First published in *IEEE Personal Communications Magazine*
Second Quarter 1994, Volume 1, Number 2

Avoiding the bottlenecks caused by centralized control functions

Control and Quality-of-Service Provisioning in High-Speed Microcellular Networks

Anthony S. Acampora and Mahmoud Naghshineh

Design and implementation of broadband networks is one of the major focal areas in modern telecommunications. With recent developments in the field of wireless, hand-held terminals, as well as in personal communications services (PCS) [1-6], integration of mobile, wireless connections in a backbone broadband network is an essential and challenging task since mobile users may need to access the communication services offered by the fixed broadband network. This implies that wireless networks must provide packet-based transport and bandwidth-upon-demand, as well as support multimedia applications. Since the radio spectrum is limited, future wireless systems will have micro/picocellular architectures in order to provide the higher capacity needed to support broadband services [7-9]. Due to the small coverage area of micro/picocells and characteristics of the multipath and shadow fading radio environment, hand-off events in future microcellular systems will occur at a much higher rate as compared to today’s macrocellular systems, and control of such systems will introduce a new set of challenges.

We can view wireless/mobile connections as consisting of paths (or routes) through the broadband backbone network; and radio links between the mobile, wireless terminals and base stations (or access points) which are the interface of mobile users to the fixed backbone network. When the quality of a radio link between a wireless terminal and its access point degrades, a new access point with acceptable quality must be found (hand-off), and network control functions of both the fixed and wireless network need to be invoked. In the backbone network, hand-off requires the establishment of a new route, which transports the packets destined to (or originated from) the wireless terminal to (or from) the new access point. Here, network call processing functions need to be invoked in order to set up such a route and ensure that the newly established route maintains acceptable quality-of-service (QOS) to both the wireless connection and to pre-existing calls sharing links of the new route. Furthermore, to execute hand-off, the network call controller must first ensure that the new wireless connection does not overload the new access point and then create a radio link between the mobile terminal and the new access point.

As one can see, a substantial number of call processing and control functions of the fixed and wireless network must be invoked to complete a hand-off event. If such control functions are performed in a centralized fashion, call processing of hand-off events would impose a bottleneck on the capacity of future microcellular networks. In this article, we propose and study distributed control methodologies for high-speed microcellular networks based on a hierarchical grouping of backbone and wireless network resources. With our approach, a number of adjacent cells are grouped into a cell-cluster that is used for call setup and control of the radio links, and all access points in a cell-cluster belong to the same backbone network connection tree, to be used for call setup and control of the backbone portion of wireless connections.

Distributed Control of
The Backbone Network

Since, in a micro/picocellular network, the interface of a wireless connection to its backbone changes frequently as a result of hand-off events, network routing functions must frequently be invoked. The virtual connection tree concept as described in [10-12] reduces the call setup and routing load on the network call processor in such a way that a large number of mobile/wireless connections can be supported. A virtual connection tree is a collection of cellular base stations and wired network switching nodes and links. The root of the tree is a fixed switching node of the wired network and the leaves of the tree are mobile access points or base stations. For each mobile connection, the connection tree provides a set of virtual connections (in each direction), each providing a path from the root to one leaf. To complete a mobile connection in the backbone network, a fixed virtual connection is created for that connection from the tree's root node back to either a wired network port (if the connection is to a fixed port) or to the root of some other connection tree (if the connection is to another mobile user).

Figure 1 illustrates a virtual connection tree in a high-speed packet-switched network that supports mobile cellular radio connections.

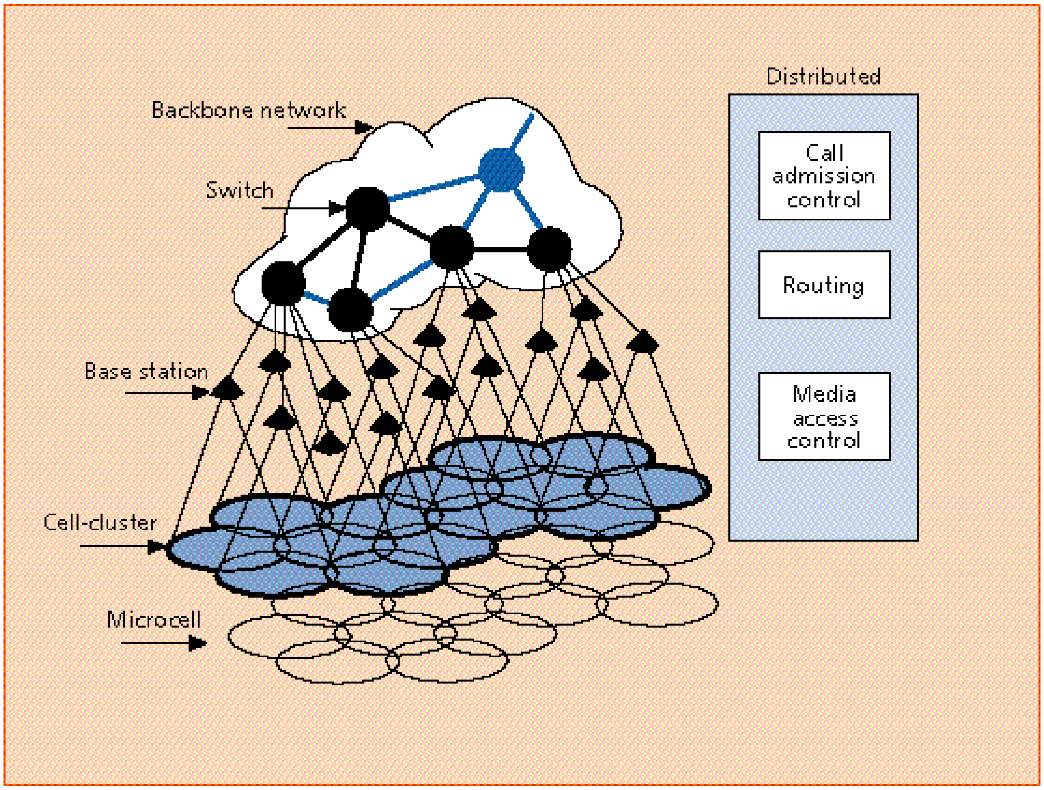

Figure 1. A virtual connection tree.

