UniNet: A Hybrid Approach for Universal Broadband Access Using Small Radio Cells Interconnected by Free-Space Optical Links

Anthony Acampora, *Fellow, IEEE*, Scott H. Bloom, *Member, IEEE*, and Srikanth Krishnamurthy

***Abstract*—A new type of broadband access system is proposed for providing high-quality, bandwidth-upon-demand telecommunication services to the home and office. Communication terminals attach to the network via short radio links, and users can roam freely within a house or building unencumbered by the availability of wired "telecommunications outlets." Basic service is extended through small, high-capacity radio cells; the base stations are interconnected via short, highly focused free-space optical links in a multihop mesh arrangement; and the need for new broadband access cabling is totally surmounted. A tiered arrangement of radio cells further extends service to both out-of-building pedestrian and vehicular users, and service is universally available.**

Benefits of the approach are described, and issues involving reliability, availability, capacity, and hand off are identified and addressed. Opportunities to refine the basic approach are suggested for further study.

***Index Terms*—Broadband networks, local access, wireless communications.**

I. Introduction

FOR the past several years, the telecommunications industry has witnessed an explosive growth in the demand for: 1) nonvoice types of services (driven by so-called multimedia traffic, and suggestive of some unspecified combination of low- and high-speed data, voice, image, and video), and 2) service to nonstationary end terminals [1]–[5]. Despite this demand, three primary technical challenges have impeded the creation of a telecommunications access platform adequate to meet this demand. The first of these involves inadequate capacity afforded by the copper wires which typically connect homes and offices to core switches within the telecommunications network, and the cost of upgrading subscriber and local loop facilities. The second involves the limited availability of radio spectrum to meet the demand for nonstationary service. The third involves the need to guarantee universal service availability.

In this paper, a new systems approach, which we call UniNet (for universal access network), is proposed which surmounts all three broadband access challenges. The approach uses three tiers of radio cells, with some possibly large number of lower tiered cells nested within each higher tiered cell. To avoid extensive rewiring, short, free-space optical links interconnect the base stations of the lowest tiered cells.

The greatest fraction of access traffic would be handled by cells in the lowest tier. Here, each of a large number of geographically small, high-capacity radio cells (picocells) is responsible for serving some small number (perhaps one) of homes and/or offices. Most service subscribers would, in most cases, attach to "the network" via their home or office base station and, through this base station, enjoy complete freedom to roam within the building and its immediate surroundings. A cluster of contiguous picocells would thereby serve a large population of users, each of whom is served via a home or office picocellular base station. A packet mode of wireless access, similar to if not identical to the asynchronous transfer mode (ATM), supports the bandwidth-upon-demand needs of multimedia traffic [6], and imbedded virtual connection trees maintain quality-of-service guarantees while permitting rapid, decentralized handoff of live connections among adjacent picocells [7]–[10].

The picocellular base stations, themselves, are interconnected by a dense mesh of highly focused, free-space optical links which are physically short in length (under several hundred feet). As will be shown, the shortness of these links, and their highly focused nature, provides excellent margin against fog and other atmospheric disturbances (essentially 100% availability). Also, although the transmitters and receivers of the free-space optical links must be spatially aligned, the tolerances are such that the links can easily withstand extreme mechanical disturbances such as strong wind force.

By means of these free-space optical links, traffic generated within (or delivered to) any picocell would, in general, be relayed among a sequence of base stations in a multihop arrangement, eventually entering (or leaving) the wired network at a local end office. The picocellular base station itself is a small stand-alone unit containing a power supply, antenna, radio equipment, baseband processing equipment, a small packet switch (needed to relay traffic), and several optical transceivers, each aimed at a different one of the neighboring picocellular base stations. The richness of the optically based interconnecting mesh, and its alternative routing capabilities, serve to balance the traffic among the optical links, and also vastly improve system reliability.

Manuscript received August 1997; revised February 1998.

A. Acampora and S. Krishnamurthy are with the Department of Electrical and Computer Engineering and the Center for Wireless Communications, University of California at San Diego, La Jolla, CA 92093 USA.

S. H. Bloom is with AirFiber, Inc., San Diego, CA 92121 USA.

Publisher Item Identifier S 0733-8716(98)04805-7.

The second tier of cells is more conventional in appearance. Each "standard" cell covers an area with a diameter measured in terms of miles or tens of miles, and a variable number of picocells (ranging from zero to, perhaps, several thousand) is contained within its footprint. The base station of a standard cell attaches directly to an end office. Each standard cell serves three purposes: 1) extension of service to any location not served by a picocell, 2) extension of service to any vehicles traveling at a speed too great to be accommodated via picocellular handoff, and 3) provisioning of an alternate means to access the service office in the event of a picocellular failure and/or interruption of all of its optical links. Since a standard cell handles only the overflow traffic not served by its subtended picocells, it is envisioned that its traffic burden would remain modest.

The highest tier contains a contiguous raster of megacells, each with a diameter of several hundred miles. The megacells: 1) provide access from locations covered by neither a picocell nor a standard cell, and 2) ensure universal availability. It is envisioned that the megacells would be created by a constellation of low earth orbit (LEO) satellites. Once again, since a megacell handles only the overflow traffic from its subtended standard cells and picocells, its traffic burden should remain modest.

Since the footprint of a picocell is so small, its bandwidth is shared by only a small number of users (possibly one), and each subscriber thereby enjoys broadband service. By reusing the radio spectrum sufficiently often, the problem of limited availability of spectrum is surmounted. Furthermore, since each picocellular base station is served by its free-space optical links, new buried broadband cabling apparatus is unnecessary, and the capacity constraint of the preexisting copper wiring is surmounted. Finally, the standard cells and megacells extend service to regions not covered by picocells, improve reliability, and serve high-velocity vehicles for which handoff among picocells might be problematic; universal service is thereby assured with no blackout regions.

We note that, not only is broadband access extended to homes and offices, but, moreover, the access enjoys the additional virtue of being tetherless. Furthermore, pedestrians and vehicles are readily served.

We will assume that our system is based on asynchronous transfer mode (ATM) transport. Although this assumption is not essential, ATM supports the type of virtual connectivity and instantaneous, on-demand access to bandwidth preferred for serving newly emerging types of nonvoice traffic. Thus, all interfaces (radio at all three tiers, free-space optical, and serving office) are assumed to be ATM. For our sample system, the radio links support a peak data rate of 20 Mbits/s, and no optical link is operated at a rate greater than 155 Mbits/s. Also, although it is not mandatory, we will further assume that the backbone wired network is based on ATM.

In Section II, we present the basic system approach in greater detail. Section III contains a discussion of some optical mesh routing considerations, and Section IV presents a description of handoff and the embedded virtual connection tree. Section V contains representative descriptions of the radio and optical links, along with some representative capacity and optical link margin calculations.

II. System Description

The superposition of large radio cells on top of smaller cells in a three-tiered hierarchy is depicted in Fig. 1. Each picocell (tier 1) is created by a small, premises-based base station. Subscribers normally attach via their home or office base station and, through this base station, enjoy complete freedom to roam within the home or office environment and, to some extent, externally to the building as well.

If a subscriber should leave the area served by the home or office base station, service would be handed to either the base station serving the adjacent building, if there were no gap in coverage, or to the base station of a standard cell, the service domain of which spans several or many picocells. Finally, if a user should enter some environment serviced by neither a standard cell nor a picocell, then service would be extended via a megacell, whose range of coverage spans several or many standard cells. Service among megacells is contiguous.

The primary purpose of tier 1 is to provide adequate capacity to serve most users, thereby dramatically reducing the demand placed upon tier 2, and base stations at tier 2 connect directly to the core network switches, as shown in Fig. 1. However, since tier 1 base stations are residential or office units, new means must be provided to interconnect these to the core network. As shown in Fig. 2, these new means include short, free-space optical links, such that the tier-1 base stations are interconnected in a rich optical mesh. Each tier-1 base station thereby serves two functions. First, it accepts information from (and delivers information to) each terminal unit within its footprint via a radio-based air interface. Second, it serves as a cooperating *relay* station, accepting information (in the form of addressed data packets, possibly ATM cells) from inbound optical links, and routing each onto the correct outbound optical link.

As shown in Fig. 3(a), tier-1 base station equipment includes several optical transmitters and receivers, the radio and baseband equipment, a small electronic packet switch (perhaps an ATM switch), and a small control processor to manage the base station and serve as the local representative of the network control system. The radio equipment includes all apparatus needed to accept/deliver signals from/to the switch and deliver/accept these to/from the antennas. Fig. 3(b) is a physical representation, showing the base station as it might be installed in the attic of a home. The packet switch is needed so that, in addition to serving as source and sink of the traffic placed onto the optical network, each base station can relay (hop) packets from one optical link to the next, along a route (set of links) chosen at connection setup time to maintain quality-of-service (QoS) guarantees, until each packet arrives at an entry/exit port of the core network (for packets traveling along the access network toward the core network), or at the terminating tier-1 base station (for packets being distributed from the core network to the tier-1 base stations of the access network). Multihopping is important to enable the use of very short free-space optical links as needed to ensure freedom

