

Case Study: Resilient Backbone Design for IPTV Services

Meeyoung Cha*, Gagan Choudhury†, Jennifer Yates†, Aman Shaikh†, and Sue Moon*
*KAIST †AT&T Labs – Research

mycha@an.kaist.ac.kr, gchoudhury@att.com, jyates@research.att.com, ashaikh@research.att.com, sbmoon@cs.kaist.ac.kr

ABSTRACT

This paper presents a case study of designing resilient backbone network for supporting IPTV services within a single network provider. We first introduce the architecture and the characteristics of IPTV traffic. We then explore the design space in terms of IP versus optical technologies, hub-and-spoke versus meshed service layer topology, dual-homed versus ring access, use of multicast, routing, and fast failure restoration. From this design space, we propose a number of design instances and evaluate them according to the capital expense they incur. We demonstrate significant benefits of multicast in reducing capital expense for broadcast TV, illustrate that our particular switched optical network design requires less capital than the IP based designs, and show that ring access to the backbone is more attractive for broadcast TV, while dual-homed access being more attractive if we have high volume of realtime Video on Demand (VoD) traffic.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design; C.4 [Performance of Systems]: [Design Studies, Reliability, availability, and serviceability]

General Terms

Algorithms, Design, Management, Reliability

Keywords

Internet Protocol TV (IPTV), Network Provider, Network Design, IP Network, Optical Network, Capital Expense

1. INTRODUCTION AND MOTIVATION

The telecommunications market is rapidly evolving to offer commercial-grade live broadcast TV and Video on Demand (VoD) over IP – known as Internet Protocol TV (IPTV) [9]. World-wide, the number of IPTV subscribers was reported to surpass 4 million in 2005. Scaling this to mass markets in today's highly competitive environment necessitates an extremely reliable and cost effective network infrastructure – all the way from the central head ends where the video is sourced, to the customers.

Nation-wide delivery of broadcast TV to video head ends has traditionally been achieved via a satellite-based infrastructure. However, satellite capacity exhaust and competi-

*These authors are supported by KOSEF through AITrc.

tive pressures are providing network providers with the opportunity to carry video on their terrestrial infrastructures [9]. Today, these terrestrial networks typically equate to IP and optical backbones.

IPTV introduces a number of interesting challenges for the research community. One fundamental question is: **what is the best network architecture for supporting IPTV?** In this study, we limit ourselves to the problem of backbone network design, and do not focus on distribution out towards the customers. Even though IP forms the basis of the architecture for most network providers, this does not necessarily translate into using an IP backbone for distributing traffic over the wide area. An alternative is to use an optical infrastructure to distribute the traffic. Beyond technology decisions, network designers must also incorporate topology, routing, and performance considerations in designing a reliable and cost effective distribution network.

In this paper, we provide what we believe to be the first detailed comparison of different architectural alternatives for supporting IPTV within a single network provider. We focus on the capital expenditure associated with these designs. Our work is not a complete exposition on IPTV service designs; nevertheless, we aim at grasping the essential high-level design principles and trade-offs by rigorously exploring the design space.

One purpose of this paper is to address the design space to support IPTV services in the backbone distribution network. These design choices are discussed in light of technologies, hierarchy, routing, and failure restoration. Another purpose is to present how the design space can be realized in an operational backbone network. Especially, we focus on a few key architectural alternatives, namely: 1) integrating IPTV services with an existing IP based network; 2) constructing a dedicated overlay on top of an existing IP based network; 3) constructing a new point-to-point interconnected flat IP network; and 4) integrating with an existing switched optical network. Accordingly, we propose a number of design instances and evaluate them according to the capital expense they incur. We also show the efficacy of using multicast over unicast for broadcast TV.

From a performance point of view, there is a tight interplay between the application and the network design. The application can be designed so as to best utilize the available network; similarly, the network must have adequate performance to support the application. If the application is highly sensitive to loss (*e.g.*, little or no buffering is employed), then the network must have minimal loss and support ultra-fast failure recovery. Discussing detailed requirements of IPTV services is beyond the scope of this paper, and we refer to the details of the ongoing work in [9]. How-

ever, we consider the performance requirements to the extent that they impose specific fast failure recovery and availability constraints that impact our network designs.

