Time-Frequency-Code Slicing: Efficiently Allocating the Communications Spectrum to Multirate Users

Mark J. Karol, *Fellow, IEEE,* Zygmunt J. Haas, *Senior Member, IEEE,*
Clark B. Woodworth, *Member, IEEE,* and Richard D. Gitlin, *Fellow, IEEE*

***Abstract*— The time-frequency-code slicing technique allows multiple users with different data-rate requirements access to a communications resource in a manner that is cost effective over a wide range of access rates. For instance, with a time-frequency slicing (TFS) approach, users are assigned different portions of the frequency spectrum (e.g., on a slot-by-slot basis), granting them access to a fraction of the shared resource that is commensurate with their needs and their own end equipment. Users with high-data-rate requirements can "grab all the bandwidth" when no one else needs it. Also, by efficiently packing the time-frequency space, better system utilization is attained. For the specific case of TFS, we compute the reduction in blocking probability achieved under the constraint of a single transmitter/receiver per user. As an example, consider the case of 70% traffic load with ten frequency bands and 15 time slots per frame. Using the traditional allocation scheme in which users can be assigned only a single-frequency band per time slot, there is a 10% blocking probability for new connections that request 14 "time-frequency slices." The TFS technique reduces this blocking probability to well below 0.01%.**

***Index Terms*— Access protocols, code-division multiaccess, frequency-division multiaccess, GSM, IS-54, time-division multiaccess.**

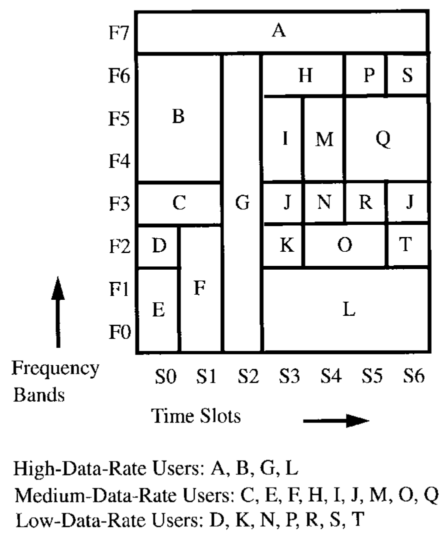

Fig. 1. An example of a time-frequency-sliced system with continuous frequency usage across frequency allocations of the same user.

I. Introduction

Many communication systems (e.g., wireless cellular networks, fiber-optic-based LAN's, satellite systems, and cabled distribution networks) have requirements for: 1) a variety of access rates to support a wide variety of applications; 2) low-cost access for users with low-data-rate requirements; and 3) high spectral efficiency. While time-division multiple-access (TDMA) systems allow both users with high-data-rate requirements and low-data-rate requirements to share the communications bandwidth (e.g., by assigning more time slots per frame to the high-data-rate users), they require high-rate access, even for the users with low-data-rate requirements, which increases their cost and complexity. Alternatively, if the frequency spectrum is partitioned into frequency bands [i.e., frequency-division multiple access (FDMA)], then the maximum bandwidth available to an individual user is limited (unless a user has multiple transmitters that allow it to access several frequencies at the same time)—even if the user desires a large peak bandwidth only for a short period of time.

To support users of "arbitrary" access rates and retain low-cost access for users with low-data-rate requirements, approaches using "field coding" [1] and "universal time slots" [2] have been proposed. In these approaches, users are allowed to transmit at their own desired rate during their assigned time slots. These techniques are spectrally inefficient, but are suitable for optical media, where bandwidth is abundant. However, in some cases, such as radio, the medium is quite precious, and techniques that make efficient use of the transmission spectrum are necessary.

To achieve better spectral efficiency while preserving inexpensive access, we propose dividing the time-frequency-code space into slices, which are allocated to users according to their transmission requirements [3]. This time-frequency-code slicing approach: 1) attains better spectral use than universal-time-slot approaches [2]; 2) supports a variety of access rates; 3) allows users with low-data-rate requirements to keep low-cost endpoints; and 4) requires only a single transmitter/receiver pair per user.

