Proceedings of IDPT'99
4th World Conference on Integ. Design & Process Technology
July 1-5, 1999, Kusadasi, Turkey

ISSUES AND TECHNIQUES IN NETWORK-BASED DISTRIBUTED HEALTHCARE: ADVANCED NETWORK TECHNOLOGY

Stan McClellan* & Gary Grimes
Center for Telecommunications Education & Research
Electrical & Computer Engineering
Univ. Alabama at Birmingham
Birmingham, Alabama, 35294
Email: smcclell@uab.edu

Ken Burst
Broadband Laboratory
BellSouth Technology Assessment Center
Bellsouth Telecommunications
Hoover, Alabama 35244
Email: kburst@btac.com

ABSTRACT

With the advent of the "next generation Internet" and various related technologies, we observe important tradeoffs between emerging network capabilities and design related application requirements for network-based distributed healthcare systems. One of the aspects of these emerging network capabilities is *packet-switched routing*; *Packet-switched routing* is a particular aspect of "next generation"network structure which will have a profound impact on the deployment of advanced, network-sensitive medical applications. Here, we explore some of the underlying aspects of *packet-switched routing* technologies along with inherent capabilities and constraints which should be considered when designing network based systems.

INTRODUCTION

A significant obstacle to effective deployment of advanced networked medical applications is the failure to adeqately consider network capabilities and design parameters in system development. Historically, the network component has been the primary limitation for optimization of network-based medical system architectures. Factors contributing to this limitation include sparse network service availability, lack of economic incentive for including network cosiderations, inadequate insights into the relationship between application requirements and network capabilities, and inadequate understanding of network performance issues.

With the advent of the "next generation Internet" and related enabling technologies, we observe important tradeoffs between network capabilities and design related application requirements for network based distributed healthcare systems. These tradeoffs are intimately related to the design of systems which rely on Internet-based telecommunications to transport data between practitioners, patients, and/or remote locations. *Packet routing* is a particular aspect of "next generation" network structure which will have a profound impact on the development and deployment of such applications. The underlying constraints, requirements, approaches, details, and evolutionary design which will determine modern routing capabilities are of significant interest in many arenas.

EVOLUTION vs. DESIGN: INTERNET ROUTING

As with all large projects, the design of fundamental or critical components faces tradeoffs. In the distributed Internet, *routing protocols* are one of the most fundamental and critical technologies. The fact that the Internet exists and is somewhat stable lends credence to the idea of large-scale technological evolution. The Internet, to the credit of its designers and developers, has been able to maintain backward

---

*Corresponding Author. This paper was prepared in affiliation with David A. Conner, David Green, Gregg Vaughn, Jay Goldman, and Murat Tanik of the UAB Center for Telecommunications Education & Research. The authors also wish to acknowledge the contribution of Harold E. Crane of ATI.

compatibility with "the old" while introducing "the new", and at the same time remain reasonably stable. This has been accomplished in spite of many problems which are directly attributable to its tremendous popularity and expansion. The question at hand is this: will the Internet continue this evolutionary process, or will new designs replace the commodity Internet in a sudden paradigm shift?

The reliable, flexible structure of the commodity Internet is based, in large part, on the dynamic distributed routing protocols which determine "next hop" packet forwarding. These protocols, like the Internet Protocol (IP) are fundamental elements of the connectionless, best-effort paradigm of hop-by-hop routed networks. Unfortunately, the network requirements mandated by non-commodity users, such as distributed medical applications, as well as the congestion due to mass-migration of commercial traffic onto the Internet is challenging the "evolutionary" nature of these networks in 3 general areas: (Dumortier, 1998)

Routing protocols and topological complexity

Current-generation IP routers use routing protocols (ie. data exchange mechanisms) to distribute topology and reachability information which is used to compute optimal paths for "next hop" forwarding. With continued explosive growth of the Internet, the computational complexity of this procedure will become unmanageable. An approach to managing complexity is to sacrifice topological accuracy in favor of predictable structure. Such condensed topology indicators, where reliance on hierarchical structure is used to minimize information loss, are an area of active research. This ambiguity does not bode well for large-scale implementation of Internet-based distributed medical systems.

Quality of Service (QoS)

The complexity of current Internet routing mechanisms is based on a single dimension of traffic service characteristics. This single dimension optimizes reachability by minimizing "hop count". The introduction of higher-order, application-specific requirements such as "latency" or "guaranteed bandwidth" into routing algorithms requires significant changes to route computations as well as to the protocols which distribute topology information. Additional telemetry data (ie. higher volume, faster updates) are needed for each routing node to maintain an accurate perspective of the topology (ie. more complex computation). Distributed medical applications are likely to rely **very heavily** on absolute guarantees of QoS. As a result, an understanding of the mechanisms responsible for data delivery is incumbent on designers and system integrators (McClellan & Burst, 1999).