The rest of the paper is organized as follows. In §2, we introduce the overview of IPTV services. In §3, we present a case study of design space exploration and present a number of pragmatic design alternatives. Assessment of the designs follow in §4, and we conclude in §5.

2. OVERVIEW OF IPTV SERVICES

We start out by describing the service level architecture and traffic characteristics of IPTV services.

2.1 Service Architecture

The IPTV service architecture we envision includes a backbone and multiple regions [5]. We assume that there are two locations from which IPTV traffic is sourced – these are known as the *Super Hub Offices*, or *SHOs*. Two SHOs provide redundancy to ensure reliable video transmission, even in the face of catastrophic failure of one of the SHOs. In our study, we assume that both SHOs are always live to ensure rapid recovery in the event of a SHO failure.

Video streams transmitted from the SHOs are received at the *Video Hub Offices*, or *VHOs*, where the video streams may be further processed (*e.g.*, advertisement insertion) and then transmitted out towards the customers. VHOs also store local video content to support VoD. Each VHO consists of video equipment and two routers (for redundancy), where these routers interconnect the backbone with the regional network. For simplicity in this paper, we collectively refer to the SHOs and VHOs as *service nodes* and refer to the routers within a service node as *service routers*. Figure 1 illustrates the service architecture.

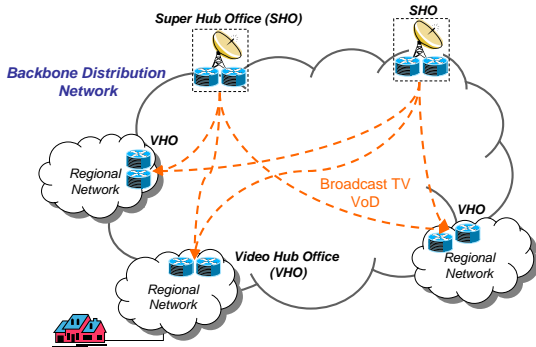


Figure 1: Service architecture of IPTV

This paper focuses on the backbone part of the service architecture, *i.e.*, the network that connects SHOs and VHOs. The design of the regional networks for connecting VHOs to end hosts (or customers) is outside the scope of this paper.

2.2 Characteristics of IPTV Traffic

We consider three types of IPTV traffic being carried over the backbone network: broadcast TV (realtime), VoD downloads (non-realtime), and realtime VoD (realtime). Broadcast TV traffic consists of traditional TV, High Definition TV (HDTV), and CD quality music channels. To be competitive with today’s cable and satellite TV offerings, IPTV must offer hundreds of TV channels, which are typically 2 to 6 Mb/s or 6 to 12 Mb/s, depending on whether they are standard or high definition TV signals. This translates into

multiple Gb/s of aggregated capacity toward each VHO. Although this in itself is not a very large value, if we multiply this by tens of VHOs, it becomes a massive amount of traffic. Since identical content is to be distributed across all VHOs, we consider the use of multicast capabilities to minimize the communication cost (*e.g.*, bandwidth consumed) [3].

VoD content is sent to each individual user as a real-time dedicated stream. However, to minimize the amount of traffic that must be carried across the backbone, we assume that popular VoD content is stored at the VHOs, and sourced from the VHO as and when requested by the end users. When new content becomes available to the SHOs, it is pushed from the SHOs to the VHOs during off-peak periods. These transfers do not require realtime delivery, and bulk-transfer applications (*e.g.*, ftp) can be used to ensure reliable delivery. Since the traffic is assumed to be carried during low utilization periods and does not require realtime transmission, it has minimal impact on the network design and architecture. We thus do not take non-realtime VoD traffic into account while designing the backbone network.

If the service provider decides to offer a vast variety of VoD content, it may not be cost effective to store the entire content at every VHO. From an economic stand-point, it generally makes more sense to store popular VoD content at the VHOs ahead of time, but source more esoteric content from the SHOs in realtime, as and when customers request it. Sophisticated cache management algorithms may be used to increase the hit ratio of the requested VoD content. However, some percentage of the VoD content will have to be sourced from the SHOs to serve realtime customer requests. We assume unicast delivery for those VoD content, and take the bandwidth requirement of such traffic into account in our backbone design analysis. It is worth pointing out that the realtime VoD downloads are expected to have greater traffic variability during peak usage periods compared with the broadcast TV traffic.