At the time a mobile connection is admitted to a connection tree, the call setup procedure is executed in two steps. First, the fixed portion of the virtual connection is established between the root of the tree and the appropriate fixed point of the wired network (the fixed user terminus or the root of a destination tree, as described above). This fixed portion is maintained as long as the mobile stays within the connection tree. Second, within the connection tree, two sets of connection numbers are assigned to that mobile connection (one in each direction) with one member of each set used to define a path from the root to one of the leaves, and the routing tables of the switches within the connection tree are appropriately updated to include the new connection numbers. At any given time, only the two virtual connections (one in each direction) between the access point chosen by the mobile and the root of the virtual connection tree are actually in use as a result of the mobile's decision to choose that access point as evidenced by the routing information included in its packet headers (the virtual connection tree is similar to a point-to-multipoint connection, except that only one path of the tree is in use at any given time).

When a mobile user already admitted to a virtual connection tree wishes to hand-off to another base station in the same virtual connection tree, it simply begins to transmit packets with the connection number assigned for use between itself and the new base station. Using the pre-established path between the new base station at the root of the tree, that mobile's packets will flow to the root, across the fixed portion of the network, and to their ultimate destination. In this way, the call processor is not involved in the hand-over. In the reverse direction, the first packet to arrive at the root from a given mobile connection that bears a new connection number is properly interpreted (in hardware) as evidence of a hand-over. Using its knowledge of the connection number assigned for that mobile connection from the root of the tree to the new mobile access point, the routing table at the root switch is appropriately updated so that packets flowing from the root to that mobile receive the connection number appropriate to the deliverance of those packets to that new mobile access point. An example for the connection tree setup and routing protocol is given in Fig. 2.

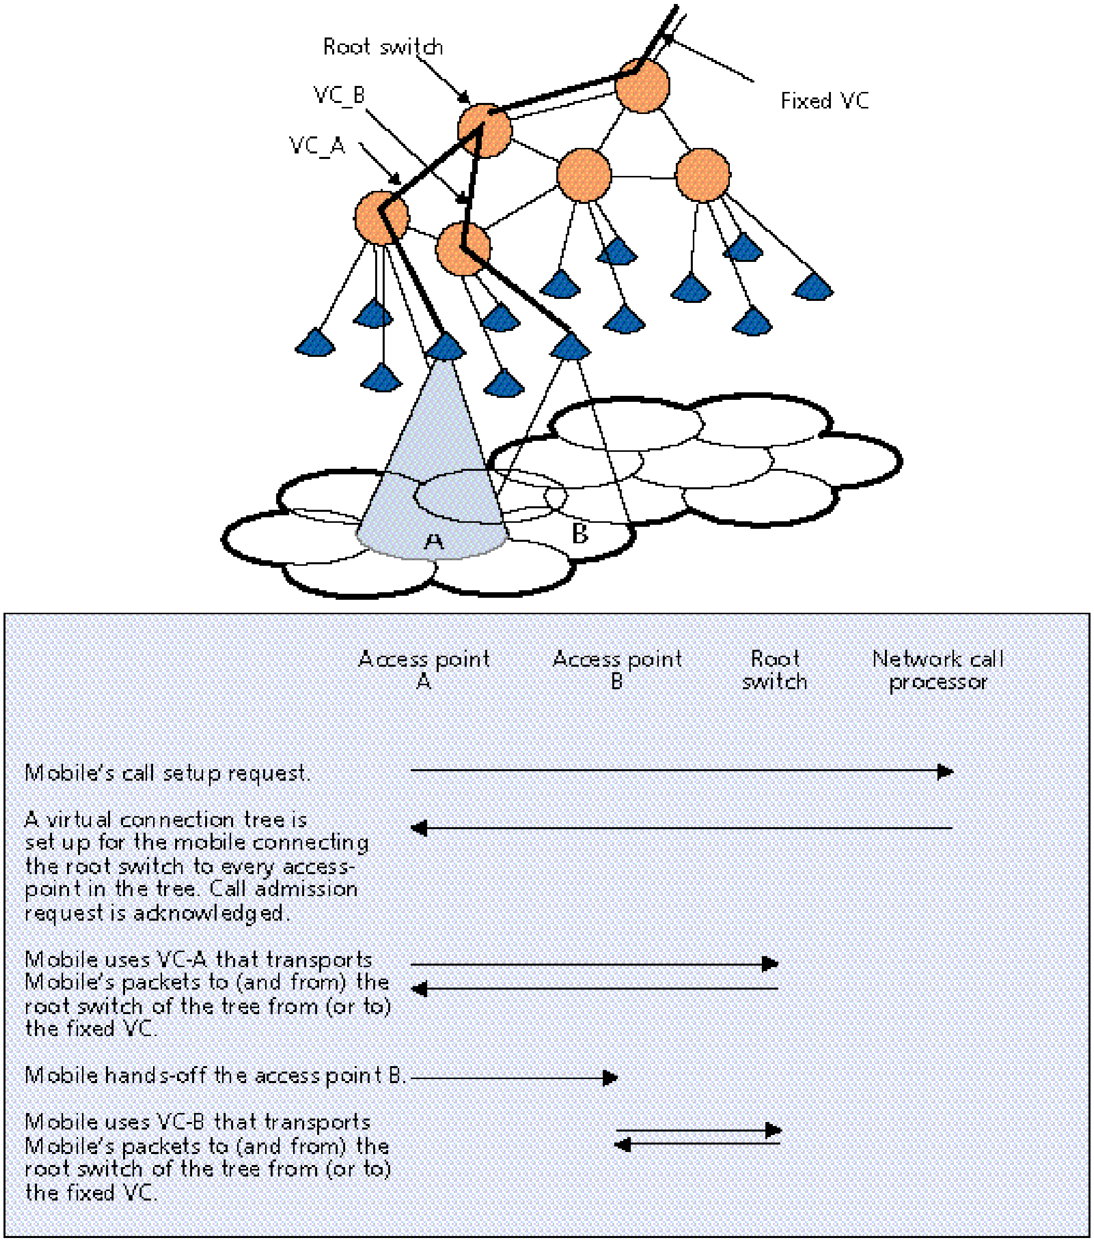

Mobile's call setup request.