Traffic volume and forwarding performance

As traffic volumes increase and link data rates adjust upward to compensate, forwarding performance (normally a "layer 3" consideration) keeps pace by implementing algorithms in hardware-driven, connection-oriented modes called "shortcut routing" (ie. a "layer 2" consideration). Current-generation routers examine and process the header of each packet to determine "next hop" forwarding. Even if new, efficient routing algorithms reduce the computations and memory accesses required for the well-defined basic forwarding operation, additional poorly-defined burdens such as policy-based filtering and QoS assurance metrics will push hardware/software architectures to their limits (Metz, 1998; Dumortier, 1998). Additionally, implications of inadequate or unreliable advanced network service will translate into serious considerations for medical system designers.

The changes in IP routing over the years is a large-scale example of the evolutionary development of a system with an incomplete a-priori specification. Furthermore, the large installed base of legacy technologies in the Internet has created an "incumbency factor" which must be considered. Taken together, these circumstances result in an absolute requirement for tradeoffs between "optimum design" and "evolution". In the following sections we discuss some of the technologies and approaches which are likely to form a basis of the "next generation" Internet.

SHORTCUT ROUTING

The terminology "shortcut routing" has been used to describe capabilities available with the integration of connection-oriented network links in a connectionless data path. With shortcut routing, network-layer (or, "layer 3") data packets can be forwarded directly at the link layer ("layer 2"). So, with appropriate configuration, data streams can be transmitted *directly* between nodes in the forwarding path which support connection-oriented links. In this fashion, intermediate hop-by-hop processing is bypassed and latency characteristics are partially optimized.

The issues associated with integration of connection-oriented and connection-less technologies, particularly with respect to the evolution of routing protocols in the Internet, are interesting from a "design" perspective. These issues are even more important when considered in the context of distributed, Internet-capable medical systems. The predominant technologies which are vying for inclusion in the eventual solution are both products of extensive design and testing. However, the perspective from which these designs were initiated is quite different. A contender for the "connectionless" paradigm is the ubiquitous suite of Internet Protocols (IP). The primary strength of IP is the flexibility with which modifications can be integrated and design responsibilities distributed among the user community. A contender for the “connection-oriented” paradigm is the relatively new technology of Asynchronous Transfer Mode (ATM) which is primarily a product of industrial coalition. Thus, the integration of IP and ATM technologies to achieve a capability for shortcut routing in the Internet has two distinct flavors which derive from their respective proponents: “layered” (or “independent”) and “integrated”.

A Layered Approach

The “layered” approach to integrating IP and ATM in the Internet maintains a distinction between each technology’s routing protocols. The predominant IP routing protocol is known as Open Shortest Path First (OSPF), whereas ATM uses the Private Network-Network Interface (PNNI). The independence of these protocols in layered solutions leads to a scattering of layer-2 and layer-3 connectivity information between separately maintained databases. The result is that hop-by-hop IP forwarding paths must be mapped to ATM using a “crossover” protocol known as the ATM Address Resolution Protocol (ATMARP).

OSPF OSPF is a link-state routing protocol developed for large, router-based networks. Its primary function is to exchange network topology information among all of the routers in the network. Large routing networks are broken into areas, and these areas are interconnected by area border routers and backbone networks to form a two level hierarchy. With OSPF, each router has a complete view of the topology for its area, but an incomplete view of the topology for other areas. OSPF is a very efficient routing protocol using less than 1 percent of link bandwidth and less than 2 percent of CPU capacity for routing calculations. Large networks of routers are able to converge on a topology in a matter of seconds because routing changes are flooded across the network and calculated in parallel. Since every router in the network has an overall view of the network topology, each router is able to calculate the shortest (best) path across the network. OSPF is aware of link bandwidth and uses this information to calculate this shortest path. However, each router is still required to read the header and make a forwarding decision for each IP packet which is forwarded.

Unfortunately, routers are not able to take advantage of the native QoS capabilities of ATM. ATM switches in the backbone may be used to pass OSPF link-state information across a network to other routers, but routers make all forwarding decisions. As a result, there are no shortcuts across the network for large or important transmissions between two hosts. (Moy, 1998)