3. DESIGNING THE BACKBONE NETWORK

We focus on designing the backbone network interconnecting the service nodes. We start our discussion with the main axes of the design space and various options available along each axis. A set of design instances are realized by combining options available along each of these axes.

The first axis deals with the technology used for interconnecting the service nodes. We consider two alternatives: layer 1 (optical) and layer 3 (IP/MPLS) technologies. By optical network, we refer to an infrastructure consisting of optical components that provide fixed bandwidth “pipes” (or links) interconnecting the service routers. On the other hand, by IP network, we assume that traffic between the service nodes is routed over intermediate IP routers.

The second axis is the service layer topology. We consider two options: hub-and-spoke and meshed topologies. In a hub-and-spoke topology, the SHOs are directly connected to every VHO. In this topology, the service nodes only source/sink traffic. In a more highly meshed (or ring based topologies), service nodes carry “through traffic” in addition to acting as sources/sinks of traffic. Thus, traffic destined to each VHO may pass through multiple intermediate VHOs before reaching its destination.

Finally, we consider the option of using multicast [3] capabilities to reduce capacity required for broadcast TV. As mentioned earlier, identical broadcast TV content is streamed

from SHOs to all VHOs. This indicates that multicast can be used to reduce the overall bandwidth consumed by such traffic. In this work, we also investigate how much value multicast capabilities provide us, compared to using unicast.

Having discussed the main axes for designing the backbone, we next present how to combine various options to come up with design instances. Since the technology has the biggest impact on what other options we can choose and the resulting cost, we divide our discussions into IP and optical based backbone designs.

3.1 IP Based Network Designs

We start by considering a network provider carrying IPTV traffic over an IP infrastructure. At one extreme, we consider a dedicated IP network constructed purely to carry IPTV traffic. This has the advantage in that the design can be customized to support IPTV services, and that IPTV traffic is not mixed with traditional Internet traffic, thereby isolating this sensitive traffic from the perils of the public Internet. The other extreme is to use a single common network to carry all traffic – including Internet, VPN, and IPTV services. Such an approach simplifies network management, in terms of having only a single network to design and operate, but requires careful performance management to ensure that the IPTV traffic receives high priority forwarding and is isolated from the vagaries of the Internet. Intermediate solutions also exist, for example, overlaying a dedicated topology on top of an existing infrastructure. This can be achieved by using the common backbone routers, but with separate links to carry IPTV traffic, providing a level of isolation for the highly sensitive IPTV traffic.

3.1.1 Integration with an Existing IP Backbone

We assume a scenario in which a network provider already has an existing IP network over which they will incorporate the new IPTV demands. As mentioned before, only the backbone links are shared and access links are dedicated for IPTV traffic. Utilizing an existing infrastructure enables rapid deployment of the new services, with minimal overhead and efficient utilization of the network resources. IPTV traffic is special in that it is mostly uni-directional (from SHOs to VHOs) and requires high bandwidth. Moreover, in the case of VoD, we expect high traffic variability during the peak usage hours. Utilizing a common infrastructure offers the potential to share bandwidth between applications. For example, Saturday night may be a high load period for IPTV, as couples sit down to watch their favorite movies, but is typically not a high utilization period as far as business applications are concerned. Figure 2 illustrates the integrated IP based network architecture. Service routers are connected to backbone routers via dedicated access links (*e.g.*, via a high bandwidth SONET circuit).

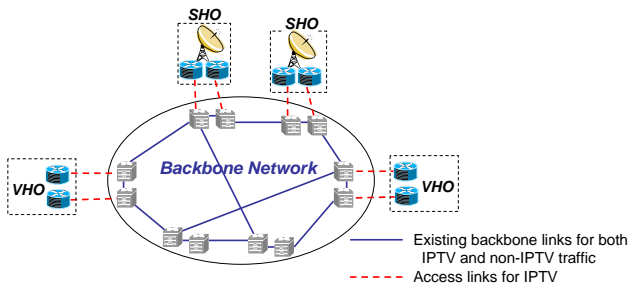


Figure 2: IP-based integrated design

One of the challenges of using an integrated network to support both Internet and highly sensitive IPTV traffic is that the IPTV traffic is at risk of being impacted by the Internet traffic (*e.g.*, congestion, DOS attacks). Measures should be taken to minimize such negative impact, for example: provide adequate priority forwarding of the IPTV traffic (using Diffserv [7], *e.g.*); implement measures to isolate IPTV traffic from DOS attacks; and provide IPTV traffic with preferential fast failure recovery (faster than that of best effort traffic) where feasible.