A virtual connection tree is set up for the mobile connecting the root switch to every access-point in the tree. Call admission request is acknowledged.

Mobile uses VC-A that transports Mobile's packets to (and from) the root switch of the tree from (or to) the fixed VC.

Mobile hands-off the access point B.

Mobile uses VC-B that transports Mobile's packets to (and from) the root switch of the tree from (or to) the fixed VC.

Figure 2. An example for connection tree set up and operation.

Whenever the mobile connection reaches the boundary of a connection tree, it seeks admission to a new connection tree. We refer to this procedure as the *virtual connection tree hand-off*. At this point, the network call processor must again become involved. However, since the geographical coverage of a virtual connection tree is large compared to the size of the radio cells which comprise the tree, the rate of connection tree hand-over is assumed to be acceptably low and manageable by the call processor. To prevent the connection of a mobile situated at the boundary of two connection trees from oscillating between the two, connection trees would overlap in space (i.e., some base stations might belong to two trees) such that, as a mobile approaches the geographical boundaries of its current tree, it hands over to the new tree and appears in the new tree safely within the new tree's interior. Thus, it is highly unlikely that the connection will immediately seek to again hand-over to yet another tree.

In contrast to today’s fixed networks, where the user-network-interface remains unchanged throughout the connection life-time, backbone network resources allocated to a wireless connection change very frequently as a result of hand-off from one access point to another. Hence, the problem of optimal resource allocation to wireless connections is inherently different from that of resource allocation to static connections. Resource allocation and admission control of the backbone portion of a wireless, mobile connection can also be effected through use of the virtual connection tree. Each virtual connection tree consists of a number of virtual connections (or routes) between the root of the tree and access points, and only one pair of these virtual connections is used by a wireless connection at any given time. Thus, by considering the probability that a mobile connection utilizes a specific path of the connection tree, the time distribution of that path, and the connection traffic parameters, existing call admission methodologies (as discussed in [14-19] for example) can be adopted such that an adequate QOS is provided to a mobile connection as long as it remains within the area covered by a virtual connection tree. The portion of a mobile connection from the root of its connection tree to a fixed destination (or root of another connection tree) remains static, and hence existing call admission control methodologies can be applied over this portion.

Distributed Control of Wireless Resources

As far as the wireless spectrum is concerned, overload conditions might occur if the communication needs of a number of wireless terminals populating a small area exceed the total capacity of all access points within their reach. We refer to this situation as a radio congestion state. Depending on the load of the wireless network and the manner in which the radio resource is shared among mobile users (fixed channel allocation, dynamic channel allocation, etc.), a radio congestion state might be encountered at hand-off, an event that will result in either the termination of the connection, large delays, and/or packet loss. Thus, in the new paradigm of high-speed wireless networks, call control functions are required to keep the probability of a congested radio state suitably low. In addition, these wireless call control functions need to be simple enough that a high rate of mobile connection hand-offs can be managed in a short time, such that the hand-offs are not noticeable to the mobile users. Moreover, call control functions must be based on the class of each wireless connection, and must differentiate between new calls and hand-off calls. In the past, several solutions were studied to reduce the radio congestion probability in cellular networks. Some of these solutions are based on power control [20], dynamic or hybrid channel allocation schemes [21, 22], and layered architectures [23, 24], which try to minimize the congestion probability while providing high radio spectrum utilization efficiency. The above approaches do not impose any explicit control on the admission of new calls, and as the load of the cellular network increases, the radio congestion probability of such systems becomes unacceptably high. There exist other solutions that treat new calls and hand-off calls differently by reserving a number of guard channels at each base stations for hand-off calls [25], or by giving such calls service priority over the new calls [26]. These solutions lower the congestion probability of hand-off calls, but use only local state information (number of calls in the cell where a new call is initiated) for accepting a new call as opposed to the global state information (number of calls in the neighborhood of the cell where a new call is initiated). In [27], it is shown that by taking the global state information (as opposed to local state information) the wireless resources can be utilized more efficiently. Most importantly, the above-referenced studies consider only one class of wireless traffic.

Class-Based QOS in Wireless Networks

We consider three classes of wireless connections, differentiated on the basis of the action initiated when a radio congestion state is encountered. They are: 1) real-time connections, 2) connection-oriented data connections, and 3) message-oriented, delay-insensitive traffic. For real-time or class I connections such as voice or video, the connection must be dropped if the mobile moves into a congested area where no wireless channel is available. Hence, we define the QOS metric of a class I connection to be the hand-off dropping probability. Due to their strict delay requirements, class I connections are given priority over the other two classes of traffic.

Class II calls are data connections that support applications requiring reliable transport. Transmission Control Protocol (TCP) might typify such connections. In contrast to class I connections where the encountering of a congested radio state results in the termination of the call, class II connections can be "put on hold" under such conditions. We refer to this situation as an overload state wherein the total wireless capacity available in an area is smaller that the total instantaneous capacity required by wireless terminals in that area. A class II connection that is in an overload state will suffer from packet delay (or loss) much higher than normal conditions. As a result, the end-to-end connection control entity (e.g., transport protocol) will observe the congestion and will reduce or stop the flow of the packets in the network. Once the congestion (or overload) state is terminated, the normal flow of packets in the network is resumed. Thus, we define the QOS metric of a class II connection to be the probability of being in such an overload state. The duration of the congestion state is also another QOS metric that may be important.