PNNI PNNI is a protocol specification developed to provide route discovery, routing, and signaling between autonomous networks. An autonomous network in this context is defined to be an ATM switch or group of switches which are operated and administered by the same entity. Network topology is discovered by arranging the overall topology into a hierarchy. Switches are grouped together into peer groups, and these peer group networks are grouped together to form larger and larger peer groups within the hierarchy. This hierarchical arrangement is extremely scaleable. Each switch in the hierarchy contains a complete overview of the network topology, although this view is somewhat limited in detail. The view is obtained by each switch or network advertising the state of its links to group leader switches or networks which in turn pass this information to the higher levels of the hierarchy. This link state information includes traffic parameters and QoS parameters. Switches use this topology and the link state parameters to route calls over links that will support the level of service required. When a call setup is initiated, the switch uses its topological view to create a designated transit list (DTL), or map of how to traverse the network to the destination switch. This DTL is then included in the call setup request and forwarded from switch to switch along the designated path. The topology known by the originating switch contains detailed information for its own peer group, but less detail for other peer groups along the designated path. Therefore, it is necessary for each switch along the designated path to add more detailed information about its peer group topology as the call setup request is passed from switch to switch. Since the information in a switch’s topology is incomplete, the actual characteristics of the connection are not known until the connection is completed. When it is found that a connection cannot support the desired level of service, PNNI is capable of re-routing the call using a process called *crankback*. PNNI signaling is capable of either point-to-point, or point-to-multipoint call processing. The signaling is responsible for functions such as call setup, call release, adding parties, dropping parties, providing information about call status and crankback processing (Onvural & Cherukuri, 1997).

ATMARP ATMARP is used to translate a destination IP address to an ATM address so packets can be routed properly across the ATM network. ATMARP and its sister protocol inATMARP are used in what is known as Classical IP over ATM (CIOA). In CIOA, the ATM network is devided into logical IP subnets (LIS). Each LIS has an ARP server to translate IP addresses into ATM addresses and to register routers. The inATMARP protocol is used to identify and record the IP addresses of an ATM-connected router when it registers as part of a LIS. In this paradigm, routing functions are relegated to routers and the ATM network simply provides transport services to the routers via Virtual Circuits. CIOA has no capability by itself of using or interpreting ATM QoS (Minoli & Minoli, 1998; Laubach, 1994). This approach, as in all “bandaid” approaches to evolution, results in route computations which don't have access to globally-known topology information. Consequently, end-to-end routes are fragmented, and not subject to overall optimization for throughput, reachability, or cost. (Dumortier, 1998)

Integrated Approaches

The "integrated" approach to combined routing can be divided into 2 distinct classes: *Multi-Protocol Label Switching* (MPLS), where topology information is collected by "slave" layers and translated into a form compatible with a single "master" layer, and *Integrated PNNI* (I-PNNI), where existing complementary approaches are melded together.

The characteristics of the ATM Forum's technology-specific topology distribution protocol (PNNI) have been described previously. I-PNNI extends PNNI with a hop-by-hop mode. In this fashion, both the existing layer-3 and layer-2 routing protocols are discarded in favor of a completely re-designed, fully-integrated, multilayer routing protocol. With I-PNNI, all nodes have full access to topology information for both layers. As a result, routing calculations can be jointly optimized using fully disclosed data. Unfortunately, although the design decisions and assumptions for I-PNNI are logically and theoretically superior to MPLS in many instances, the I-PNNI approach has some significant flaws. For example, the assumption of ATM as the only supported layer-2 technology, while a convenient simplification to the problem, neglects the real-world constraints of existing (legacy) network infrastructures. Additionally, the delays associated with design, implementation, and introduction of an abrupt paradigm shift into a large, distributed system often results in accelerated risk exposure due the non-stationary nature of technology.

In either MPLS or I-PNNNI, the *integrated* approach is more complex and risky than the *layered* approach. However, optimization based on the strengths of various node types (IP-only, ATM-only, combined IP-ATM) as well as the available forwarding modes (connection-oriented, connection-less) can produce a highly scalable and flexible solution to next-generation routing requirements. Further, in the context of distributed medical systems, the global optimization of routes and latencies is a crucial consideration.

MPLS: A REALITY-CONSTRAINED TRADEOFF

As a candidate for next-generation IP routing technology, MPLS addresses the requirements for QoS and availability in backbone networks by simplifying the packet forwarding function (Hagard & Wolf, 1998). This simplification is accomplished by pushing complex analysis and processing of per-packet routing information to the ingress point, instead of repeating the process at intermediate or "interior" nodes (Metz, 1998). This function can be viewed as a compromise between "signalling" as in true connection-oriented networks and "packet forwarding" of the commodity Internet.

The core forwarding function can be a significant bottleneck in conventional IP networks because of complex, per-packet calculations[1] combined with the increasing link data rates due to massive utilization. Conventional approaches to efficient forwarding for high-speed links include hardware acceleration (Hagard & Wolf, 1998). Unconventional approaches involve unique route-lookup algorithms which simplify the "longest match" comparison between the variable-length destination network prefix and the entries in the forwarding table (Kanakia, 1998). Unfortunately, as the dimensionality of the data requirements increases to allow QoS, etc. and forwarding algorithms require high-order information, these inflexible or "hardened" approaches will be somewhat limiting.