We assume that routing for unicast traffic is based on traditional Interior Gateway Protocol (IGP) algorithms such as OSPF [10] or IS-IS [2]. These algorithms route traffic over the shortest path between source and destination, based on the administrative weights assigned to each network link. For multicast traffic, we assume that there is a single multicast tree including the two SHOs and all the VHOs. The SHOs are assigned a common anycast source address so that each VHO is routed to the nearest SHO as it joins the multicast tree. Thus, traffic is transmitted from the nearest SHO to each VHO, and each VHO receives a single copy of the multicast traffic at any given point in time. Should a SHO fail, then the VHO is automatically re-routed to the surviving SHO. The multicast routing is realized using the PIM-SSM (protocol independent multicast - source specific multicast) protocol [1, 6], which uses a shortest path multicast distribution tree (MDT) rooted at the source. Receivers join the tree along the reverse shortest path based on the sum of the administrative link weights.

Rapid failure recovery is important in minimizing the negative impact of network outages longer than 50 to 100 milliseconds. Note that if we only use traditional IGP routing protocols and PIM-SSM, failure recovery may take up to a few seconds. While such delay may be acceptable with many applications, IPTV application requires well below 1 second failure recovery in order for the end users to not feel the impact at all. The IPTV application servers can also assist by incorporating buffers, so that the application can ride out short-term outages. Buffering necessitates delaying the signal transmission, which in many broadcast applications has no effect on the customer's satisfaction. However, for an avid sportsperson, for example, delaying a live broadcast is simply not acceptable – especially if one tries to synchronize across multiple broadcast mediums (*e.g.*, radio and TV).

It is thus critical that the backbone network be able to rapidly recover from network outages. We consider two enhancements in providing fast recovery against link failures, the most frequent type of failures in the network: 1) optical layer recovery (*e.g.*, SONET protection switching [8]), in which optical layer failures such as fiber cuts are recovered rapidly at the optical layer with sub-50ms protection; and 2) fast re-route (FRR), which sets up an alternate route that avoids the failed link [12]. In FRR, as soon as the link failure is detected, all the traffic on the failed link is diverted over a pre-established backup tunnel. At a later time as the IGP and PIM-SSM re-converge, traffic switches back from the backup tunnel to a new optimal path dictated by the routing protocols.

Finally, we consider two types of *access connections* between the service nodes and backbone network nodes: dual-homed and ring. Each node in the network, be it a backbone or a service node, is comprised of two routers to protect against a single node failure. In a dual-homed case, each

service node is connected to a backbone node such that the two service routers and backbone routers connect in parallel for fault tolerance. In a ring case, the service nodes in close geographic proximity to one another are connected as a ring, where the ring is also connected to the two routers in one of the backbone nodes. Figure 3 illustrates the two scenarios. Note that there exists a link connecting the two service routers of a VHO in the ring access. These two access mechanisms translate to providing a dual-homed and mesh-based service layer topologies, respectively.




Figure 3: Two types of access connections

So far, we have explored the design space in light of technology, hierarchy, routing, access connection, and failure restoration. This results in four designs which are summarized in Table 1, where each design is provided a unique name and the technology used. *Link-cap.* field shows whether the capacity of backbone links are shared with other Internet applications or dedicated for IPTV traffic, and *access* field shows the type of access connection. *Fast-failover* field shows what type of protection is used at backbone against link failures.

Table 1: Designs integrating with IP backbone

Design	Layer	Link-cap.	Access	Fast-failover
Int-IP-HS	IP	shared	dual-homed	protected links
Int-IP-HS-FRR	IP	shared	dual-homed	fast re-route
Int-IP-Ring	IP	shared	ring	protected links
Int-IP-Ring-FRR	IP	shared	ring	fast re-route

3.1.2 Dedicated Overlay on Top of an Existing IP Based Network

As an alternative IP layer design, we consider a network provider overlaying a dedicated topology on top of an existing infrastructure. This is implemented by using common backbone routers but dedicated links that support only IPTV traffic. The main advantage of this design lies in performance management since links are dedicated for IPTV traffic. Figure 4 illustrates the idea, where solid thin and thick lines represent the backbone links for the existing (non-IPTV) traffic and IPTV traffic, respectively. Service nodes are connected to backbone nodes via dedicated access links.