For convenience, we will focus initially on the example given in Fig. 1. It illustrates a time-frequency-sliced system in which users modulate a signal in one or more of the frequency bands on a slot-by-slot basis. A *unit slice,* which

Manuscript received December 15, 1995; revised October 29, 1996.

M. J. Karol, C. B. Woodworth, and R. D. Gitlin are with Bell Labs, Lucent Technologies, Holmdel, NJ 07733 USA.

Z. J. Haas is with the School of Electrical Engineering, Cornell University, Ithaca, NY 14853 USA.

Publisher Item Identifier S 0018-9545(97)05107-4.

consists of one frequency band allocation for one time slot, is the minimum amount of resource available to a user. Users with high-data-rate requirements can modulate a signal that covers several frequency bands. Also, high-data-rate users can "grab all the bandwidth" when no one else needs it. We call these users *high-bandwidth* users in contrast with *low-bandwidth* users that modulate a signal in only one (or a small number of) frequency band(s) per time slot. Note in Fig. 1 that in each time slot, each user's frequency allocations are contiguous (e.g., high-data-rate user B is assigned frequency bands F4, F5, and F6 during time slots S0 and S1), but time-slot allocations need not be contiguous (e.g., user J is assigned frequency band F3 during time slots S3 and S6). Because users can cover several unit slices (in one or both of the time and frequency dimensions)—without wasting the "guard bands" between unit slices—we refer to the proposed technique as time-frequency-code "slicing" rather than time-frequency-code "division."

A key observation we make is that through scheduling, the frequency spectrum can be filled more efficiently than the universal-time-slot approach allows because several low-bandwidth users can be scheduled to transmit on different frequencies in the same time slot. During other time slots, a smaller number of high-bandwidth users (perhaps only one) may be scheduled to transmit.

In a time-frequency-sliced system, a time-frequency allocation can be assigned on a "permanent" basis (for the duration of a connection without reassignment), in a repeating time frame (e.g., for circuit-switched applications), or the time-frequency allocations can be announced, for example, on a slot-by-slot basis for ATM (packet) systems. If the time-frequency allocations are indeed scheduled on a slot-by-slot basis, then a demand-assignment media-access protocol such as DQRUMA [4] can be used to collect users' requests and coordinate their transmissions.

Although Fig. 1 illustrates an example in which high-bandwidth users require their frequency slots to be assigned contiguously, other modulation schemes, for example, multitone [5], may allow the adjacency requirement to be relaxed. Examples are provided in Section II. Other variations on the proposed scheme that include dividing the resource space into a time-code space or a time-frequency-code space are discussed in Section III. After discussing these specific examples of the general time-frequency-code slicing technique, we then, for illustration, compare in Sections IV–VI the performance of time-frequency-sliced systems with the performance of a traditional assignment (TA) scheme. In the TA, time and frequency are divided into *fixed* time and frequency slots, and users can be allocated a single frequency band only in any time slot.

Although we discuss only the performance advantages of time-frequency-sliced systems with contiguous frequency usage (as illustrated in Fig. 1), similar advantages hold for the other variations of sliced systems as well.[1] Specifically, assuming a circuit-switched time-slotted system in which there are F frequency bands and N time slots per (periodic) frame,

[1] Appendix A contains a brief outline of the performance analysis with time-frequency slicing (TFS) and noncontiguous frequency assignments.

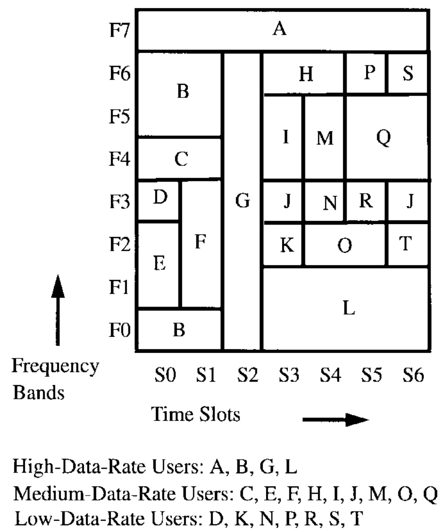

Fig. 2. An example of a time-frequency-sliced system with noncontiguous frequency assignments.

we compute the reduction in blocking probability attainable by using a variable bit-rate transmitter/receiver to dynamically increase/decrease the bandwidth of the transmitter/receiver. In general, though, it is not necessary to equip all users with variable bit-rate transmitters/receivers. Nevertheless, the time-frequency concept allows users, in practice, to have their own fixed-rate transmitter/receiver.