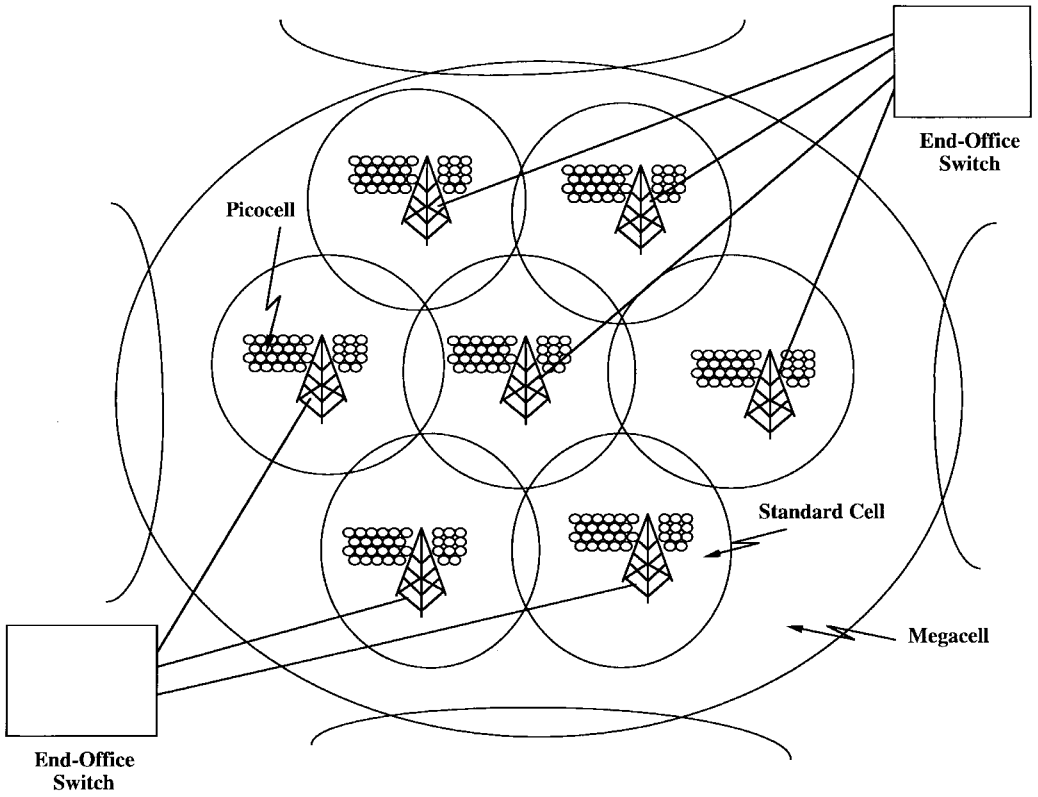

Fig. 1. Three-tiered radio access system.

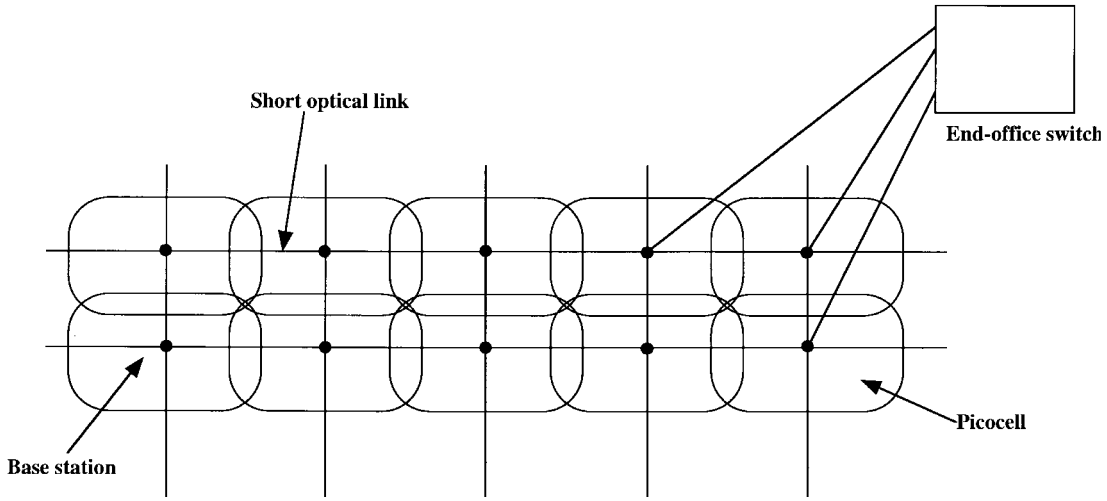

Fig. 2. Interconnection of picocells by dense optical mesh.

from atmospheric disturbances and good link availability. The dimensionality of the packet switch is at least 5×5 so that, in addition to the signals arriving from/sent to the radio interface, at least four optical transceivers can be supported. Note that this small switch is an *electronic* switch; both radio and optical packets are first converted to baseband electronic signals prior to switching, as in most multihop arrangements.

The multihop architecture is well known within the field of multiwavelength optical networks, but its application to free-space optical networks is novel, as is the overall approach whereby tiers of base stations are created, with the lowest tier interconnected by a new free-space optical network which avoids local wireline bottlenecks [11], [12]. The actual interconnection pattern of the optical network might consist of a recursive grid, a quasi-rectangular mesh in which nesting of access stations (tier-1 base stations in our case) into sublevels is allowed to easily enable the addition of new access stations (modular growth) without disturbing any more than one preexisting link [13], [14]. An additional benefit of the recursive grid is its compatibility with scalable routing algorithms, meaning that the computational complexity associated with the establishment of a multihop route for a new virtual connection scales linearly with the number of nodes in the grid, that is, the computational complexity per node is independent of the

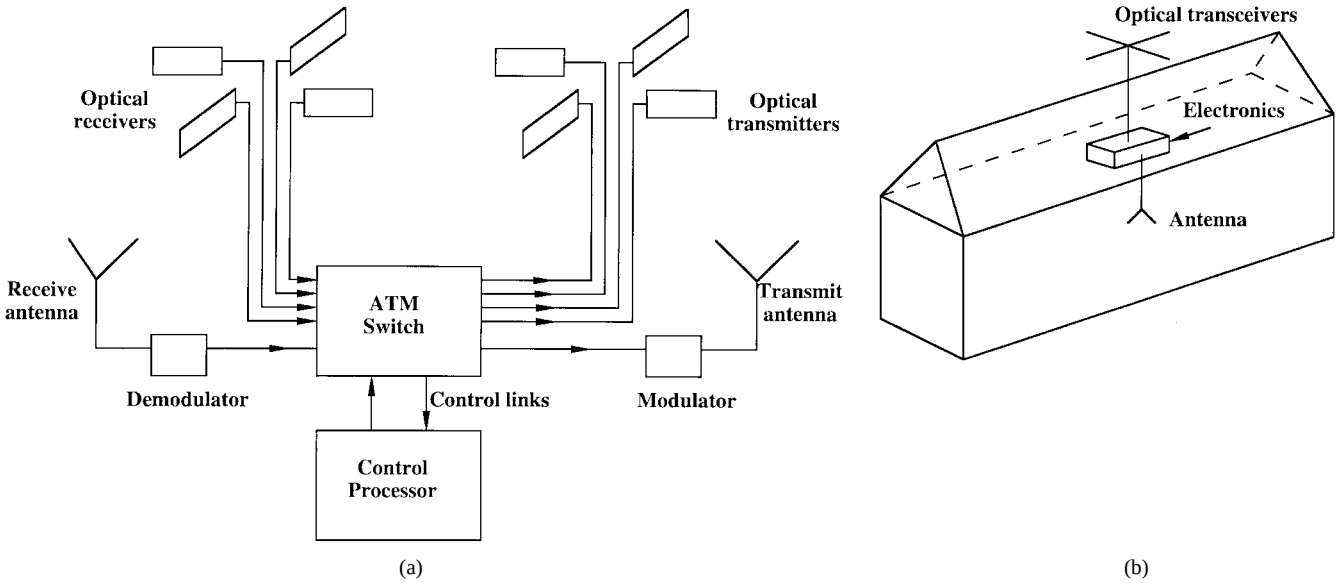

Fig. 3. (a) Base station apparatus needed for each picocell. (b) Physical representation.

number of nodes. Since, for the purposes of this paper, any topology can be chosen for the optical mesh, the interested reader is referred to the literature for additional information concerning the recursive grid.

Multihop routing for our radio access system is easiest to explain assuming fixed-point service, that is, no handoff; routing with handoff will be covered in Section IV. For fixed-point service, the route to or from a given base station taken by the packets associated with a given virtual connection flow is comprised of a sequence of links chosen by an admission controller at call setup time. When choosing this route, the admission controller must guarantee that all quality-of-service (QoS) objectives are met. This means not only that each optical link in the sequence can accommodate the new virtual connection without unacceptable QoS degradation, but also that the terminal radio cell can accommodate the new virtual connection. If handoff is supported (among clusters of picocells and between picocells and standard cells), then, as will be explained in Section IV, the admission decision must also guarantee that the overall traffic intensity (new plus preexisting calls) presented to the cell cluster remains acceptable.

Signaling for new connection requests is handled, in a conventional fashion, by means of a permanent signaling virtual channel connection existing between each base station and a control computer located within the end office.

Finally, to vastly improve the dependability of the free-space optical mesh, a set of back-up virtual routes is established at the time that the admission decision is made such that, in the event of a link failure or transient interruption of the optical beam, alternate routing can be instantaneously effected with minimal information loss. In general, virtual resources must be reserved on each optical link to accommodate instantaneous alternate routing, and an interesting research issue which will not be addressed in this paper involves the selection of primary and alternate routes such that, in the event of a single link failure, the additional traffic burden placed on the surviving links is minimized, that is, the virtual resources reserved for failure recovery should comprise only a small fraction of a link's capacity. This suggests that primary and alternate routes be chosen such that, in the event of a link failure or transient disruption, the rerouted traffic is "smeared" over as many disjoint paths as possible.

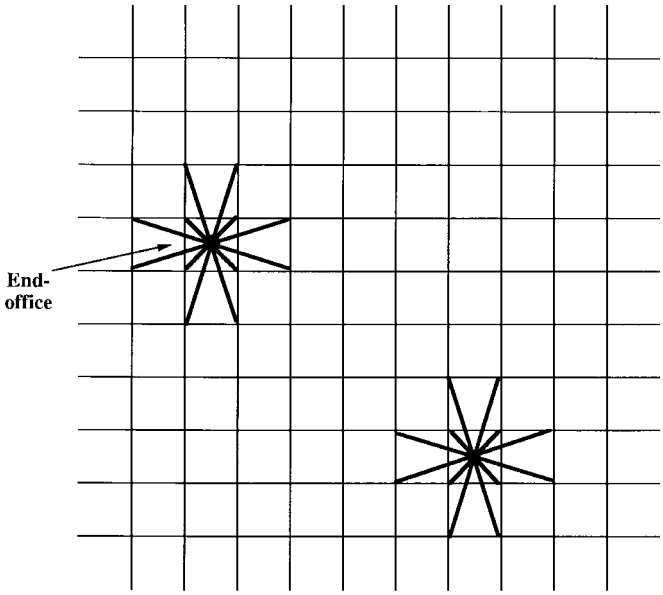

Fig. 4. Optical mesh, including end-office links.

