Abstract

Driven by the twin forces of industry-wide deregulation and the explosive demand for Internet access and bandwidth-intensive multimedia services, broadband local access has emerged as one of the key issues in modern telecommunications. In this article we describe a broadband local access network consisting of small, densely spaced packet-switching nodes interconnected by focused free-space optical links in a multihop mesh arrangement. Each switch serves a client, which may be an office building (containing, for example, conventional PBXs and LANs), a picocellular base station, or both. It is the responsibility of our local access network to economically and reliably extend broadband local access service (perhaps OC-3 or OC-12 for building LANs and PBXs; perhaps several tens of megabits per second to base stations) from an infrastructure end office or fiber ring add/drop multiplexer without requiring the installation of new buried optical cabling. Computed is the capacity of the multihop mesh, defined to be the maximum number of virtual connections which can be delivered to the infrastructure access point such that, independent of the traffic distribution among clients, all quality of service guarantees are maintained.

A Broadband Wireless Access Network Based on Mesh-Connected Free-Space Optical Links

ANTHONY S. ACAMPORA, UNIVERSITY OF CALIFORNIA AT SAN DIEGO
SRIKANTH V. KRISHNAMURTHY, HRL LABORATORIES

Throughout the 1990s, the Internet has increasingly evolved into a major means of communicating, accessing information, and conducting business, and the range of new multimedia services accessible via the Internet is strongly driving the demand for residential and commercial broadband local access. At the same time, deregulation of the telecommunications industry has created new opportunities and competitive forces, which are further accelerating the need to deploy broadband local access networks. Finally, as broadband access takes root, it is overwhelmingly likely that existing voice-oriented wireless services must be expanded to enable high-quality multimedia services to mobile users [1-4].

One of the main challenges in providing broadband local access is the need to deploy adequate capacity to ensure quality of service (QoS) guarantees to end users. Fiber-oriented access systems can surely provide adequate capacity, but are economically disadvantaged due to high installation costs. Wireless local loop (WLL) has become the focus of increasing attention because of its low capital and maintenance costs and ease of deployment [5]. Broadband WLL (BWLL) retains the advantages of WLL while providing the high capacity required to support the delivery of broadband multimedia services to homes and offices, and can potentially connect pedestrian and vehicular mobile subscribers as well.

In this article we consider a new strategy for reliably and economically deploying a broadband access platform. The basic idea is illustrated in Fig. 1, which is described later in greater detail. At its core, our broadband access network consists of a "continuum" of geographically dispersed but densely spaced small packet-switching nodes. Each node serves a client, which may be a building containing private branch exchanges (PBXs) and LANs, a picocellular base station, or both, and would ideally be mounted on top of that building or base station. In addition to providing access to the occupants of the homes/offices within its client building, a node which serves a picocellular base station enables service to mobile users in its vicinity.

The packet-switching nodes are interconnected by a dense mesh of focused bidirectional free-space optical links, each fully capable of withstanding atmospheric and mechanical disturbances by virtue of its short physical length [3]. Traffic generated by the clients of these nodes is ultimately relayed by the multihop optical mesh to an access (entry/exit) port of the wired infrastructure (e.g., a fiber ring add/drop multiplexer or an end-office switch).

In the past, consideration of free-space optical links for high-data-rate point-to-point applications have generally produced disappointing results: these relatively long links are prone to frequent failure caused by fog attenuation, wind, and mechanical misalignment. However, with our new approach, wide-area access regions are readily, reliably, and economically served. Key to the reliable operation of a low-cost free-space optical mesh is adequate density of switching nodes: if the density is sufficiently high, the length of each optical link will be sufficiently small that fog attenuation (the dominant impairment) is negligible, and mechanical tolerances are loose. Accordingly, the link margins can then be considerably relaxed, and the optical link expenses associated with tight link margins (pointing accuracy, optical beamwidth, focusing, high-power lasers, sensitive photo-receivers) can be eliminated. By way of example, data from the national climactic data center suggests that in many cities, a link availability of 99.999 percent is achieved at a bit error rate of 10^{-12} and a data rate of 622 Mb/s (OC-12) if the link length is smaller than 200 m; these numbers assume reasonable (but not extreme) pointing accuracy and a 20 mW laser transmitter. Depending on climate, 99.999 percent availability may be achievable at signifi-