Applications such as paging, which rely on the network's best-effort delivery of single messages and which are delay tolerable, use class III service. Class III traffic utilizes the extra capacity of the wireless network not consumed by class I and II connections. Packets of class III traffic are stored in the mobile terminal until wireless channel resources (which are not serving class I or class II connections) are available. Hence, we use the average queuing delay of class III packets as their QOS metric. The attributes and QOS metrics of the three classes of wireless traffic are summarized in Table 1.

Table 1. Attributes and QOS matrics for different classes of wireless traffic.

Regarding wireless resources, the purpose of call admission control is to limit the number of in-progress wireless calls such that, once a wireless call is admitted, the probability of its encountering a congested radio state is acceptably low as to provide the required QOS. This is done by blocking new wireless call setup requests when the number of existing calls has reached this limit. Thus, the new-call blocking probability is another QOS metric that needs to be considered for all classes of wireless traffic.

Class-Based Wireless Call Admission

Our approach to class-based wireless call admission is based on the cell-cluster concept. A cell-cluster is a group of adjacent cells which belong to the same connection tree. Whenever a new mobile connection seeks admission to a cell-cluster, the cell-cluster call controller will admit (or reject) the call based on the class of the mobile connection and the number of wireless connections of each class already admitted to that cell-cluster. By taking into account bandwidth requirements, hand-off rate, and call duration time statistics, the call admission decision can then be made, e.g., to guarantee the QOS metrics. It is important to note that once a wireless connection is admitted to a cell-cluster, it can freely hand-off from one base station to another without the involvement of the call controller, and that the probability of encountering a radio congestion state is limited to a guaranteed pre-defined level for each class of wireless connection.

Based on the total mobile connection load in a given geographical area (such as the area covered by a number of neighboring cell-clusters), the number of base stations and the coverage area of each need to be designed in such a way that under normal conditions, the new-call blocking probability is suitably low. However, it is important to guarantee that once a mobile connection is admitted to a cell-cluster, it will enjoy its prescribed QOS during its connection life time. The following example should clarify these concepts. Let us consider a homogeneous system with fixed-channel allocation supporting only class I calls, wherein each base station can support up to 20 calls. The average call duration is 0.5 units of time, and the average time spent communicating with any particular base station before handing-off to another is 0.1 units of time. Moreover, let us perform call admission control based on cell-clusters consisting of 20 cells (or base stations), each of which can support up to 20 real-time wireless connections (C = 20). New-call requests are rejected if 320 calls are admitted to the cell-cluster.

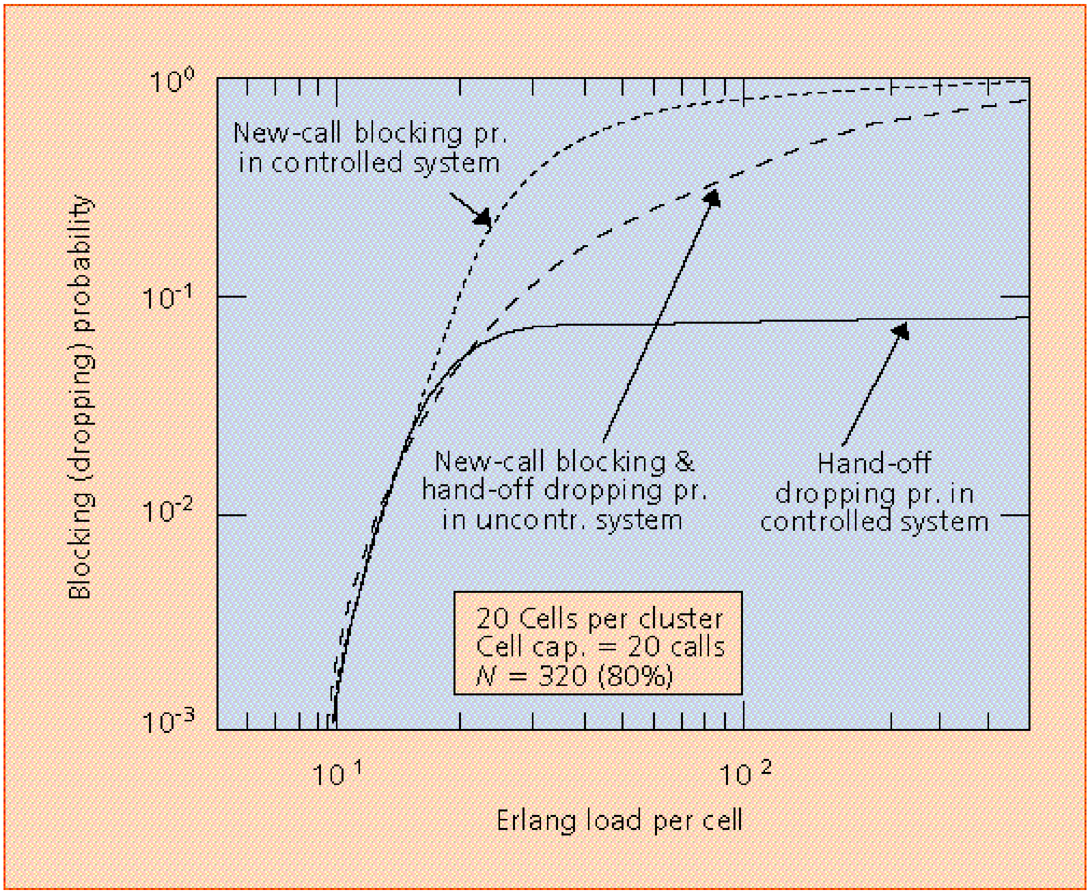

Figure 3. The QOS provided by performing wireless call admission control compared to an uncontrolled system.