III. Routing in the Optical Mesh

Of primary consideration to the flow of traffic along the links of the optical mesh is the avoidance of traffic bottlenecks or "hot spots." The situation is as depicted in Fig. 4, drawn for a simple topology in which the optical mesh consists of a rectangular grid, and each cross point represents a single tier-1 base station. Shown in Fig. 4 is a portion of the tier-1 access network, the purpose of which is to deliver traffic generated within the picocells to the end offices, and to

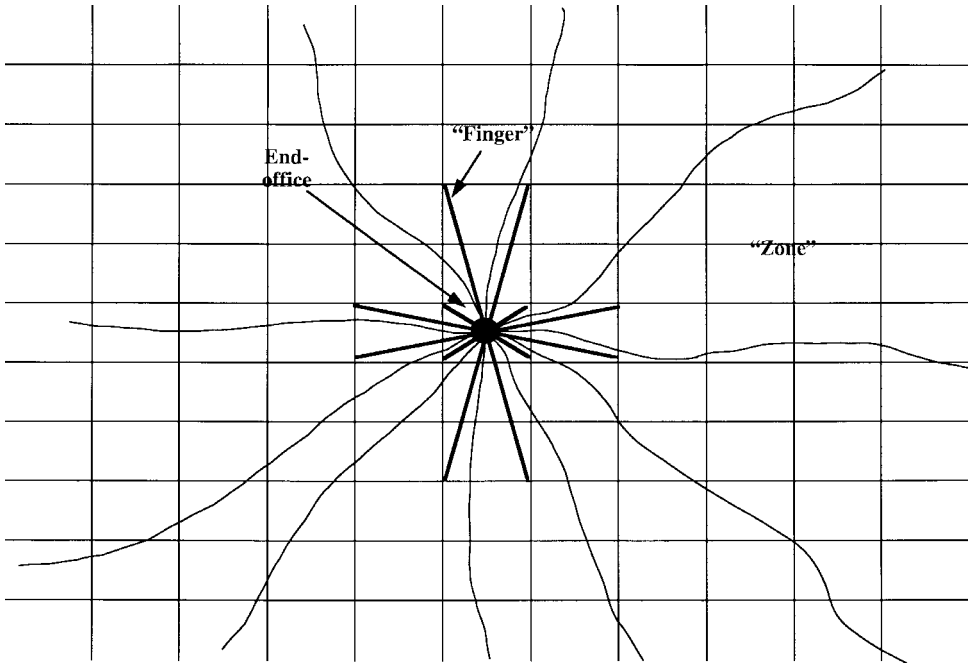

Fig. 5. Optical mesh, showing equal-traffic "zones."

distribute traffic from the end offices to the picocells. Each end office attaches to some multiplicity of neighboring tier-1 base stations by means of short free-space optical links (note that the base stations which are connected to an end office need an additional optical transceiver beyond the number needed by other base stations). To maintain their dependability, the optical links must be kept short, and the number of base stations connected to an end office is correspondingly limited. The objective is to route each virtual connection such that the flow of traffic is evenly distributed among the limited number of optical links leading directly to/from an end office.

Some simple calculations are in order. Suppose that all picocells generate the same level of traffic intensity C, and suppose that each end office is responsible for N base stations. Let each end office be equipped with L optical links. Then, assuming that traffic can be "balanced" among these L links, each must handle a traffic intensity $C_L = NC/L$. (Note that, in general, other optical links will handle a lesser traffic intensity since, in our multihopping arrangement, the traffic handled by the optical links grows as traffic is relayed toward the end offices.)

The above provides a "best case" relationship among the number of optical links needed per end office (L), the multihop capacity deflator (C_L/C, the ratio of the required link capacity normalized by the delivered user or picocell capacity), and the number of picocells supportable by an end office. It is the objective of routing in the optical mesh to "smear" the traffic such that this "best case" applies. While we should view such "optimum" routing to be an open research problem, some simple observations and heuristics will be offered.

A. A Possible "Zoned" Approach

First, it is possible to define a "zone" associated with each "finger," or optical link, emanating from an end office. Referring to Fig. 5, a zone is defined to be a set of base stations chosen such that: 1) traffic associated with all base stations of a given zone enters/exits the end office through a common finger, and 2) the same traffic intensity is associated with all zones. The objective of the routing algorithm, then, is to choose a route for each new call attempt such that all traffic associated with a given zone enters/exits the end office through that zone's finger. Such zones may be permanently defined based on the average traffic intensity presented by each picocell. Alternatively, to achieve better routing efficiency, zones may be dynamically defined in response to instantaneous traffic patterns. For example, in the latter case, the zone associated with a currently lightly loaded finger might be adaptively enlarged to include a greater number of base stations, while the zone associated with an adjacent heavily loaded finger might be reduced to include a smaller number of base stations, such that, for *future* calls, it is more likely that a route will be chosen through a lightly loaded finger. Then, as new virtual connections are made and old ones terminated, the boundaries of the zones will continuously adapt in an attempt to maintain load balance among the fingers.

Note that, with such a "zoned" approach, traffic originating or terminating within a base station at the boundary of two zones might be bifurcated to best achieve load balance. In a similar spirit, traffic originating or terminating within a base station on the boundary of two zones served by two different end offices might be bifurcated among those offices (furthermore, when adapting the "footprint" of a zone to the prevailing traffic, a given base station can be reassigned, for future calls, to a different end office if this will facilitate load balancing). Note further that the geometrical boundaries of the zones may be quite irregular. It is the intent to capture a common traffic intensity within all zones (for permanently assigned zones), and to tailor the zonal boundaries to the

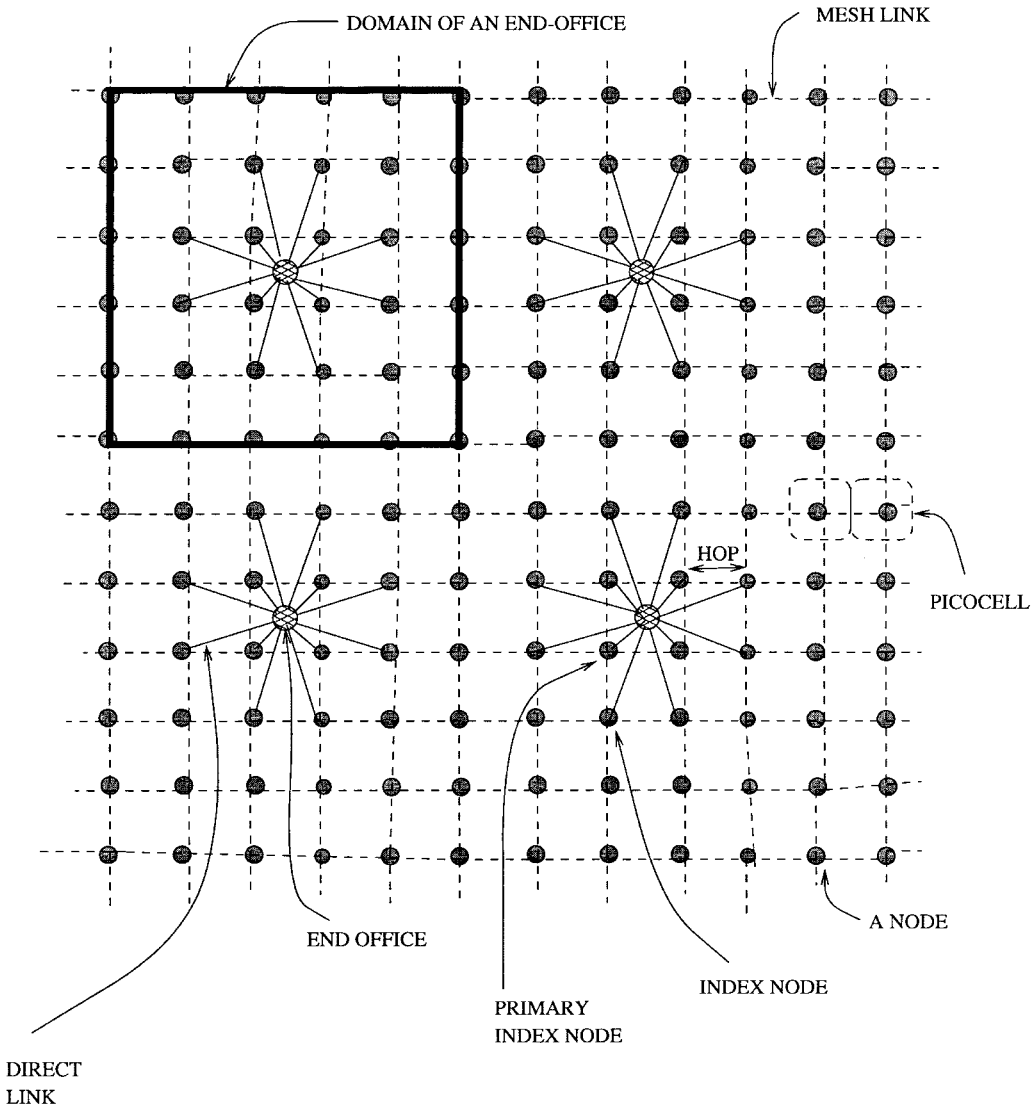

Fig. 6. Rectangular mesh.

remaining available capacity of the fingers (for dynamically assigned zones), and both the traffic generated within a given picocell and the geometrical deployment of picocells are very much demographically dependent. In general, one should not expect uniform spacing of the picocells as might be (incorrectly) implied from Fig. 5, which is intended only to show a connectivity pattern among cells, but does not speak at all to their spacing. Finally, note also that the richness of the mesh, the possibility of bifurcating traffic among fingers and end offices, the ability to enter/exit an end office via any of several fingers, and the ability to select alternate routes all provide substantial protection against disruption of an optical link or set of links (either as a result of equipment failure, atmospheric impairment, or transient interruption of an optical beam due, perhaps, to a migrating bird!). It is envisioned that, as part of fault management, a diagnostic routine is continuously executed which permits the base stations to rapidly sense disruption of an optical link and to initiate corrective action (alternate routing). Not only can the surviving optical links be used to communicate management and control commands during the restoration phase subsequent to link disruption, but, if necessary, the radio interfaces can *also* be used to deliver these fault management messages.