The MPLS concept relies on the definition of "overlay" label-switched paths (LSP) through a best-effort network. These LSPs form "tracks" or "classes" into which ingress routers forward packets based on the "best match" between data requirements and path availability. The limitation of available paths can be viewed as a finite number of lanes on a highway. While bottlenecks may occur in the "fast lane", at least all traffic doesn't stall in a single lane. As a packet traverses a forwarding node, the short, fixed-length "track label" is used to indicate the exit port. In this fashion, IP traffic can be forwarded efficiently since subsequent routers can bypass costly per-packet calculations by relying on the label or tag which indicates the correct "track" (Hagard & Wolf, 1998; Metz, 1998). A result of this compromise between best-effort and application-specific QoS is a sort of quantization of the possible QoS mechanisms. As a result, some distortion must be introduced in the servicing of individual data streams.

The label-switching approach is particularly amenable to ATM-based forwarders since ATM virtual path and virtual circuit indices have a similar structure. However, since IP is the assumed encapsulation, the concept behind MPLS route optimization is to translate all topology information into a common IP-based reference. As such, all layer-2 (ATM) configuration and routing intelligence must be "mapped" into layer-3 (IP) terms. This mapping is not obvious or completely accurate. Also, the resulting approach obviates existing, standardized ATM Forum protocols for User-Network Interface (UNI) and Network-Network Interface (NNI). Instead, traditional IP routing protocols such as OSPF are used for route computation, and the results are "stamped" into layer-2 devices by new pseudo-signalling protocols such as

---

[1] For example, OSPF uses Dijkstra's algorithm on distributed topology and reachability information to produce weighted per-hop forwarding rules. (Moy, 1998)

the Label Distribution Protocol (LDP) (Dumortier, 1998; Metz, 1998). This approach is clearly not optimal from the perspective of individual data streams, because OSPF is not QoS-aware. If "native" ATM signalling is required, MPLS and UNI/NNI can co-exist independently in a mode known as "ships in the night".

A strength of MPLS is its independence from the link layer technology. Much like its IP roots, MPLS can function over heterogeneous network infrastructures. In fact, the MPLS approach to integrated next-generation IP routing is in many regards inferior to the somewhat technology-specific I-PNNI. However, the *backward compatibility* of MPLS which allows the continued use of *legacy technologies* in a complex, distributed system is a vital component in reality-constrained optimization. The cost of this "nondestructive" integration is a mismatch between requirements and capabilities.

CONCLUSION

The dramatic growth in usage of the Internet and the "incumbency factor" of legacy components demands that IP routing be integrated with ATM in a compromise configuration. This compromise will take advantage of ATM's capabilities while remaining backward compatible with IP routing in what some might consider a suboptimal design. However, the constraints of legacy infrastructure and topological complexity require performance compromises to achieve broad-scale functionality and continued growth. As a result, "next generation" Internet routing may be incompatible with certain requirements of distributed medical systems. It is incumbent on medical system designers to understand these compromises as well as the capabilities of the evolving Internet. Such considerations must be incorporated in system designs in multiple ways. In particular, special attention to the interaction between application-specific traffic requirements and the "real performance" of deployed technologies is imperative in mission critical medical applications.

ACKNOWLEDGMENT

This work was sponsored, in part, by grants from the National Institutes of Standards and Technology Advanced Technology Program (70NONB581011) administered by SCRA, and by Bellsouth Telecommunications. This material is also based upon work supported by the National Science Foundation under Grant No. 9729500.

REFERENCES

Dumortier, P. 1998. Toward a new IP over ATM routing paradigm. *IEEE Communications Magazine*, **36**(1), 82–86.

Hagard, G., & Wolf, M. 1998. Multiprotocol label switching in ATM networks. *Ericsson Review: The Telecommunications Technology Journal*, November, 1–12.

Kanakia, H. 1998 (December). Next-Generation Routing Technology for the IP/ATM Network. *In: The IP and ATM Conference*.

Laubach, M. 1994. *RFC 1577: Classical IP and ARP over ATM*. IETF Network Working Group.

McClellan, S., & Burst, K. 1999 (July). Issues and Techniques in Network-Based Distributed Healthcare: Quality of Service and Bandwidth. *In: Proc. 4th World Conference on Integ. Design & Process Technology*.

Metz, C. 1998. Ingredients for better routing? Read the label. *IEEE Internet Computing*, **2**(5), 10–15.

Minoli, D., & Minoli, E. 1998. *Delivering Voice over IP Networks*. Wiley.

Moy, J. T. 1998. *OSPF: Anatomy of an Internet Routing Protocol*. Addison-Wesley.

Onvural, R. O., & Cherukuri, R. 1997. *Signaling in ATM Networks*. Artech House.