In Fig. 3 we compare the QOS provided by a controlled system to the QOS provided by an uncontrolled system. The horizontal axis represents the Erlang load of new calls per cell, and the vertical axis represents the QOS in terms of new-call blocking or hand-off dropping probability. The dashed line represent the hand-off dropping probability (which is very close to the new-call blocking probability) in an uncontrolled system. The solid line represents the hand-off dropping probability, and the dotted line represents the new-call blocking probability in the controlled system. It is important to note that cluster-based wireless call admission control does not noticeably reduce the utilization efficiency of the base stations under normal conditions, since the new-call blocking probability of the uncontrolled system is very close to the new-call blocking probability of the controlled system. Only under conditions of heavy load does the control mechanism increase the blocking of new calls so that with high probability the wireless spectrum remains available to serve already-admitted calls, thereby limiting the hand-off dropping probability to a prescribed maximum level.

An important measure of performance is the utilization efficiency of scarce wireless resources when the cell-cluster-based call admission control strategy is used to maintain QOS guarantees. As an example, consider a system supporting only class II calls. Here, we define utilization efficiency R as the number of calls N that may be admitted to a cell-cluster subject to some QOS guarantee, normalized by the "raw" capacity of the cell-cluster containing B base stations, each capable of handling C connections ($R = N/(BC)$). This is plotted in Fig. 4 as a function of C for a cell-cluster containing $B = 8$ base stations and a guaranteed overload probability of 1 percent. Also plotted in Fig. 4 is the expected overload period (normalized by the mean time between hand-off events for any mobile) for the same system and guaranteed service quality. We note that as the capacity per base station increases, cell-cluster utilization efficiency improves and the mean overload period diminishes, implying that for large C, the instantaneous connection load is approximately balanced among the base stations. Results such as there clearly demonstrate the benefits statistical multiplexing.

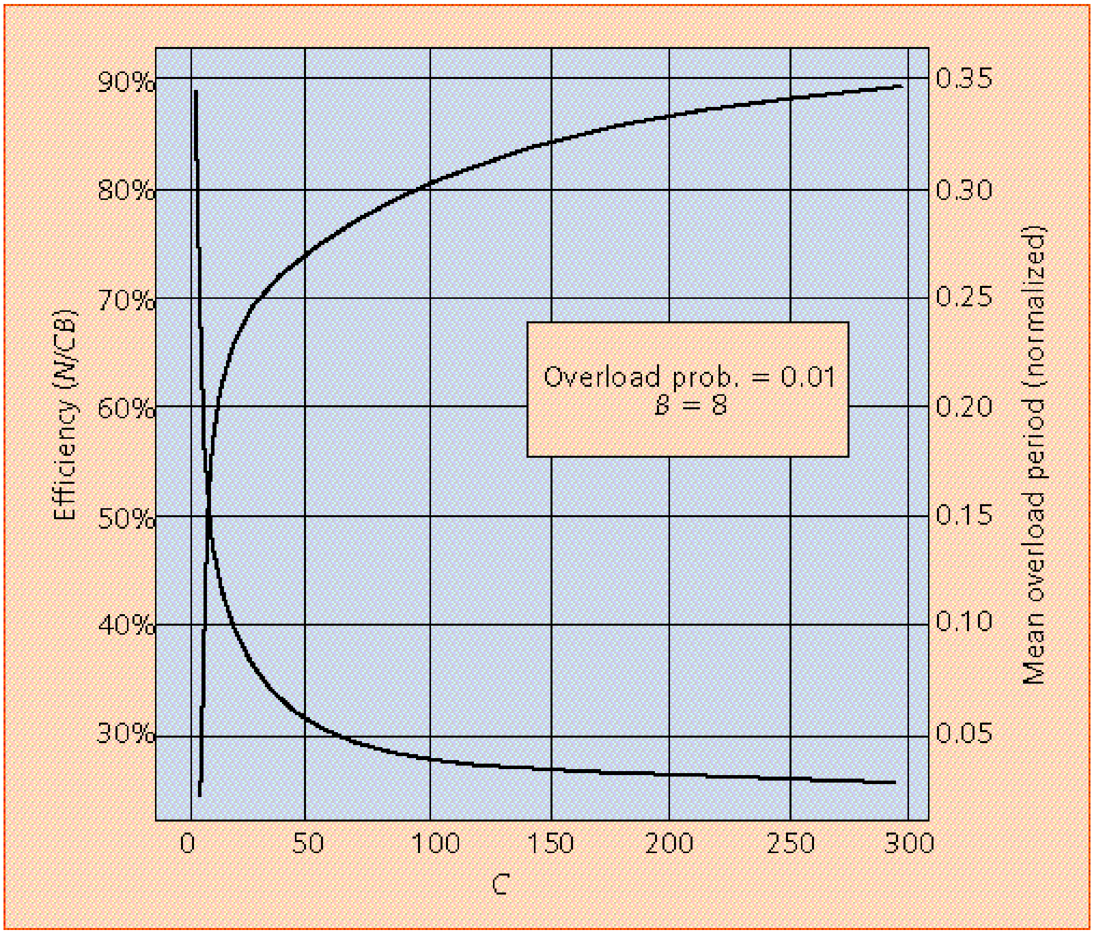

Figure 4. Utilization efficiency and mean overload period in a cell-cluster containing eight base stations and overload probability equal to 1 percent.

These results can be extended to multiple classes of wireless traffic [28]. For example, depending on the call characteristics and desired QOS for each class of traffic, the resulting call admission region might appear as shown in Fig. 5. Whenever a wireless connection seeks admission to the cell-cluster, it is admitted if the system is operating within the call admission region. This region can be calculated and saved in tables so that a computationally intensive task does not need to be performed "on the fly" for every call admission request.