Figure 4: IP-based Design with dedicated overlay

The backbone links used for carrying the IPTV traffic

can be optically protected or not – independent of whether optical layer protection is used on the existing IP infrastructure. Access links may be dual-homed or interconnected in a ring topology. This results in four designs summarized in Table 2.

Table 2: Designs with IP-based dedicated overlay

Design	Layer	Link-cap.	Access	Fast-failover
Ded-IP-HS	IP	dedicated	dual-homed	protected links
Ded-IP-HS-FRR	IP	dedicated	dual-homed	fast re-route
Ded-IP-Ring	IP	dedicated	ring	protected links
Ded-IP-Ring-FRR	IP	dedicated	ring	fast re-route

3.1.3 Flat IP Network (No Backbone)

We consider constructing a new dedicated IP network purely to carry IPTV traffic. Here, we do not use an IP backbone network, but rather connect the service routers directly using point-to-point links. High-bandwidth links between two service routers can be established directly over dense wavelength division multiplexors (DWDMs) [11]. Therefore, we term this design as “flat IP network” or “flat IP over DWDM network.” Note that in our case, all the IP designs are built on top of layer 1 optical network. The distinction between IP and optical designs lies in which layer multicasting and routing are employed.

In designing such a network, we need to identify how the service nodes are interconnected. We consider here a meshed topology, where the service nodes carry “through traffic” in addition to acting as sources/sinks of traffic. We use a simple heuristic to design the topology: we divide the service nodes into a small number of communities of interest based on geographical proximity. Within a single geographic region, we connect the service nodes in a ring topology so that no single failure can disconnect the network (*i.e.*, traffic can be re-routed around the ring in case of a failure). More specifically, the two IP routers within each service node are connected, and also the IP router in one service node is connected to other IP router in different service nodes (forming intra- and inter-regional ring connections). Thereby, end-to-end paths from SHO to VHO will go through a series of intermediate service routers.

DWDMs do not traditionally provide multicast capabilities, and certainly not in the equipment deployed in large telecommunications networks. We thus assume that multicast capabilities are not available at the optical layer for this design instance. However, as the service nodes are carrying through traffic, we utilize multicast routing in the service routers themselves. That is, routing is determined by PIM-SSM for multicast and IGP for unicast at service routers.

Using this design, IPTV traffic is completely isolated from other Internet based traffic that the existing network may carry. Thus, additional performance measures are not required (although, we may still give high priority to IPTV traffic among others). Also, service nodes themselves are directly interconnected, and thus have no access structure. Providing fast failure recovery mechanism is still important in this design, and can be done by providing protection at the optical layer (which translates to using protected links between the service routers) or by relying on FRR in the service routers. This results in two design instances as summarized in Table 3.

Table 3: Direct inter-connected optical design

Design	Layer	Link-cap.	Access	Fast-failover
P2P-DWDM	IP	dedicated	none	protected links
P2P-DWDM-FRR	IP	dedicated	none	fast re-route

3.2 Optical Based Network Designs

The conceptually simplest way to connect service routers is to provide direct links between them. In contrast with the previous approaches, the paths from SHOs to VHOs are directly connected by logical links (without traversing through intermediate IP routers). The link interconnecting is established over an optical infrastructure, consisting of optical components such as dense wavelength division multiplexors (DWDMs) and optical cross-connects (OXCs) [11]. In the following, we consider one optical network alternative.

3.2.1 Integration with an Existing Switched Optical Network

We consider a network provider using its switched optical network to interconnect the service routers. We assume the optical backbone consists of OXCs. The term here is used somewhat liberally, to cover both all-optical and electronic switching. In either case, we assume that the OXC can switch any input wavelength/fiber to any output wavelength/fiber. If the OXC is electronic, then it can also support time-division multiplexing [8]. Connections between the service routers are established across one or more OXCs.

Optical networks defined above have traditionally been designed to provide point-to-point connections, such as those required to interconnect two IP routers. However, multicasting capabilities are now being considered for optical networks [13]. Optical networks can multicast their signals by replicating an input signal from a given input port onto two or more output ports at an OXC.