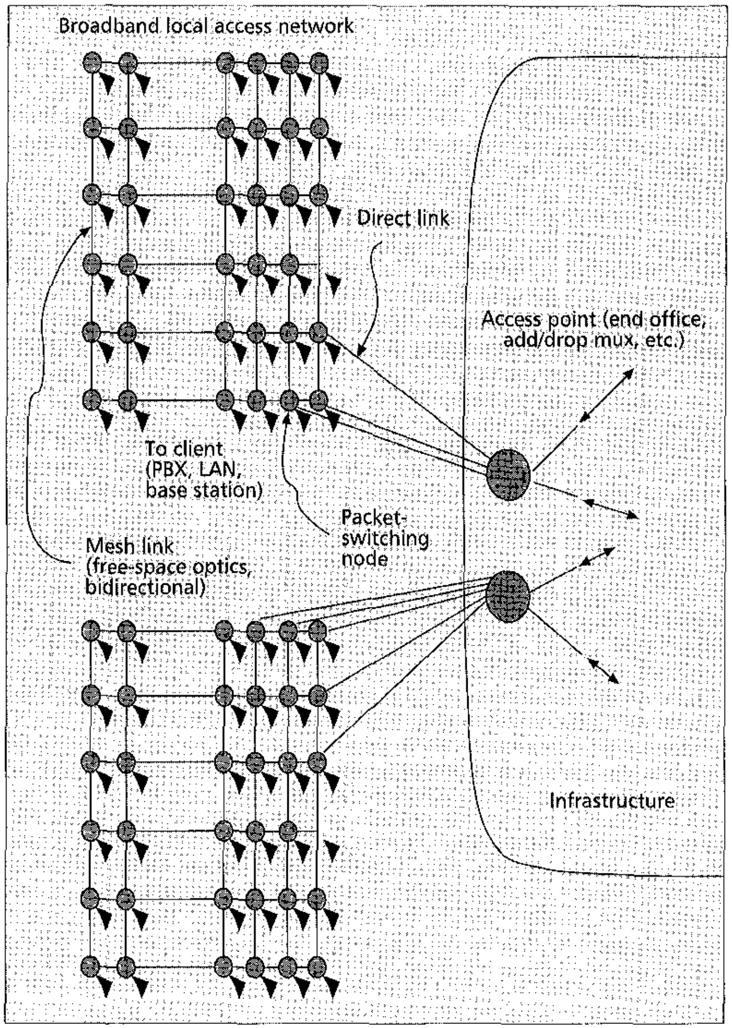

■ **Figure 1.** *The basic architecture for the broadband access network.*

cantly greater range. Alternatively, at 99.99 percent availability, many more cities can be served with link lengths of at least 200 m [6].

The topology of the mesh network may be represented by a rectangular grid, as shown in Fig. 1, where each crosspoint of the grid represents a packet-switching node. The mesh attaches to the wired network access point by means of some small number of direct links from those nodes geographically closest to the access point. Thus, each access point serves some relatively small region via the free-space optical links which span that region. Note that although the broadband access network takes the form of a rectangular grid, the clients served do not necessarily conform to this regular geometric pattern; the mesh shown in Fig. 1 merely illustrates how the packet-switching nodes are interconnected by free-space optical links, and logically this connectivity diagram does indeed take the form of a rectangular grid.

Thus, a relatively wide area can wirelessly be served from each network access point. The small packet-switching nodes are deployed with sufficient density that each optical link remains suitably short. Each packet switch is of maximum dimensionality 5 x 5 (four free-space optical crosslinks, one drop, not shown, to the wired building or picocellular base station).