B. Conditions for Avoiding "Hot Spots" on the Mesh

One problem associated with the zoned approach just described is its loss of trunking efficiency; if the finger associated with a given zone is loaded to capacity (meaning that no additional virtual connections can be added without unacceptably degrading the QoS guarantees), but adjacent fingers are under utilized, it is not possible (at least not with fixed zones) to accept additional traffic generated within the zone associated with the fully loaded finger at another finger. Thus, virtual connections that *might* have been accepted by an under utilized finger will be blocked. Consequently, it is useful to examine conditions under which better trunking efficiency might be realized. In fact, we can find some simple conditions under which a routing algorithm exists such that the loading of all optical cross links not leading directly to an end office (*mesh links*) is always less than the loading on an optical link that

does lead directly (*direct links or fingers*) to an end office. In such a case, simple Erlang blocking formulas can be applied, independent of the detailed traffic distribution among picocells.

To begin, consider Fig. 6 drawn for a rectangular mesh. Suppose that each direct link has a capacity of C_L virtual cells, and each picocell generates traffic of value $C \leq C_L$. (Alternatively, each picocell may terminate traffic of value $C \leq C_L$; it is readily shown that for either traffic inbound to an end office or outbound from an end office, each bidirectional mesh link is used in only one direction, and therefore, by symmetry, the Erlang blocking formulas are independently applicable for each direction.) Further, assume that each mesh link can carry traffic load of C_L. Let us define each base station (which forms a picocell) to be a node. The nodes directly connected to an end office are called *index nodes* (Fig. 6). The unit distance between two adjacent nodes is called a *hop*. Let us impose the constraint that a node can communicate only with its closest (in terms of the minimum number of hops) end office. Then, the rectangular mesh may be subdivided into *domains*. A domain consists of an end office and all of the nodes communicating with it. A node may belong to more than one domain if it is equidistant to more than one end office. The index node, geographically closest to the end office it is associated with, is called the *primary index node*.

Let us now consider Fig. 7, in which there are exactly four direct links to the end office, i.e., the maximum traffic intensity which can be handled by the end office is $4C_L$. The traffic generated within the domain must be routed to the end office in the domain. In each *quadrant* of the domain, concentric loops can be drawn as shown in Fig. 7(a). A *loop* connects nodes along the boundary of a rectangle. If, in any quadrant, the generated traffic is less than or equal to C_L, we can use those mesh links within the quadrant not associated with the mesh links comprising the loop to first route the traffic to the outermost loop in the quadrant. The direct link attached to this loop can then be used to route the traffic to the end office. If, on the other hand, the traffic generated in a given quadrant exceeds C_L, i.e., is between C_L and $4C_L$, then traffic in excess of C_L can be redirected to neighboring quadrants using the mesh links between the quadrants (not shown in Fig. 7). Note from Fig. 7(b) that there will be at least four mesh links leading out of a quadrant, and hence, it is always possible to redirect excess traffic of up to $3C_L$ to the neighboring quadrants, thereby uniformly distributing the traffic among the quadrants. As long as no node generates more than C_L units of traffic, and as long as the total traffic is less than $4C_L$, then the only blocking encountered is due to the direct links being fully loaded. The call-blocking probability can then be computed using the Erlang blocking formula given by

$$P_b(\rho, M) = \frac{\dfrac{\rho^M}{M!}}{\displaystyle\sum_{k=0}^{M} \frac{\rho^k}{k!}} \tag{1}$$

where the number of available circuits is M and the load is ρ. In the above case, if the load generated in the domain is ρ_D, the call-blocking probability is given by $P_b(\rho_D, 4C_L)$.

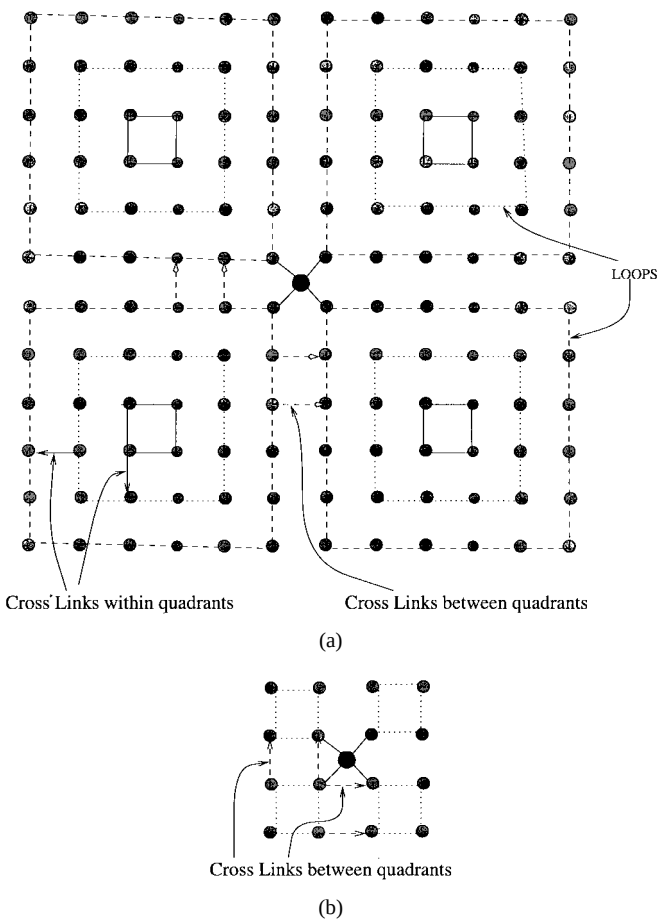

Fig. 7. Case 1: traffic restricted within a domain.

We next impose the constraint that all traffic generated within a quadrant must reach the end office by using only those direct links which lead into that quadrant (Fig. 8). Furthermore, the route taken by any call is constrained to lie entirely within the quadrant. Let K be the number of direct links from the end office into a quadrant. It is essential that the value of K be less than or equal to five in order to avoid blockage on the mesh links (to be discussed later).

Now, consider a quadrant of size $R \times R$. Let $N = \lceil K/2 \rceil$, where $\lceil x \rceil$ is the closest integer greater than or equal x. Suppose that $R \geq N$. In such a case, it is always possible to draw $M = \lceil R/2 \rceil$ loops in the quadrant (sometimes a loop may consist of just a single node). We label loops as shown in Fig. 8(a) and (b), 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* [see Fig. 8(b)]. 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, then i is at most $(R-1/2)$, and the number

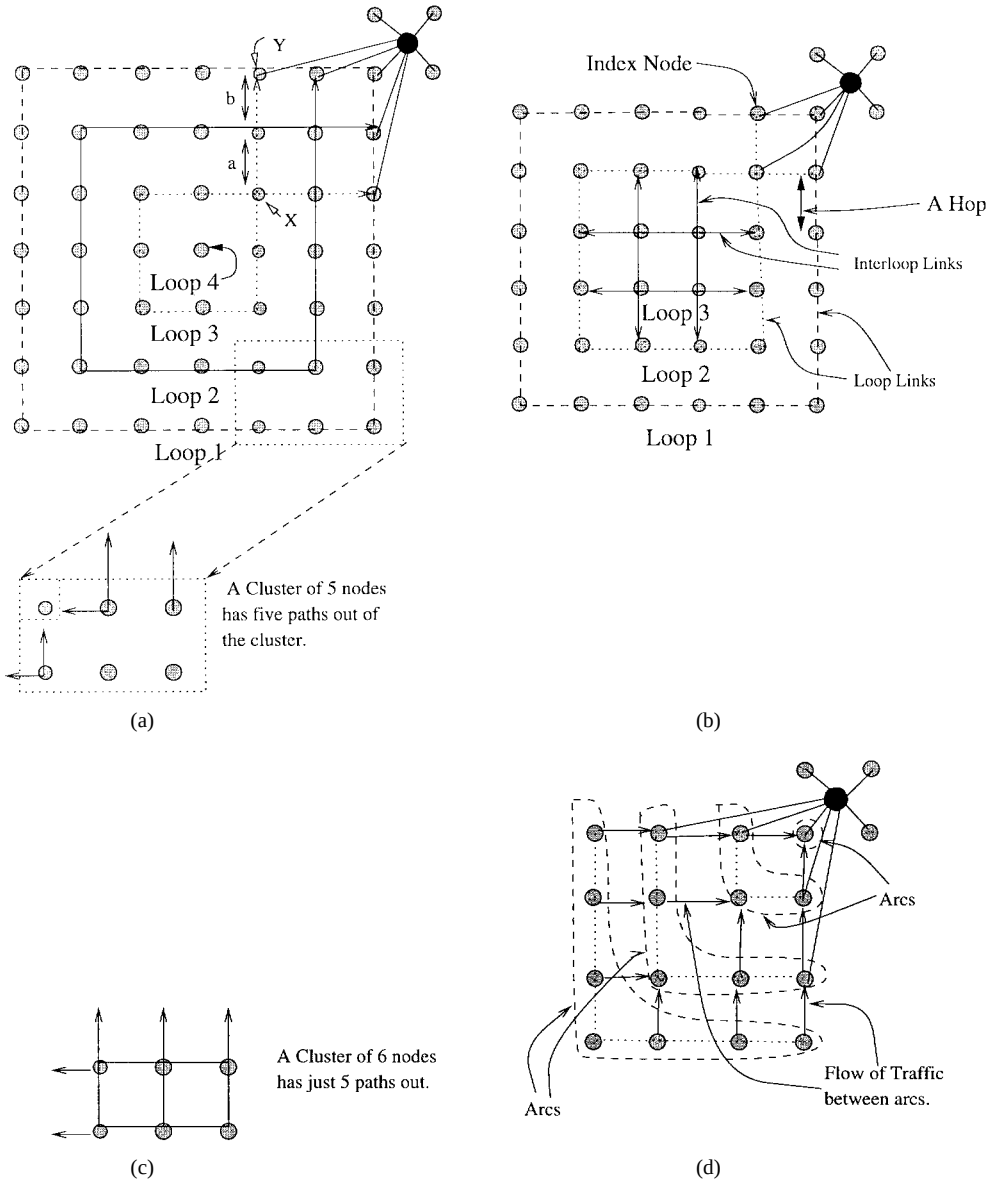

Fig. 8. Case 2: traffic restricted within a quadrant.