To effectively make use of the multicast capabilities offered by the switched optical network, we assume a hub-and-spoke logical topology for the service node interconnection. Specifically, we assume that each SHO establishes a multicast tree including all the VHOs, simultaneously multicasting the same signal to all VHOs. Each VHO is thereby simultaneously connected to both SHOs. The OXC to which each VHO is connected is then used to select the higher quality signal from the two SHOs.

Using the above approach, failure recovery is achieved by rapidly switching between the two signals received by the OXC directly connected to each VHO. We assume here that the optical network does not re-route internally around a failure. Thus, the paths from the two SHOs to a common VHO must be physically-diverse from one another to ensure that traffic can be received even when there is a major fiber cut. This introduces some additional complexity into the MDT calculation. In this paper, we use an Integer Programming (IP) based approach to create the trees, ensuring the necessary diversity requirements are met.

Table 4: Designs in optical network

Design	Layer	Link-Cap.	Access	Fast-failover
Opt-Switched	optical	time-divisioned	dual-homed	disjoint paths

4. EVALUATION OF DESIGNS

We analyze an IPTV architecture consisting of two SHOs and 40 VHO locations across the US. For fair comparison, we overlay the service layer on a common backbone network topology for both IP and optical network designs. This topology is representative of an operational backbone network consisting of approximately 100 nodes and 200 links. However, for the flat IP design with no backbone, as described before, we use a simple heuristic to complete the design – the outcome being a sparsely connected network de-

signed specifically for carrying only the IPTV traffic. This contrasts with the other network designs which are in no way optimized for the IPTV traffic as they also support a wide range of other services.

We omit detailed discussion on how each design is realized due to lack of space and simply note that all design alternatives are realized efficiently to carry the total traffic in the network (both multicast and unicast IPTV traffic, and in the case of the integrated IP design, other non-IPTV traffic).¹ In all of the designs, we consistently assume a high level of redundancy, by incorporating two SHOs and two routers within each service node. We also ensure that there is sufficient capacity between the service nodes so that the design can survive any single failure, including any link, line card, or router (or optical switch) failure.

We compare the capital expenditure across designs. We separately refer to the cost due to the backbone side as the *backbone cost* and denote the cost of the network-facing service routers as the *access cost*. The cost of a network design is represented as the cost of all of the transport cost relative to the distance of links and the router/optical switch port cost at the two ends of the links. Parameters for transport and port cost vary across different technologies (*e.g.*, port/link vendor/type/capacity). In our evaluation, we assume that all the links are OC48 (2.5 Gb/s) and select realistic parameters accordingly.




Figure 5: Cost comparison of multicast and unicast

We commence by evaluating the effectiveness of multicast in reducing the network capacity required. Figure 5 illustrates the relative cost of carrying broadcast TV traffic using multicast versus unicast technologies on both IP and switched optical networks². The Y-axis represents the capital expenditure, scaled to hide the proprietary nature of the data. The degree to which multicast benefits depends on the number of multicast endpoints (the 40 VHOs in this case) and the backbone topology. In general, multicast technologies will provide more significant gains as the number of VHOs increases. In the realistic topologies considered here, the graph demonstrates that multicasting reduces the backbone network cost by more than a factor of three compared with unicasting. The reductions in network cost appear comparable for both the optical and IP based designs.

Next, we compare the capital expenditure across all of the designs. Figure 6 shows the relative cost across designs for two different multicast and unicast loads. The cost is normalized such that the cost of Opt-Switched for carry-

¹By “efficiently carrying traffic”, we mean minimizing end-to-end latency and the total cost of designs under various failure scenarios.

²For the IP network, both multicast and unicast routing paths are determined according to Int-IP-HS-FRR design. For the optical network, multicast and unicast paths are determined in a fashion described in Opt-Switched design. Costing of all cases incorporates the extra capacity required to service the traffic under each single failure scenarios. Using other types of IP designs yield a similar result.




Figure 6: Cost comparison across design instances

ing a multicast load of 1 Gb/s to each VHO is 1.0. Please note the different y-axes in the two plots. In this figure, we make the following observations. First, for all of the loads analyzed (including loads not published here), the optical network appeared more economical than the IP network. This is primarily due to the lower port cost of an optical switch than an IP router – the optical port cost being approximately one third of the cost of an equivalent rate router port. In the optical designs, we have assumed an electronic switching fabric that can support time-division multiplexing (TDM) and virtual concatenation (VCAT) [4], so that we can allocate bandwidth within the backbone in multiples of STS-1 (54 Mb/s) granularity. This allows most effective use of network capacity, but it is important to note that not all technologies deployed today can support this.