In this article we find the capacity of such a broadband access network. We define the capacity as follows. Consider a single traffic class with some known traffic descriptor. Let the data rate in each direction for each optical link be such that, for this given traffic class, C virtual connections can be carried, each enjoying some guaranteed QoS. Similarly, let each direct link be capable of carrying up to C virtual connections. Suppose that node i's client originates and terminates T_i virtual connections where, for each node, $T_i \leq C$. Then the capacity of our broadband access network is defined to be the maximum number of QoS-guaranteed virtual connections that our broadband access network is always capable of delivering to an access point, independent of how the calls are distributed among clients (as long as no client generates/receives more than C calls, where C is the single-link capacity).

Note that we consider access network capacity only for a single traffic class, and only for symmetric traffic. Consideration of capacity regions for multiple traffic classes and asymmetric traffic are beyond the scope of this article, and are in fact subjects of ongoing research. Results are as follows. Referring to Fig. 1, let there be K direct links leading to a single access point. Then the capacity of the broadband access network is given by

$$C_{NET} = \begin{cases} KC, & 1 \leq K \leq 5 \\ 5C, & K > 5 \end{cases} \tag{1}$$

Proof

Consider a mesh of size R x R. Let $N = \lceil K/2 \rceil$, where $\lceil x \rceil$ is the smallest integer greater than or equal to x. Suppose that $R \geq N$; in such a case, it is always possible to draw $M = \lceil R/2 \rceil$ loops in the mesh (sometimes a loop may consist of just a single node). We label loops as shown in Fig. 2a and 2b, with the innermost loop bearing the highest index number; that is, loop $i + 1$ would always lie within loop i. Mesh links which belong to a loop are called *loop links*, and the mesh links interconnecting loops are called *interloop links* (Fig. 2b). For this case, we can redistribute the calls by first routing all calls generated by nodes in the inner $M - N$ loops to loop N using interloop links (note that these do not belong to any loop). Since the number of interloop links from loop $i + 1$ to loop i is equal to $4(R - 2i)$, it is possible to route all calls from the innermost $M - N$ loops onto loop N. Note that when R is even, the maximum value of i is $(R/2) - 1$, and for this value of i the number of links out of the innermost loop is eight. If R is odd, i is at most $(R - 1)/2$, and the number of links out of the innermost loop will be four. However, since there is exactly one node in the innermost loop when R is odd, it is possible to route calls from the innermost $M - N$ loops to loop N. As an example, in Fig. 2b it is possible to route up to $4C_L$ calls (the maximum number which can be generated in loop 3) from loop 3 to loop 2. Thus, at the conclusion of this process, all the traffic is concentrated on loops 1–N (the outermost loops). Now, if $K \leq 4$, the total traffic generated is less than or equal to $4C_L$. Since we have not yet flowed any traffic onto the interloop links interconnecting the outermost two loops, we can now use these links to uniformly distribute the traffic among these two loops. (Note that for this process, we cannot use any interloop link which will subsequently be used to attach an inner loop to the direct link serving that loop. For example, referring to Fig. 2 a, interloop links a and b cannot be used to redistribute traffic among the loops since they will subsequently be used to attach loop 3, via node X, to node Y, which is an *index node* or a node which has a direct link attached to it). Furthermore, having accomplished this, neither loop carries traffic in excess of $2C_L$ (note that loops are bidirectional).

If $K = 5$, the traffic intensity may be as great as $5C_L$, and in such a case, using an identical procedure, one may uniformly distribute the traffic among loops 1, 2, and 3 (the outermost three loops).

Since it is possible to carry up to $2C_L$ calls on each loop (C_L in the clockwise direction and the other C_L in the counterclockwise direction), we can successfully route the entire traffic to the *corner node* on the loop closest to the end office without encountering blockage on the mesh links. The interloop links connected to this node lead directly to an index node (Figs. 2a and 2b). Thus, using these links the traffic can now be routed to the end office.