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. 8(b), 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 of 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. 8(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.) 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 1–3 (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 counter-clockwise direction), we can successfully route the entire traffic to the *corner node* on the loop, closest to the endoffice, without encountering blockage on the mesh links. The interloop links connected to this node lead directly to an index node [see Fig. 8(a) and (b)]. 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 $5C_L$. This is illustrated in Fig. 8(c) and (d). Let all of the traffic be generated by a cluster of nodes at a corner of the quadrant. If a maximum load of $5C_L$ is generated by a cluster

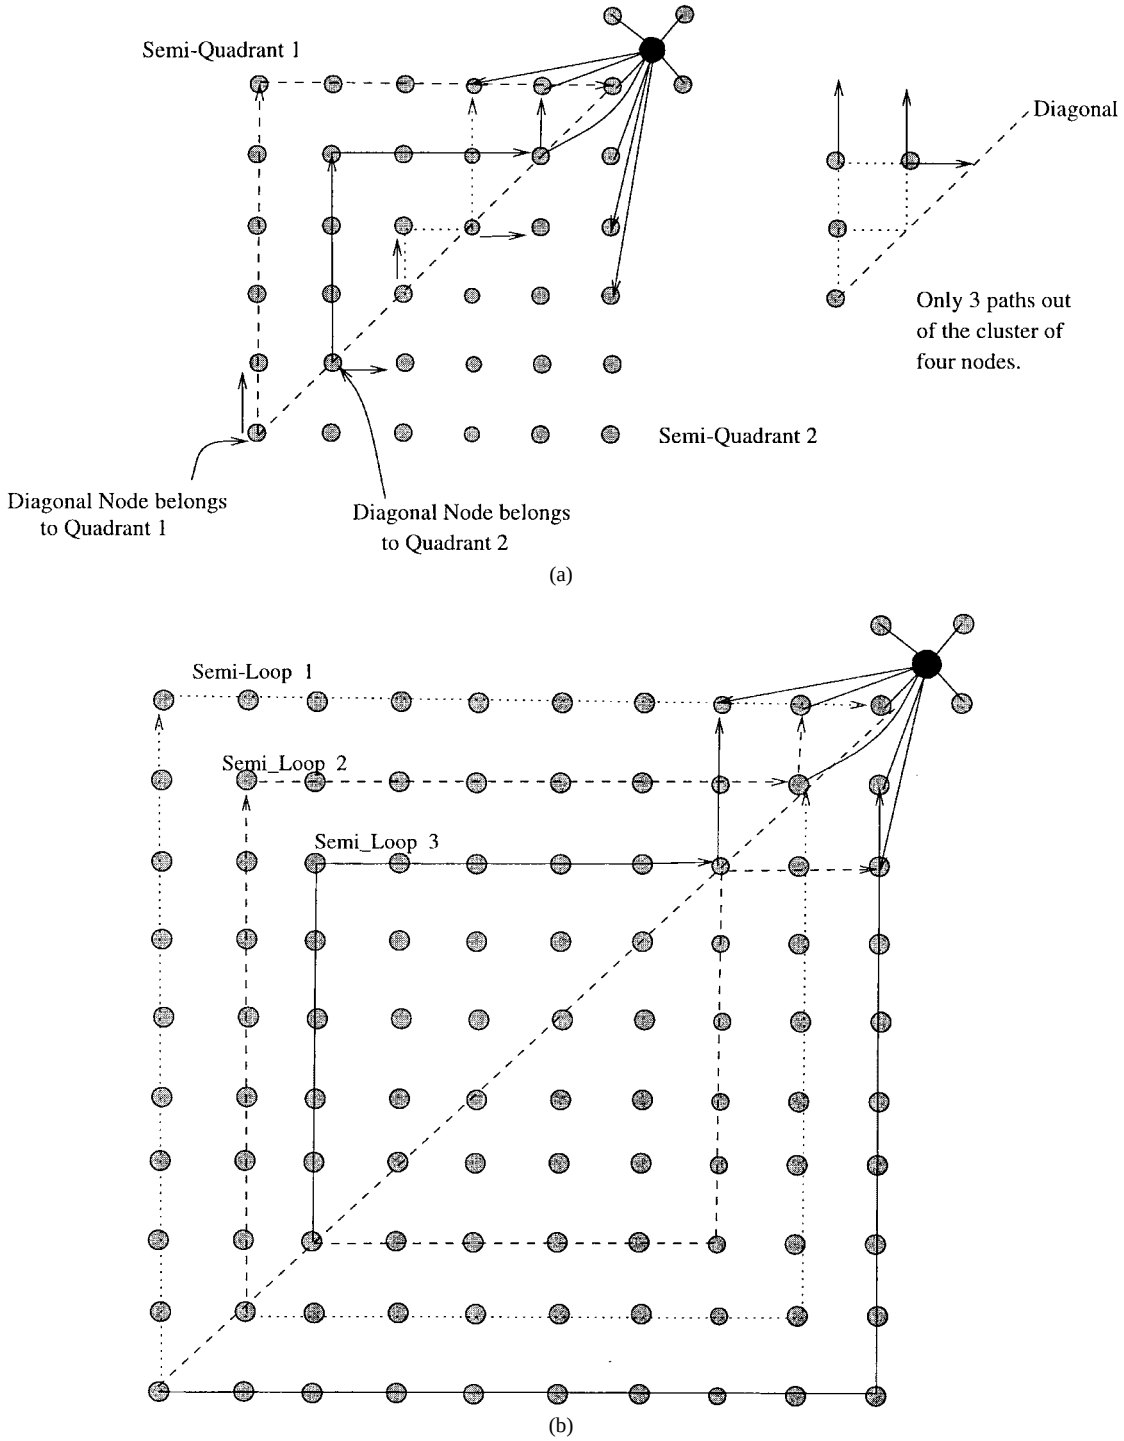

Fig. 9. Case 3: traffic restricted within a semi-quadrant.

of five nodes in a corner of the quadrant, then, since there are five mesh links leading out of this cluster, it is always possible to redistribute the load [Fig. 8(c)] 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 quadrant as shown in Fig. 8(d), 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 quadrant is greater than five, given that the traffic within the quadrant is less than $K\ C_L$, one cannot guarantee that there will be no blockage on the mesh links. (Note that this, in turn, implies that there may be up to 20 links total leading out of an end office.)

Now, suppose $R \leq N$. For this case, one may subdivide the quadrant into *arcs* as shown in Fig. 8(d) [there are four arcs shown in Fig. 8(d)]. Note that the inner N arcs (an arc may consist of a single node, i.e., the primary index node) are connected to index nodes. It is readily apparent that all of 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 a traffic of $j\ C_L$. 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 quadrant, and are symmetrically arranged as shown in Fig. 8(a), then call blocking can be computed using the Erlang blocking formula. Let the traffic intensity generated in a quadrant be ρ_q. Referring to (1), the probability that a call is blocked is then equal to $P_b(\rho_q, 5C_L)$.

If the number of direct links from the end office into a quadrant is six (i.e., the total number of links leading out of the end office is now 24), one can eliminate blocking on mesh links by imposing the additional constraint that routes taken by calls generated within a *semiquadrant* [see Fig. 9(a) and (b)], to reach the end office, must be constrained to lie entirely within the semiquadrant. Adjacent diagonal nodes belong to different semiquadrants. One can draw *semiloops* within a semiquadrant as shown in Fig. 9. Since the number of direct links is six, at least three semiloops can be drawn in a semiquadrant (in the limiting case, a semiloop may have just a column of nodes). If the maximum traffic generated within a semiquadrant is less than or equal to $3C_L$, it is possible to distribute the traffic among the three outermost semiloops such that the traffic carried by any semiloop is less than or equal to C_L. Note that if the traffic is greater than $3C_L$, one may encounter blockage on the mesh links. For example, let the traffic generated be $4C_L$. If all of the traffic is generated in the cluster of four nodes at the corner of the semiquadrant adjacent to the diagonal [Fig. 9(a)], then, since there are only three links leading out of this cluster, C_L calls must be blocked. Thus, the maximum load which can be routed to the end office without encountering blockage on mesh links within a semiquadrant is $3C_L$. Referring to (1), the probability of a call being blocked under such constraints is therefore $P_b(\rho_s, 3C_L)$, where ρ_s is the total traffic intensity generated in the semi-quadrant.

IV. Cell Handoff

Routing in the optical mesh, as described in Section III, assumes that, for its entire duration, a given virtual connection flows to/from the base station from which it was originally generated over a fixed set of optical links. In reality, since users are free to roam among radio cells, the virtual connections must be "handed off" among radio cells as a user "travels" among a sequence of cells. At the time of handoff, it is required that: 1) a new route be found leading from the newly serving base station back to the end office, and 2) new routing instructions be provided to the switch contained in each base station along the new route. Implicit in the selection of a new route is the ability to maintain QoS over each link in the new route and within the radio cell accepting responsibility for the handoff call. Ensuring that QoS objectives are met can present a critical real-time processing challenge, especially when handoffs occur very frequently such as might be expected in a picocellular system. Delivering new routing instructions to the switches, after QoS has been ascertained, is a further challenge. Both challenges are adequately met by the virtual connection tree.