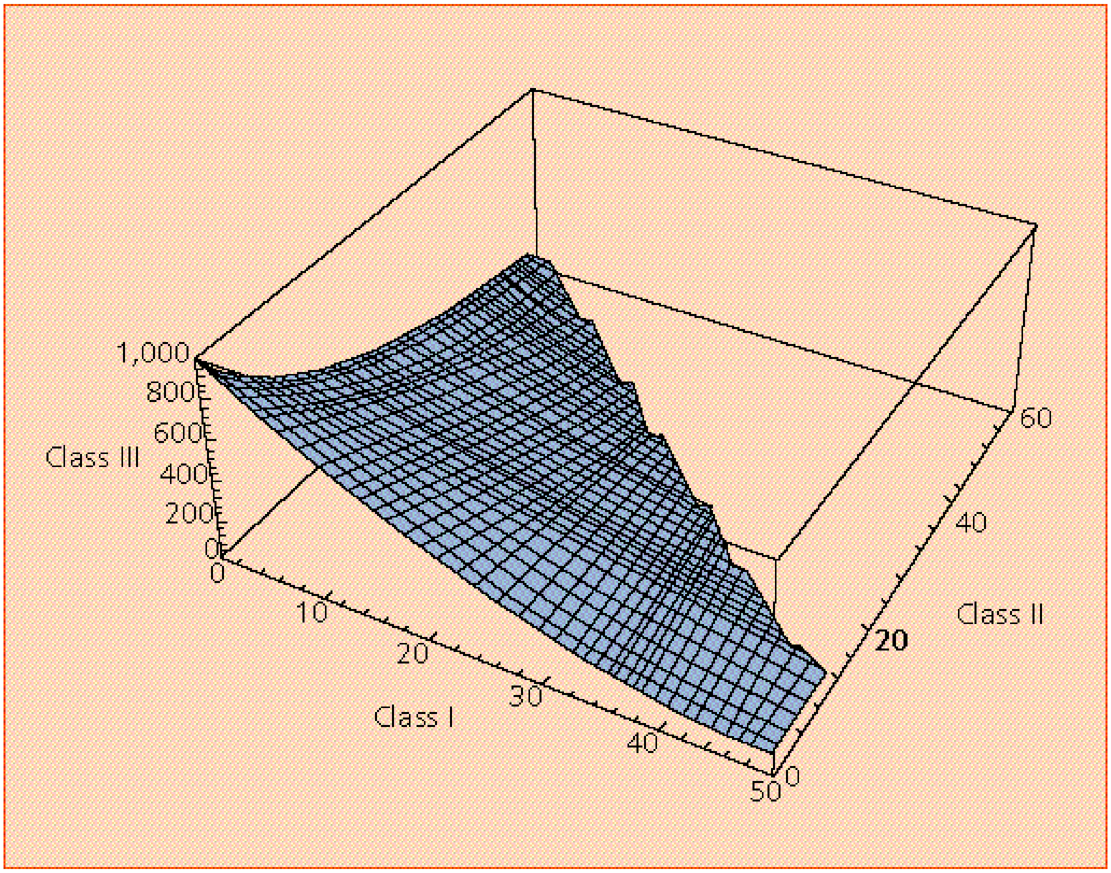

Figure 5. Multi-class wireless call admission region.

Reliable Radio Link and its Implications on QOS

QOS provisioning and efficient control of wireless connections, as discussed above, can be used to ensure that network resources are available with sufficiently high probability to enable mobile hands-off from one access point to another. Another important factor in providing an adequate QOS to a wireless connection is the provisioning of a highly reliable radio link between the wireless terminal and its access point. This is an essential and challenging task in high-speed micro/picocellular networks due to the presence of dispersive multipath fading, which can result in unacceptable high link outage. With the continuing decrease in the cost of digital signal processing hardware, antenna diversity techniques can be deployed in wireless systems to combat multipath fading and also to suppress interference, thereby increasing overall system capacity. In fact, a system with M antennas can achieve an M-fold increase in capacity (compared to systems without spatial diversity) [30] which will further increase the statistical multiplexing efficiency of the wireless resources. We envision a microcellular system in which the space diversity is achieved with an antenna array at the base station and only a single antenna at the mobile, thereby allowing the cost to be shared among many users.

In the following, we describe a radio system organization and signaling scheme for microcellular wireless networks that deploy spatial diversity techniques in conjunction with virtual connection tree and cell-cluster concepts for distributed control of network resources. Figure 6 shows a timing diagram of the various signals involved. Each transmission frame is divided into three sections: the frame marker field, the signaling field, and the communication field. All transmission frames are of equal length, and their boundaries are delimited by frame marker fields. Time-duplexed transmission is used, i.e., the same channel is time shared for base-to-mobile and mobile-to-base communications. The modified polling media access scheme described in [31] is used within the communication field, whereby each base station regularly polls each mobile for which it is currently responsible, to obtain from each information on instantaneous communication needs. Replies to the polling also provide information needed to adapt the base station array to each mobile, thereby maintaining low link outage. The base station then schedules transmission time for mobile-to-base and base-to-mobile communication. Permission to transmit is sequentially conferred to each mobile by means of a permission token.