Note that it may be impossible to distribute traffic uniformly among the outer three loops if the total load generated is in excess of $5C_L$, as illustrated in Figs. 2c and 2d. Let all the traffic be generated by a cluster of nodes at a corner of the mesh. If a maximum load of $5C_L$ is generated by a cluster of five nodes in a corner of the mesh, then, since there are five mesh links leading out of this cluster, it is always possible to redistribute the load (Fig. 2c) among the outer three loops as described earlier. However, if a maximum load of, say, $6C_L$ is generated by a cluster of six nodes in a corner of the mesh, as shown in Fig. 2d, it is impossible to route this traffic to the end office without blocking up to C_L calls since there are only five links (at least six links are necessary) leading out of the cluster. Thus, if the number of direct links, K, from the end office into a mesh is greater than five, given that the traffic within the mesh is less than KC_L, one cannot claim that there will be no blockage on the mesh links.

Now, let $R \leq N$. For this case, one may subdivide the mesh into *arcs*, as shown in Fig. 2d. Note that the inner N arcs (an arc may consist of a single node, i.e., the primary index node)

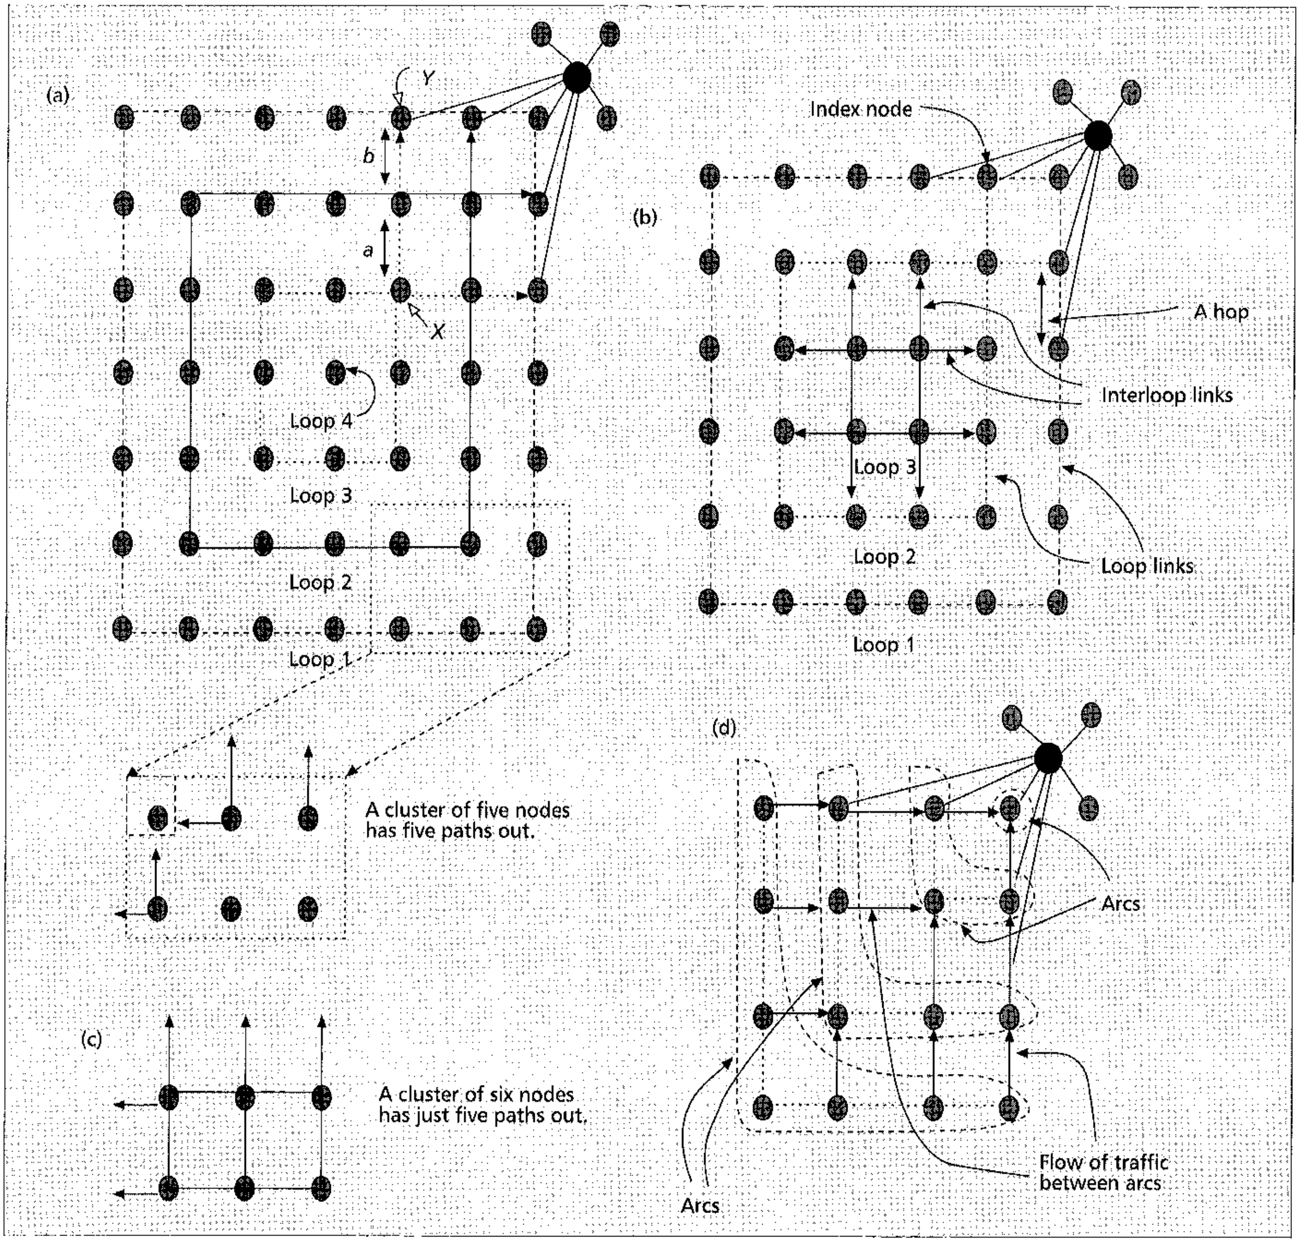

■ **Figure 2.** *Traffic entering an end office.*

are connected to index nodes. It is readily apparent that all the traffic may be routed from the outer arcs to the inner N arcs, and redistributed among these N arcs (using the cross-links between arcs) such that if an arc is connected to j index nodes, it carries traffic of jC_L. The total traffic on an arc may then be routed to the end office through the index nodes on that arc. It is to be noted that, again, if $K > 5$, it may be impossible to deliver the traffic to the inner N arcs, due to conditions similar to those described in the previous paragraph.