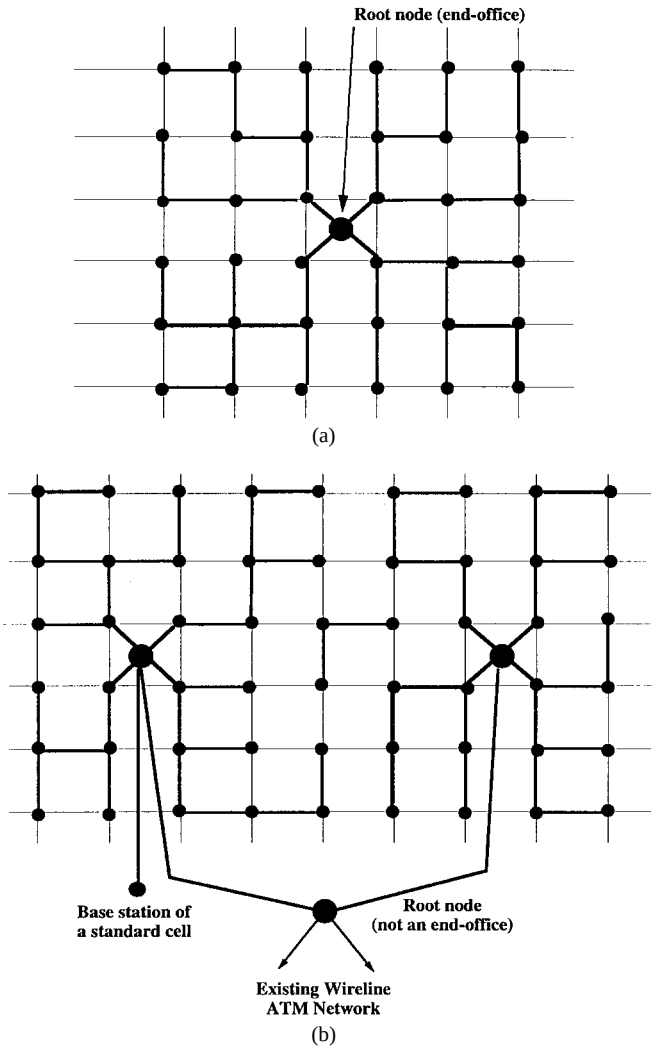

Fig. 10. (a) Virtual connection tree. (b) Alternative virtual connection tree.

The basis of the virtual connection tree is the creation, at call setup time, of a *set* of virtual connections for that call, each originating from a root node and each terminating in a different base station or leaf (actually two "trees" are set up: one leading to the root node, one leading from the root node, such that duplex connections can be handled). At call setup time, an admission controller determines whether or not a new call can safely be admitted to the tree. For a new call to be admissible, it must be determined that, on a statistical basis, and as a result of user mobility, the likelihood that too many calls will exist within a leaf of the tree, or flow on any given branch (link) of the tree, is acceptably low. If too many calls exist within a cell or flow on some given branch, then that cell or branch is said to be in overload, meaning that the QoS objectives cannot be met. Overload, then, either causes an unacceptable (but transient) degradation in service (i.e., delay or lost packet objectives are temporarily not met) or causes the call which caused overload to be dropped (as might occur, for example, for real-time traffic such as voice). In

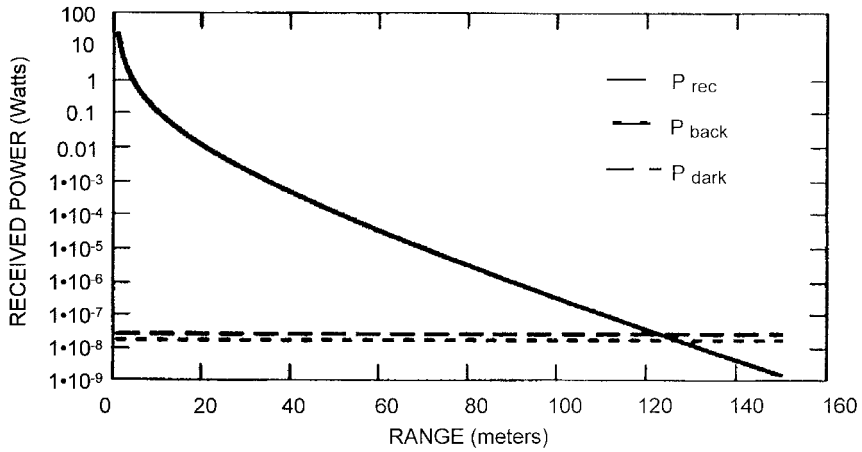

Fig. 11. Optical link budget.

essence, at call setup time, the admission controller is making a guarantee: each new user can freely roam among all leaves (base stations) in the tree and, only occasionally (that is, with some guaranteed low probability) will unconstrained motion of the users cause an overload condition to arise. To maintain this guarantee, the admission controller blocks new call requests at some predetermined threshold.

If a newly requested virtual connection is admitted, a route in each direction will be chosen by that user.

As applied to our hybrid radio-optical link broadband access network, a connection tree might appear as shown in Fig. 10(a). Note that the root of the tree is an end office, and branches of the tree extend to every base station served by that end office. Note further that, as shown in Fig. 10(a), the "fingers" leading to each end office are "main branches" of the tree, that is, the tree is defined such that the zones of Section III-A are maintained and load balancing among the fingers is provided for. Finally, note also that since each "leaf" can be reached through any of several "branches," we are free to choose the "branches" for a newly requested virtual connection such that: 1) each base station is included in the tree, and 2) load balancing is achieved among branches at a given distance from the root. Thus, when setting up a tree for a new call, each base station served by a given end office would be included, but the path leading to each leaf would be chosen such that, as users roam among cells, all branches at a given depth from the root node carry, on average, the same traffic intensity.

It is also possible to define the root of the tree as shown in Fig. 10(b), in which case an ATM switch behind the end offices serves as the root node for some larger tree, the footprint of which includes the base stations served by several end offices. Also shown in Fig. 10(b) is the inclusion of the base station of a standard cell as one leaf in the tree, so that a user can also choose service from a standard cell if conditions warrant (i.e., if the user has roamed to some vicinity not served by any picocell).

V. Optical and Radio Link Calculations

In this section, we demonstrate that, because of the relatively short range of adjacent lasercom transceiver nodes (100–200 m) envisioned in the picocell network architecture, the picocell free-space optical link, hereinafter called the lasercom link, can have all weather availability at OC-3, and higher data rates with an inexpensive, compact, low power, eyesafe transceiver. We also provide some estimates for the target capacity of the radio picocell.

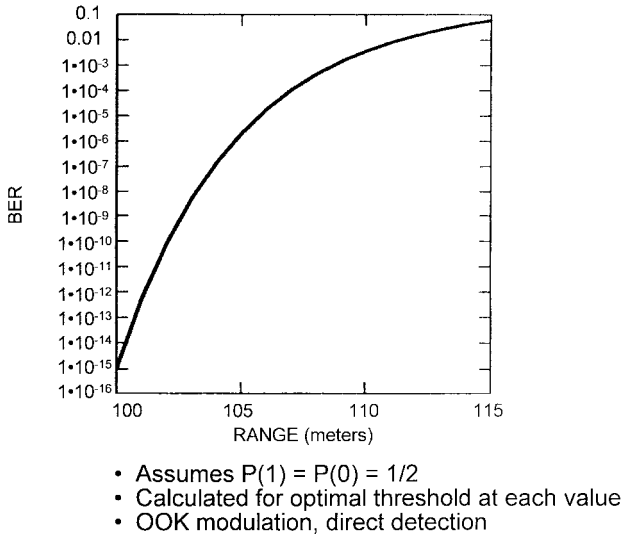

Fig. 12. Bit-error rate at optimal threshold, OOK modulation.

The overall system performance of a lasercom link is easily quantified using a link budget, the techniques being similar to those used to evaluate microwave links. There are three important parameters: transmitter power, propagation losses, and receiver sensitivity. The receiver sensitivity relates the amount of optical power needed to maintain the signal-to-noise ratio required to achieve a desired quality of service.

The received signal power can be calculated from

$$P_R = P_T \eta \frac{A_R}{d^2 \delta^2} e^{-\alpha d} \tag{2}$$

where

P_R received power;
P_T transmitted power;
η optical efficiency;
A_R area of receiving aperture;

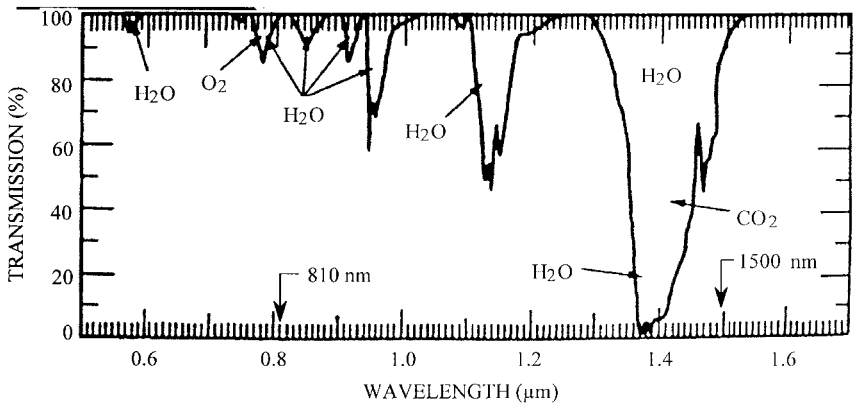

Fig. 13. Atmospheric absorption, 2 km, sea level.

δ angular divergence of the transmitted beam;
α atmospheric attenuation factor;
d distance between transmitter and receiver.

For relatively clear air, the exponential term is small, and received laser power scales with $1/d^2$ and $1/\sigma^2$; for propagation conditions of heavy attenuation, such as fog, it scales exponentially.

The amount of background radiation collected by the transceiver is dependent on the receiver's field of view and its optical bandwidth. The field of view cannot be made arbitrarily narrow due to alignment issues discussed below, and the optical bandwidth cannot be made arbitrarily small due to poor absolute transmission through most narrow dielectric filters. The received background power P_{Back} is calculated from

$$P_{\text{Back}} = \frac{0.2\ \text{W}}{\text{m}^2 \cdot \text{nm} \cdot \text{sr}} \cdot A_R \cdot F_{BW} \cdot \delta^2_{\text{FOV}} \cdot N \qquad (3)$$

where

m meters;
nm nanometers;
sr steradians;
F_{BW} filter bandwidth;
δ^2_{FOV} angular field-of-view.