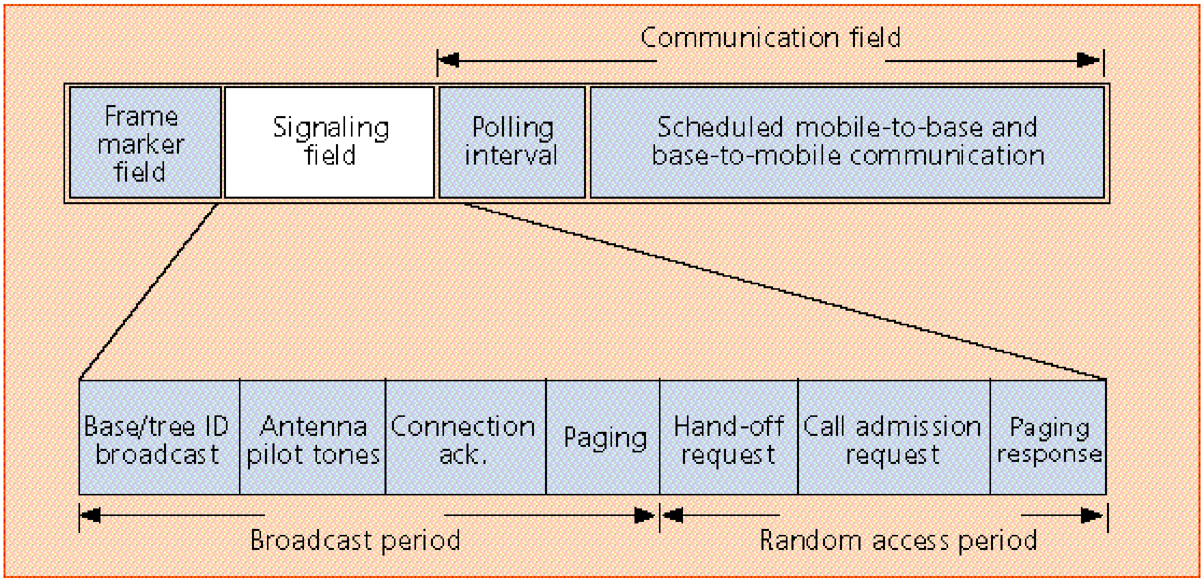

Figure 6. Overall framework for operation of the Virtual Connection Tree.

Referring again to Fig. 6, the communication field is used concurrently by all radio cells in the connection tree, with the various base stations managing the communication flow for their mobiles. The frame marker and signaling fields are used sequentially in sequential frames by the radio cells in the connection tree. The signaling field of the transmission frame is used primarily to admit a given mobile to the polling sequence of a given base station, and the frame marker field is a unique word periodically (and sequentially) broadcast from the base stations and is used by the mobiles to establish a basic time reference for start-of-frame. All mobiles that are in either an active state of communication or in a standby state available to accept calls maintain continuous frame marker synchronization. Upon powering up, a mobile will lock onto the frame marker broadcast from one of the base stations, thereby establishing its basic timing reference.

The signaling field is sub-divided into a broadcast period (base-to-mobile) and a random access period (mobile-to-base), neither of which contain any communication signals. The first segment of the broadcast field is used to uniquely identify the cell site associated with that frame, as indicated by the base station ID, along with the connection tree to which that base station belongs. This is followed by pilot tones transmitted by the various antenna elements of that base station’s array. These pilot signals, broadcast in successive frames with only one base station sending its pilots in any one frame, are processed by each mobile in the connection tree in order to determine those base stations to which a satisfactory channel can be established. By continuously monitoring these pilot signals, each mobile can decide when it wishes to initiate a hand-off (when the channel to its current base station becomes unusable as a result of user mobility and/or changing propagation environment), and to which base station it wishes to hand-off (the one to which the strongest channel can be established, any one chosen randomly from the list of those to which a satisfactory channel can be established, etc.).

Moving now to the Random Access period of the signaling field (we shall return later to the broadcast period), a mobile can request to join a given base station’s polling sequence by signaling its request during the hand-off request segment of the frame associated with that base station (remember that the frame marker and signaling fields of each frame are associated with one and only one base station). Since the hand-off request field is randomly accessed (Aloha, Aloha with capture, etc.) successful receipt is acknowledged by the base station during the connection acknowledgment of some later frame’s signaling field broadcast period.

A mobile can also request to be admitted to a new virtual connection tree and its cell-cluster (a request that must be processed by the connection tree’s admission controller) by randomly accessing the Call Admission Request segment of the desired base station’s signaling field; successful receipt of request and decision to admit are in the connection acknowledgment field of some later frame. Finally, a call may be placed to a given user by means of the paging segment of the signaling field, with the called mobile’s ID being successively broadcast by all base stations in the connection tree. When the mobile hears its page, it replies with a paging response in the signaling field of a frame associated with a base station to which a satisfactory channel can be established. Since the paging response is randomly accessed, paging continues until a paging response is produced or a time-out interval elapses (the mobile may not be located within that connection tree, or the mobile may not be in standby mode).

Conclusion

QOS provisioning in micro/picocellular networks necessitates distributed control of the backbone as well as wireless network resources. These resources need to be allocated in such a way that when a wireless terminal hands-off from one access point to another, network resources are available with an acceptably high probability. In addition, these control function need to be simple enough such that a high rate of hand-off events can be accommodated. A virtual connection tree in the backbone network, along with its associated radio cell-cluster, can be viewed as a single collective resource of the network that can be managed for efficient rerouting and call admission control of wireless connections. As we have shown, the cell-cluster-based call admission control strategy provides a guaranteed QOS as defined for each class of wireless connection. Moreover, provisioning of a highly reliable radio link is an indispensable requirement for future high-speed microcellular systems. Spatial antenna diversity techniques have shown to be quite useful for maintaining a high radio link availability in the harsh microcellular radio environment. Therefore, the combination of base station diversity techniques to maintain a reliable high-speed radio link, and simple distributed control methodologies to manage rapid cell hand-offs may emerge as key elements of the future high capacity microcellular wireless access networks.

References

[1] D. Cox, Personal Communications-A Viewpoint, *IEEE Commun. Mag.*, vol. 28, no. 11, Nov. 1990.

[2] D. Cox, Wireless Network Access for Personal Communications, *IEEE Commun. Mag.*, vol. 31, no. 12, Dec. 1992.

[3] D. Cox, A Radio System Proposal for Widespread Low-Power Tetherless Communications, *IEEE Trans. on Commun.*, vol. 29, no. 2, Feb. 1991.

[4] D. Goodman, Cellular Packet Communications, *IEEE Trans. on Commun.*, vol. 28, no. 8, Aug. 1990.

[5] D. Goodman, Trends in Cellular and Cordless Communications, *IEEE Commun. Mag.*, Feb. 1991.

[6] I. Ross, Wireless Network Directions, *IEEE Commun. Mag.*, vol. 29. no. 6, June 1991.

[7] W. C. Y. Lee, Smaller cells for greater performance, *IEEE Commun. Mag.*, vol. 29, no. 11, Nov. 1991.

[8] J. Sarnecki *et al*., Microcell Design Principles, *IEEE Commun. Mag.*, vol. 31, no. 4, April 1993.

