small ATM switches may be built cheaply using a common bus but at some point the bus becomes a bottleneck and some other architecture must be used: a multistage switch of some space-time combination or a self routing switch. The same is true of a circuit switch except that self routing would not appear to be an option. There would seem little advantage here for either approach.

5 Discussion and Conclusions

The principal issues that distinguish the two approaches are cost, transmission efficiency, and flexibility. Cost is probably the primary factor. It appears in the cost of the switch and associated equipment, and in the set-top. Standard ATM switches today are very expensive and would constitute a not insignificant fraction of the cost of a VOD system, particularly so for a small server. The cost of ATM switches is expected to fall steeply over the next few years, but we should also expect the cost of storage, the primary contributor to the cost of the server, to continue its downward march.

Much of the cost of a switch lies in the line cards, not in the switching fabric itself; this is the value of the circuit solution: its "line cards" are trivial synchronous circuits.

The cost of the set-top will be strongly influenced by whether SONET is deployed across the community network. SONET might be desired by an operator for the monitoring and management capability it provides, but it needs more consideration to determine how well the SONET's structure is matched to the needs of the cable operator.

The transmission efficiency of the ATM/SONET solution is appreciably lower than the circuit approach. This, however, will be of lesser importance as more and more efficient modulators are introduced.

Our comparison of the flexibility of the two systems revealed no ATM advantage. Bursty traffic can be combined with video streams on cell or circuit to achieve greater efficiency but at the price of increased scheduling and management problems. While both cell and circuit solutions can handle arbitrary rate video streams, the circuit approach will waste between a few percent and 10 percent of transmission capacity to do so. Because SONET only allows a few transmission rates at the low end, its use in a cable network results in less flexibility. However, a new physical layer for ATM over cable channels could be defined.

We have compared a circuit-switching solution and an ATM/SONET solution to a very specific problem: the distribution of video to the home using the existing cable network. Under our assumptions, and based on the technical issues we have addressed, no clear winner emerges although it appears that for small VOD systems or for VOD systems with a limited selection of titles, the circuit solution appears better.

In this paper we have discussed the downstream video channel. We are now challenged by the technological alternatives for the control channel and hope to explore them in the future.

References

[Bellcore89] Bellcore, Synchronous Optical Network (SONET) Transport Systems Generic Requirements, Technical Reference TR-TSY-0004400, September 1989.

[Brown93] R. Brown, Interconnects: MSOs Learn to Work Together, CED Magazine, December 1993, pp. 30-36.

[DePrycker89] M. De Prycker, Asynchronous Transfer Mode—Solution for Broadband ISDN, Ellis Horwood Series in Computer Communications and Networking, 1989.

[Hodge93] W. Hodge, S. Mabon, and J.T. Powers Jr., Video on Demand: Architecture, Systems, and Applications, SMPTE Journal, September 1993, pp. 791-803.

[Hui90] J.Y. Hui, Switching and Traffic Theory for Integrated Broadband Networks, Kluver Academic Publishers, Boston, 1990.

[Lyon93] T. Lyon, AAL2: An Idea Whose Time Has Passed, ATM Forum Technical Committee, SA&A Subworking Group, Contribution 93-1011, November 16-19, 1993.

[McDonald83] J.C. McDonald, ed., Fundamentals of Digital Switching, Plenum Press, New York, 1983.

[MPEG93] MPEG Systems Working Committee, MPEG-2 Systems Working Draft, ISO/IEC JTC1/SC29/WG11, November 4, 1993.

[National93] National Semiconductor, FRED Chipset Technical Manual, Version 2.0, 1993.

[Omidyar93] C.G. Omidyar and A. Aldridge, Introduction to SDH/SONET, IEEE Communications Magazine, Guest Editorial, Vol. 31, No. 9, September 1993, pp. 30-33.

[Robbins93] J.O. Robbins, The Cable Industry in the Post Divestiture Environment, IEEE Communications Magazine, Vol. 31, No. 12, December 1993, pp. 44-46.

[Robe93] T.J. Robe, W.E. Stephens, Proposal for Scaling the STS-1 51.84 Mb/s TC to Lower Bit Rates, ATM Forum Physical Layer Technical Committee, Contribution 93/841, September 8, 1993.

[Tobagi90] F.A. Tobagi, Fast Packet Switch Architectures for Broadband Integrated Services Digital Networks, Proceedings of the IEEE, Vol. 78, No. 1, January 1990, pp. 133-166.