Second, in IP designs, we observe that the most economical approach to achieving fast failover is fast re-route (FRR), as opposed to using optically protected links. In all of our IP designs, we assume that additional capacity is introduced in the IP layer to recover from router/port failures, independent of any other protection mechanism used. This same spare restoration capacity can be used by FRR to provide rapid failure recovery – thus, FRR only increases the required IP layer restoration capacity by a relatively small amount compared with that incorporated to handle the single router/port failures. In contrast, the additional capacity introduced at the optical layer to provide rapid optical layer protection cannot be shared with the IP layer restoration required for handling router failures, and thus effectively doubles the restoration capacity requirements in the network.

We next compare the access cost of the different solutions. The total network cost is typically dominated by access cost, with the exception of the direct DWDM interconnection of VHOs, where access cost is a meaningless concept. In the other designs, we observe that ring access is more economical than dual-homed access when carrying only broadcast TV (assuming multicast traffic). Connecting neighboring nodes in a ring topology allows the ring bandwidth to be shared by all nodes in the ring – making effective use of the multicasting capabilities. For N VHOs using ring access, we need to send only a single copy of the (multicast) broadcast TV signals around the ring. In contrast, if each of the N VHOs are dual-homed to the backbone then we need to carry a load proportional to N on the VHO access links. However, when unicast traffic is also incorporated, dual-homed access becomes significantly more cost effective. For large unicast demands, intermediate nodes in a ring are forced to carry

large amounts of “through” traffic destined for downstream VHOs, increasing the port and link costs.

Finally, in terms of the backbone cost, the integrated IP designs are consistently more economical than the dedicated designs. This illustrates the economic benefit of sharing a common infrastructure. However, such benefits must be weighed against the overheads of ensuring isolation of video traffic from the perils of the public Internet.

5. CONCLUSIONS

In this paper, we considered a particular problem of supporting IPTV services on a large backbone distribution network. Our work is based on a practical real-world setting and provided what we believe to be the first detailed examination for a single network provider. We focused on one of the critical issues in designing a network to support IPTV, namely, capital expenditure. Future work should consider performance trade-offs. Amongst our findings, we demonstrated significant benefits of multicast in reducing capital expense for broadcast TV, illustrated that our particular switched optical network design requires less capital than the IP based designs, and showed that ring access to the backbone is more attractive for broadcast TV, while dual-homed access being more attractive if we have high volume of VoD traffic that needs to be unicast in realtime.

6. REFERENCES

- [1] S. Bhattacharyya. An Overview of Source-Specific Multicast (SSM). RFC 3569, IETF, 2003.
- [2] R. W. Callon. Use of OSI IS-IS for Routing in TCP/IP and Dual Environments. RFC 1195, IETF, 1990.
- [3] S. Casner and S. Deering. First IETF Internet AudioCast. *ACM SIGCOMM CCR*, 1992.
- [4] L. Choy. Virtual Concatenation Tutorial: Enhancing SONET/SDH Networks for Data Transport. *J. Optical Networking*, 2001.
- [5] Cisco. Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband Solution Design and Implementation Guide.
- [6] Estrin et al. Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification. RFC 2362, IETF, 1998.
- [7] K. N. et al. Definition of the differentiated services field in the IPv4 and IPv6 headers. RFC 2474, IETF, 1998.
- [8] W. Goralski. *SONET: A Guide to Synchronous Optical Network*. McGraw-Hill, 1997.
- [9] Kim et al. Requirements for Internet Media Guides on Internet Protocol Television Services. Draft, IETF, 2005.
- [10] J. Moy. OSPF Version 2. RFC 2328, IETF, 1998.
- [11] C. S. R. Murthy and M. Gurusamy. *WDM Optical Networks*. Prentice Hall, 2001.
- [12] P. Pan, G. Swallow, and A. Atlas. Fast Reroute Extensions to RSVP-TE for LSP Tunnels. RFC 4090, IETF, 2005.
- [13] Y. Xin and G. N. Rouskas. Multicast Routing Under Optical Layer Constraints. In *IEEE INFOCOM*, 2004.