The value for irradiance of the sky was measured at 10° from the sun using a specially calibrated telescope/detector assembly.

The noise-equivalent power of the detector/preamp module can be calculated from a double-sided noise density usually expressed in watts/$\sqrt{\text{hertz}}$. Fig. 11 shows the received signal power, background power, and noise-equivalent power as a function of range for the following set of parameters:

telescopy diameter:	10 cm
laser power:	20 mW
laser divergence:	2 mrad
telescope efficiency:	0.5
atmospheric attenuation:	varies
data rate:	OC-3 (155 Mbits/s).

In order to determine the availability of the link in all weather conditions, an attenuation of 392 dB/km, the worst (and very rarely experienced) fog was assumed in the calculation.

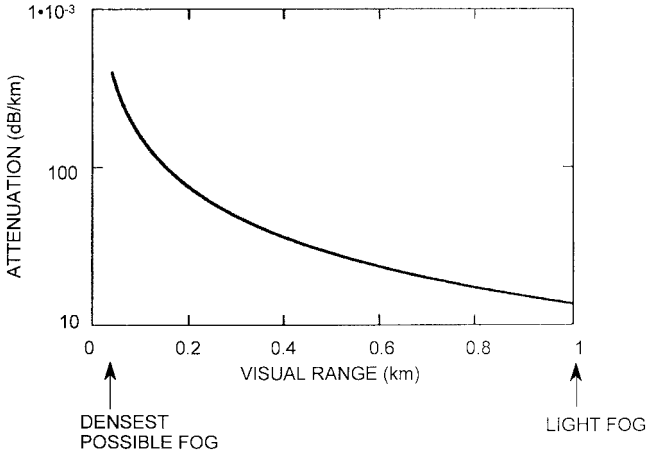

Fig. 14. Visual range compared to attenuation at 800 nm.

The bit-error rate (BER) for the p-i-n detector (a unity gain photodiode) circuit can be calculated using the previously obtained values [15]. This circuit has an overall transimpedance gain of 10^6. Fig. 12 shows the calculated BER versus range for this case. The BER was calculated assuming equally likely ones and zeros and on–off keyed (OOK) modulation with direct detection. The threshold value was chosen at each range point with the criterion of minimizing BER at that point. Clearly, the system is capable of very low (equivalent to fiber optic) BER at ranges of 100 m or less in the most attenuating condition. For the avalanche photodiode detector, the range for a given BER is increased by approximately 15 m.

Eye safety is always an issue when working with laser systems. The wavelength employed here is not inherently eyesafe (i.e., light can pass through the cornea to be imaged on the retina). ANSI standards Z131.1-1986 maintain that the maximum flux entering the eye at this wavelength is 2 mW/cm^2 [16]. A transmit laser power of 20 mW implies a transmit aperture of 10 cm^2 minimum, a criterion easily met by our conceptual transceiver.

Path loss in an optical communications link is affected by atmospheric absorption, scattering, and turbulence, however, the magnitude of each loss mechanism varies greatly. Fig. 13 shows sea level atmospheric transmission for a 350 m path as a function of wavelength calculated using MODTRAN IV. Fortunately, for available semiconductor laser diode wavelengths

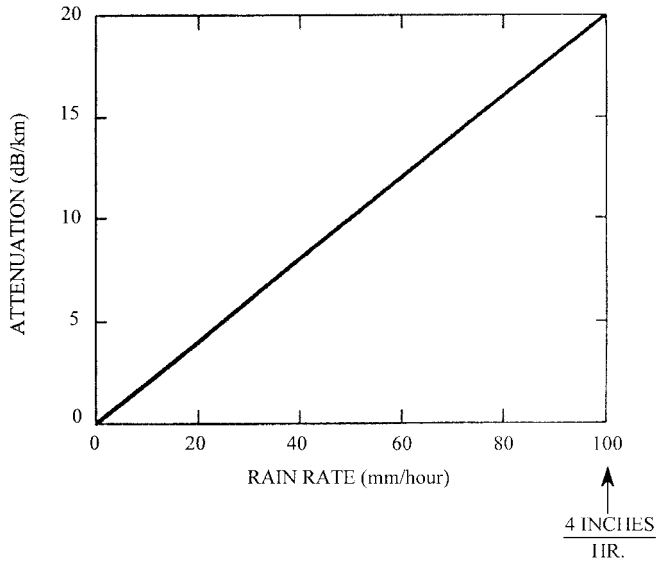

Fig. 15. Attenuation due to rain.

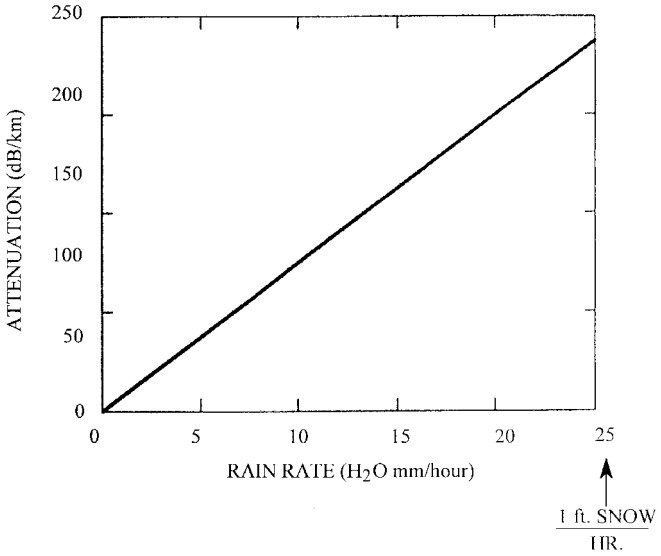

Fig. 16. Attenuation due to snow (1 in rain = 12 in snow).

of 810, 1300, and 1550 nm, path loss is negligible due to molecular absorption [17].

Path loss due to scattering is a function of optical radiation wavelength, number density, and size of scatters in the path. In order to quantify the amount of scattering attenuation along a particular path, an approximation based on Mie scattering and visibility range can be used [17]. The most severe path losses at near infrared wavelengths come about due to fog. Fog occurs when the relative humidity of an air particle approaches its saturation value. Some of the nuclei then condense into water droplets, forming the fog. Fig. 14 shows attenuation versus visual ranges up to 1 km [18]. Precipitation in the form of rain occurs when the water droplets condense up to mm sizes. This causes both scattering and water absorption losses along the path, with the size distribution of the drops determining the relative magnitude of these effects. The difference in fog and water droplet size, three orders of magnitude, accounts for the orders of magnitude difference in attenuation for typical rain (Fig. 15) versus typical fog. Additionally, the relatively large size of typical raindrops compared with near infrared wavelengths permits the drops to forward scatter a significant fraction of incident optical power. Attenuation due to snow has been measured by several groups, and is displayed in Fig. 16 for snow rates up to 1 ft/h. Typically, the path loss due to snow lies somewhere between fog and rain; however, the relationship is very complex, and is best measured experimentally.

Path loss fading due to scintillation can be very significant (25–30 dB) for long-range paths through the atmosphere. Scintillation is caused by thermal fluctuations which induce random fluctuations in the index of refraction along the path contained in the beam's cross section. This causes the air to act like sets of small prisms and lenses which deflect the beam of light. In the plane of the receiver, a speckle-like pattern appears with light and dark cells which have a characteristic size that scales like the square root of wavelength times path length. The generally accepted form for the probability distribution of the expected intensity is log normal for weak scintillations;

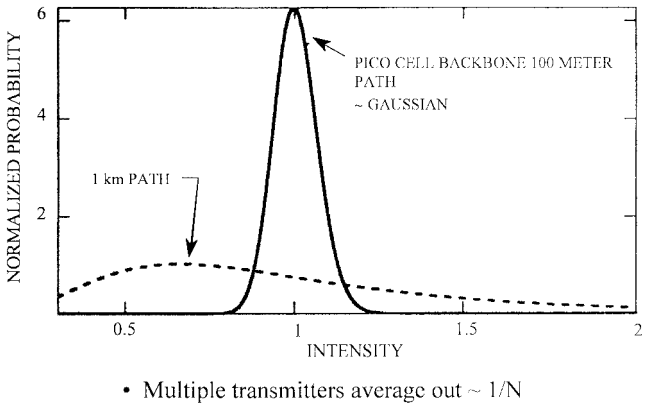

Fig. 17. Scintillation effect on beam intensity.

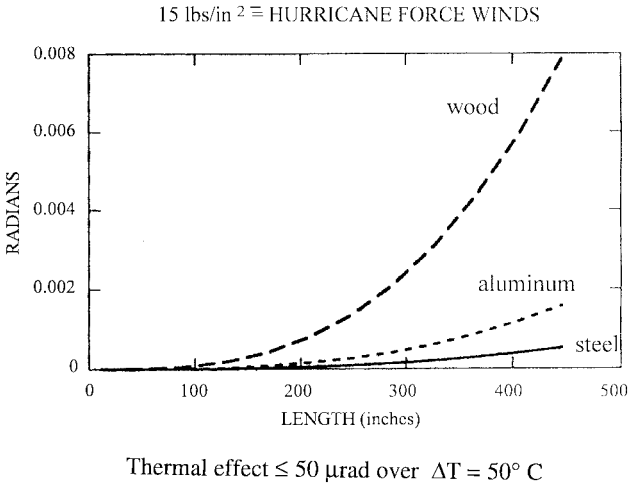

Fig. 18. Wind-bending effect on poles.

however, this tends to saturate at variances near 0.35 for heavily scintillated paths. The log-normal distribution for the intensity and variance of intensity is given by

$$P(I, \sigma_x^2) = \frac{1}{2I\sqrt{2\pi\sigma_x^2}} \exp\left[-\frac{(\ln I + 2\sigma_x^2)^2}{8\sigma_x^2}\right]. \quad (4)$$