[9] L. I. Greenstein *et al*., Microcells in personal communication systems, *IEEE Commun. Mag.*, vol. 30, no. 12, Dec. 1992.

[10] A. S. Acampora and M. Naghshineh, Method and Apparatus for Supporting Mobile Communications in Asynchronous Transfer Mode, US Patent Application, Columbia Ref: M93-009xx.

[11] A. S. Acampora, and M. Naghshineh, Wireless ATM Networks, 6th IEEE LAN/WAN Workshop, San Diego, CA., October 1993.

[12] A. S. Acampora and M. Naghshineh, An Architecture and Methodology for Mobile-Executed Cell Hand-off in Wireless ATM Networks, 1994 Intl. Zurich Seminar on Digital Communications, also to appear in *IEEE JSAC*, Special Issue on Wireless and Mobile High-Speed Communications.

[13] G. Pllini, K. Meier Hellerstern, and D. Goodman, Handover Protocols Between Metropolitan Area Networks, GLOBECOM ’92, Orlando, FL., 1992

[14] A. E. Eckberg, B-ISDN/ATM Traffic and Congestion Control, *IEEE Network*, vol. 6, no. 5, Sept. 1992.

[15] J. Filipiak, M-Architecture: A Structural Model of Traffic Management and Control in Broadband ISDNs, *IEEE Commun. Mag.*, vol. 27, no. 5, May 1989.

[16] G. Gallasi, G. Rigolio, and L. Verri, Resource Management and Dimensioning in ATM Networks, *IEEE Network*, vol. 4, no. 2, May 1990.

[17] D. Hong and T. Suda, Congestion Control and Prevention in ATM Networks, *IEEE Network*,vol. 5, no. 3, July 1991.

[18] J. Y. Hui, Resource Allocation for Broadband Networks, *IEEE JSAC*, vol. 6, no. 9, Dec. 1988.

[19] R. Guerin, H. Ahmadi, and M. Naghshineh, Equivalent Capacity and its Applications to Bandwidth Allocation in High-Speed Networks, *IEEE Sel. Areas in Commun.*, no. ISAC-7, Sep. 1991.

[20] J. Zander, Performance of Optimum Transmitter Power Control in Cellular Radio Systems, *IEEE Trans. on Vehicular Technol.*, vol. 41, no. 1, Feb. 1992.

[21] L. J. Cimini,Jr., G. J. Foschini, C.-L. I, *Distributed Dynamic Channel Allocation for Microcellular Networks*

[22] S. Tekinay and B. Jabbari, Handover and channel channel assignment in mobile cellular networks, *IEEE Commun. Mag.*, vol. 29, no. 11, Nov. 1991.

[23] C.-L. I, L. J. Greenstein, and R. D. Gitlin, A microcell/macrocell cellular architecture for low- and high-mobility wireless users, *IEEE JSAC*, vol. 11, no.6, Aug. 1993.

[24] S. T. S. Chia, Mixed cell architecture and handover, IEE Colloquium on "Mobile communications in the year 2000," no. 139, UK, June 1992.

[25] E. C. Posner and R. Guerin, Traffic policies in cellular radio that minimize blocking of handoff calls, Proc. 11th Teletraffic Cong. (ITC 11), Kyoto, Japan, Sept. 1985.

[26] S.-H. Oh and D.-W. Tcha, Prioritized channel assignment in a cellular radio network, *IEEE Trans. on Commun.*, vol. 40, no. 7, July 1992.

[27] A. S. Acampora and M. Naghshineh, Design and Control of Micro-Cellular Networks with QOS Provisions for Real-Time Traffic, Paper in preparation.

[28] A. S. Acampora and M. Naghshineh, Quality-of-Service Provisions in Micro-Cellular Networks Supporting Multimedia Traffic, Paper in preparation.

[29] A. S. Acampora and J. H. Winters, A Wireless Network for Wide-Band Indoor Communications, *IEEE JSAC*, vol. 5, no. 5, June 1987.

[30] J. H. Winters, Optimum Combining for Indoor Radio Systems with Multiple Users, *IEEE Trans. on Commun.*, vol. 35, no. 11, Nov. 1987.

[31] Z. Zhang and A. Acampora, Performance of a Modified Polling Strategy for Broadband Wireless LANs in a Harsh Fading Environment, GLOBECOM’91.

Biographies

Mahmoud Naghshineh [S’87] received a Vordiplom degree in electrical engineering from RWTH Aachen, Germany, in 1985, and B.S. (computer engineering) and M.S. (electrical engineering) degrees from Polytechnic University in 1988 and 1991, respectively. Currently, he is a doctoral candidate at Columbia University’s department of electrical engineering in New York City. He joined IBM in 1988 and he is currently working in the communication networks department of the IBM Thomas J. Watson Research Center, Hawthorne, New York, where he is with the wireless networks architecture and analysis group. He has previously worked on the design, analysis, and control of high-speed packet switched networks, as well as network design tools. His current research interests are in the area of design of network protocols for wireless/mobile computing and communications, media access protocols for PCN, design and control of high-speed microcellular networks, as well as mobile ATM.

Anthony S. Acampora [F ’88] received his Ph.D. in electrical engineering from the Polytechnic Institute of Brooklyn. He is professor of electrical engineering at Columbia University and director of the Center for Telecommunications Research, a national engineering research center in New York City. He joined the faculty at Columbia in 1988 following a 20-year career at AT&T Bell Laboratories, most of which was spent in basic research where his interests included radio and satellite communications, LANs and MANs, packet switching, wireless access systems, and lightwave networks. His most recent position at Bell Labs was director of the Transmission Technology Laboratory, where he was responsible for various projects, including broadband networks, image communications, and digital signal processing. At Columbia, he is involved in research and education programs concerning broadband networks, wireless access networks, network management, optical networks, and multimedia applications. He is a former member of the IEEE Communications Society Board of Governors. He has published more than 140 papers, holds 24 patents, and has authored a recently completed 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*. He sits on numerous telecommunications advisory committees and frequently serves as a consultant to government and industry.