Thus, we have shown that if the number of links leading directly to an end office, K, is less than or equal to five per mesh, and are symmetrically arranged as shown in Fig. 2a, call blocking can be computed using the Erlang blocking formula. Let the traffic intensity generated in a mesh be ρ_q. The probability that a call is blocked is then equal to $P_b(\rho_q, 5\ C_L)$, where $P_b(\rho, M)$ is given by the Erlang blocking formula

$$P_b(\rho, M) = \frac{\rho^M / M!}{\sum_{k=0}^{M} \frac{\rho^k}{k!}} \tag{2}$$

Conclusions

In this article we describe a new approach to broadband wireless access. Key to this approach is its use of short, inexpensive, and extremely dependable focused free-space optical links to interconnect densely deployed packet-switching nodes in a multihop mesh arrangement. Each node can then serve a client, which may consist of a building containing PBXs and LANs (for fixed-point service), a picocellular base station (for wireless service), or both. The great virtue of this approach is that very high access capacity can be economically and reliably delivered over a wide service area. As an example, the optical cross-links can easily operate at 622 Mb/s (OC-12), and the multihop mesh can, in principle, deliver this rate to some number of clients. In general, since most clients will require significantly less access bandwidth, many clients can be served by a single access mesh which attaches to the infrastructure at a single access point.