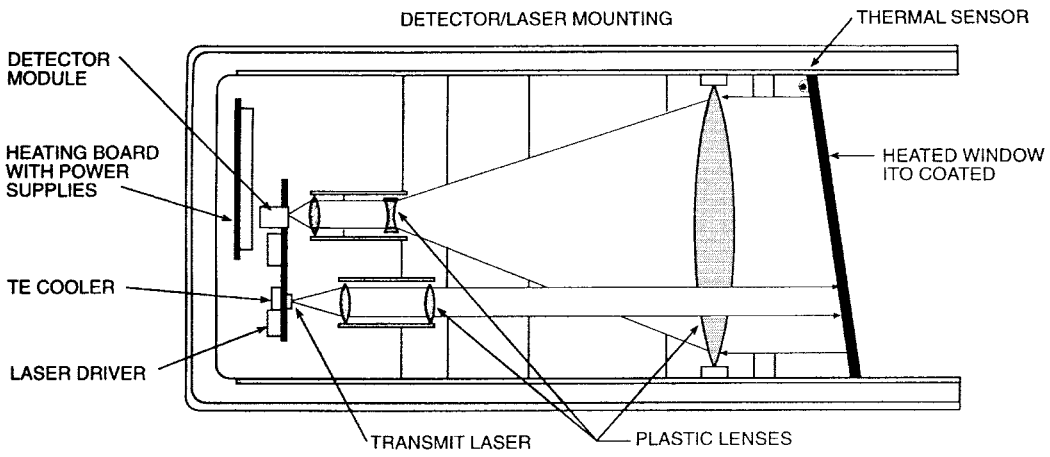

$$\blacksquare\ \text{Diffraction limited} \sim \frac{\lambda}{D} \sim \frac{10^{-6}}{10^{-1}} \sim 10^{-5}\ \text{rad}$$

$$\frac{2\,mrad}{10^{-5}\,rad} = 200 \times D.L.$$

Fig. 19. Conceptual free-space optical link for picocell backbone.

This is plotted in Fig. 17 for the calculated variances along 100 m and 1 km horizontal paths. For the 100 m case of the picocell backbone, the effect is nearly zero, and the distribution of intensity is Gaussian. For the 1 km case, the path is highly scintillated, and a typical threshold detection scheme would require nearly 25 dB additional margin of transmitted signal power. Beam wander due to changes in index of refraction is about 0.1 μrad, negligible compared to the 2 mrad beam divergence for the picocell backbone system considered here.

A final issue to consider is misalignment of the transceivers to each other due to wind or thermal expansion causing an angular displacement of either end of the link. Fig. 18 shows the calculated deflection of 10 in diameter hollow poles made from steel, aluminum, and wood, applied wind pressure of 15 lb/in^2, considered hurricane force, and no allowance for aerodynamic flow. For steel and aluminum, the deflection angle is not significant compared with 2 mrad for pole lengths near 40 ft. Since we envision much shorter poles, a much smaller diameter pipe may be used. Thermal expansion of a three–four story building over temperature ranges of 50°C causes fewer than 50 μrad of angular displacement, a negligible effect for our case.

Given the previously discussed issues, a potential conceptual design for a picocell backbone optical communications transceiver operating at the OC-3 data rate (or higher), a range of 100 m or less, all weather operation, and BER of less than 10^{-9} is shown in Fig. 19. Since the beam divergences used in this unit are over two orders of magnitude above the diffraction limit, extremely inexpensive plastic optics can be used throughout, and alignment tolerances become very forgiving, permitting inexpensive manufacture. The aperture size is 4 in, making the unit relatively small, and could be made a factor of 2 smaller without affecting performance significantly. The smaller the transceiver can be made, the easier and less expensive it is to build. Another benefit of small size and wide beam divergence is that relatively inexpensive acquisition and alignment systems can be used for potential autonomous alignment and alignment maintenance systems.

Turning now to the radio link, let us assume an available spectrum of 10 MHz, equally split between tiers 1 and 2 (tier 3, created by LEO satellites, uses a different spectrum, and will not be considered). In tier 1, since the cell diameter is so small, the path loss is low, and we can safely consider the use of some higher level digital modulation, say 16-QAM. Allowing a signaling rate of 1 symbol/s/Hz, this translates to a 20 Mbits/s peak rate/picocell. We further envision the use of a smart array processor at the base station which, coupled with admission control, allows, on average, 10 Mbits/s/picocell, that is, a peak rate of 20 Mbits/s and an effective average frequency reuse factor of 50%.

For the standard cell, more array elements would be used to provide 100% frequency reuse, but the greater path loss limits the modulation to 4ϕ-PSK, and both the peak and average rate within a standard cell are 10 Mbits/s.

VI. Conclusion

In this paper, we have presented the framework for a general-purpose wireless local access network. The use of very small radio cells with frequency reuse at the lowest tier ensures very high bandwidth per subscriber. The use of a dense mesh of short, free-space optical links in a multibox arrangement allows signals to flow between picocellular base stations and an end office, thereby eliminating any dependence on copper and/or optical cabling to the subscriber. Finally, the use of higher tiered cells ensures universal service availability (both from a radio coverage perspective and from a failure recovery perspective). The proposed system can equally well serve fixed-point, pedestrian, and

vehicular users; all access is via a high-capacity radio link, and handoff among cells is fully supported.

Although the use of free-space optical links might suggest a somewhat less than robust service availability, our sample link calculations show that the shortness of the link provides excellent immunity against even extreme weather-related impairments. (For longer spans, millimeter-wave radio links might be considered, but this might substantially escalate the cost of a subscriber's premises-based equipment.) Moreover, the richness of the optical mesh, the high data rate of each link, and the existence of higher tiered radio cells combine to provide excellent immunity against outage caused by equipment failure. Routing in the mesh can be adapted to prevailing traffic conditions such that optical link "bottlenecks" as the traffic flows toward the head end are avoided, and the approach is compatible with embedded virtual connection trees to facilitate hand off while maintaining QoS guarantees for each of several traffic classes (multimedia traffic). Exemplary capacity calculations suggest that a 20 Mbits/s peak rate and a 10 Mbits/s average rate can be provided to each subscriber, and no optical link need operate at a rate exceeding 155 Mbits/s.

Clearly, much additional work remains to "prove in" this approach, and even more work is needed to provide the knowledge base needed to enhance and optimize any system design. It is hoped that this paper will serve to stimulate interest in this approach, and we invite others who may be so motivated to join us in identifying and pursuing the various research issues that should justifiably be addressed.

REFERENCES

[1] D. Wright, *Broadband: Business Services, Technologies, and Strategic Impact.* Boston: Artech House, 1993.

[2] A. S. Acampora, *An Introduction to Broadband Networks.* New York: Plenum, 1994.

[3] *IEEE Commun. Mag.* (*Special Issue on Introducing the Internet Technology Series*), vol. 35, Jan. 1997.

[4] T. S. Rappaport, *Wireless Communications Principles and Practice.* Englewood Cliffs, NJ: Prentice Hall, 1996.

[5] *IEEE Personal Commun.* (*Special Issue on Wireless ATM*), vol. 3, Aug. 1996.

[6] M. de Prycker, *Asynchronous Transfer Mode.* West Sussex, England: Ellis Horwood, 1992.

[7] A. S. Acampora and M. Naghshineh, "An architecture and methodology for mobile-executed hand-off in cellular ATM networks," *IEEE J. Select. Areas Commun.*, vol. 12, Oct. 1994.

[8] ———, "Control and QoS provisioning in high speed microcellular networks," *IEEE Personal Commun.*, vol. 1, Second Quarter 1994.

[9] M. Naghshineh and A. S. Acampora, "Design and control of micro-cellular networks with QoS provisioning for real time traffic," *J. High Speed Networks*, vol. 5, no. 1, 1996.

[10] ———, "Design and control off micro-cellular networks with QoS provisioning for nonreal time traffic," *IEEE/ACM Trans. Networking*, submitted for publication.

[11] A. S. Acampora, "A multichannel multihop local lightwave network," in *Conf. Rec. IEEE GLOBECOM*, Tokyo, Japan, 1987.

[12] C. A. Brackett and A. S. Acampora *et al.*, "A scalable multiwavelength multihop optical network," *J. Lightwave Technol.*, vol. 11, May/June 1993.

[13] A. S. Acampora, "The scalable lightwave network," *IEEE Commun. Mag.*, vol. 32, Dec. 1994.

[14] ———, "Architectures for hardware and software scalable multowavelength networks," in *Photonics Networks.* Norwell, MA: Kluwer Academic, 1997.

[15] S. G. Lambert and W. L. Casey, *Laser Communications in Space.* Norwood, MA: Artech House, 1995.

[16] American National Standard for the Safe Use of Lasers, Laser Institute of Americ, Orlando, FL, 1986.

[17] W. G. Driscoll and W. Vaughan, *Handbook of Optics.* New York: McGraw-Hill, 1978.

[18] W. K. Pratt, *Laser Communications Systems.* New York: Wiley, 1969.

Anthony Acampora (S'68–M'68–SM'86–F'88), for a photograph and biography, see this issue, p. 818.

Scott H. Bloom (M'97) was born in Brooklyn, NY. He received the B.S. degree from the California Institute of Technology in 1983, the Ph.D. degree in physics from Tufts University in 1989, and the M.B.A. degree from the Anderson Graduate School at UCLA in 1995.

Since 1989, he has researched, designed, and built several free-space optical communications systems for space, aircraft, and terrestrial applications at ThermoTrex Corporation, Trex Communications, AstroTerra Corporation, and AirFiber Inc. These systems ranged in performance from the 150 km range, 1 Gbit/s data rate with full tracking and acquisition, to handheld binoculars with a 5 km range and 100 kbit/s data rate.

Srikanth Krishnamurthy was born in Mysore, India, in 1969. He received the M.Sc. (Hons.) degree in physics and the B.E. (Hons.) degree in electrical and electronics engineering with distinction from Birla Institute of Technology and Science, Pilani, India, in 1992, the Master of Applied Science degree in electrical and computer engineering from Concordia University, Montreal, Canada, in 1994, and the 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, NY. Currently, he is a Research Staff Member at the Information Sciences Laboratory, HRL Laboratories, Malibu, CA. His research interests are in wireless and satellite networks, ATM and IP technologies, optical networks, algorithms, queueing theory, and stochastic processes.