We show that as long as the number of concurrent virtual connections of a single traffic class is less than the mesh capacity, these virtual connections can be arbitrarily distributed among an arbitrarily large number of clients, and all will enjoy guaranteed QoS. Furthermore, if we are willing to accept constraints on the distribution of virtual connections among clients, an even greater number of virtual connections can be simultaneously supported.

References

[1] A. S. Acampora, "Wireless ATM: A Perspective on Issues and Prospects," *IEEE Pers. Commun.*, vol. 3, no. 4, Aug. 1996, pp. 8–17.

[2] M. Umehira *et al.*, "ATM Wireless Access for Mobile Multimedia: Concept and Architecture," *IEEE Pers. Commun.*, vol. 3, no. 5, Oct. 1996, pp. 39–48.

[3] A. S. Acampora *et al.*, "Uninet: A Hybrid Approach for Universal Broadband Access Using Small Radio Cells Interconnected by Free-Space Optical Links," *IEEE JSAC*, vol. 16, no. 6, Aug. 1998, pp. 973–87.

[4] W. Honcharenko *et al.*, "Broadband Wireless Access," *IEEE Commun. Mag.*, vol. 35, no. 1, Jan. 1997, pp. 20–26.

[5] V. K. Garg and E. L. Sneed, "Digital Wireless Local Loop System," *IEEE Commun. Mag.*, Oct. 1996, pp. 112–15.

[6] Nat'l. Climactic Data Ctr., "Surface Airways Hourly/Airways Solar Radiation," TD 3280/ TD 3281, Mar. 1997.

Biography

ANTHONY S. ACAMPORA [F] (acampora@ece.ucsd.edu) is a professor of electrical and computer engineering at the University of California, San Diego, and director of UCSD's Center for Wireless Communications, where he is involved in a variety of wireless access research and education programs. His background and expertise is in the general area of telecommunication networks for multimedia services. This includes system architectures, basic digital communications, digital transmission and switching, packet communications, intelligent networks, broadband networks, ATM, quality of service maintenance, routing, the Internet, multiwavelength optical networks, and all aspects of wireless communications, including cellular, PCS, broadband wireless local loop, 3G, and broadband wireless data. His current research is focused on new system approaches for broadband wireless access to the Internet, mobility management, handoff, high-speed wireless data, and wireless multimedia media access protocols for use with smart antennas, ad hoc networks for the wireless home, and systems for universal wireless access. Prior to joining the faculty at UCSD in 1995, he was professor of electrical engineering and director of the Center for Telecommunications Research at Columbia University. He joined the faculty at Columbia in 1988 following a 20-year career as a researcher and senior manager at AT&T Bell Laboratories. He received his Ph.D. in Electrical Engineering from the Polytechnic Institute of Brooklyn and is a former member of the IEEE Communications Society Board of Governors. He has published over 160 papers, holds 30 patents, and has authored a textbook entitled *An Introduction to Broadband Networks: MANs, ATM, B-ISDN, Self-Routing Switches, Optical Networks, and Network Control for Voice, Data, Image, and HDTV Telecommunications.*

SRIKANTH V. KRISHNAMURTHY (krish@hrl.com) received his M.Sc (Hons.) degree in physics and B.E. (Hons.) degree in electrical and electronics engineering with distinction from the Birla Institute of Technology and Science, Pilani, India, in 1992, a Master of Applied Science degree in electrical and computer engineering from Concordia University, Montreal, Canada, in 1994, and his Ph.D. degree in electrical and computer engineering from the University of California, San Diego, in 1997. From 1994 to 1995 he was a graduate research assistant at the Center for Telecommunications Research at Columbia University, New York. Currently, he is a research staff member at the Information Sciences Laboratory, HRL Laboratories, Malibu, California, where he is involved in research work involving wireless ad hoc networks, hybrid satellite and land mobile networks, and Internet traffic modeling. His research interests are in wireless and satellite networks, ATM and IP technologies, optical networks, algorithms and complexity, queuing theory, and stochastic processes.