Finally, Section VII summarizes the advantages of the time-frequency-code slicing technique.

II. Noncontiguous Modulation

In the example above, it might appear that high-bandwidth users require their frequency slots to be assigned contiguously, as in Fig. 1. However, other modulation schemes, for example, multitone, may allow the adjacency requirement to be relaxed. In this example, tones represent multibit symbols, and each tone toggles at a rate corresponding to the bandwidth of one frequency band. Thus, two bits can be transmitted as one quaternary symbol using two-tone modulation instead of two binary symbols.

Higher data rates are available to high-bandwidth users by signaling on a combination of tones, whereas the low-bandwidth user would modulate only a single tone. The data-rate requirements of a user determines the number of tones or frequencies allocated for that user. These tones are scheduled in possibly noncontiguous frequency bands within one or more time slots as, for example, for user B in Fig. 2. In fact, spreading the frequency allocations of a high-bandwidth user may offer some propagation benefits (e.g., a reduction in the degradation from frequency-selective multipath fading).

Note that the single transmitter–receiver advantage of the first example (Fig. 1) has to be sacrificed for the high-bandwidth users to obtain this scheduling advantage. But the base-station receiver may be simplified since only one type of (low-bandwidth) transmitter/receiver may be used. Otherwise, to accommodate different transmission rates, banks of different-bandwidth transmitters/receivers need to be provided.

Multitone systems require a linear power amplifier at the transmitter. To avoid the added cost and complexity associated with linear power amplifiers, m-ary components of the multitone system can be modulated by a constant envelope scheme, such as continuous-phase frequency-shift keying (CPFSK), for example.

Noncontiguous frequency assignment could be used with today's channelized cellular systems. For example, a multitone cellular system could provide higher bandwidth to some users. This would be accomplished by allocating multiple channels to each user with high-data-rate requirements. Since these allocations do not need to be contiguous, more users can perhaps be supported compared to the time-frequency example described in Section I.

Other variations on the proposed scheme include dividing the resource space into a time-code space or a time-frequency-code space. Articulation of the relative merits and costs of these systems is beyond the scope of this paper. In the next section, we briefly discuss various time-frequency-code slicing approaches.

III. Slicing Along the Code Dimension

In this section, we describe how the slicing technique can be used in the code dimension. We assume here basic knowledge of code-division multiple access (CDMA). For the reader's convenience, we provide some introductory CDMA material in Appendix B. (For more information, the reader is referred to [6] and [7].)

In the discussion here, we assume that the narrowband signals are transmitted with powers that are multiples of some minimum power $P_{\min}^{\text{signal}}$. Suppose first that all spreading codes have the same chip rate R^{code}. A *code slot* is defined as the noise density (per hertz) that a narrowband signal of power $P_{\min}^{\text{signal}}$ contributes at the receiver when spread with the spreading code of the rate R^{code}. Then, in the three-dimensional case, a *unit slice* corresponds to the fraction of the total allowable noise density (per hertz) that is used to transmit a signal of minimum power $P_{\min}^{\text{signal}}$ spread by a code of rate R^{code} in one frequency band during one time slot.

In general, though, the time-frequency-code slicing technique allows spreading codes of different rates. We define the maximum rate of the spreading code $R_{\max}^{\text{code}}$ as the chip rate that corresponds to the bandwidth of the entire frequency space. We further assume that the rate of any spreading code R_i^{code} satisfies

$$R_{\max}^{\text{code}} = k \cdot R_i^{\text{code}} \tag{1}$$

for some positive integer k.

In the multirate spreading case, the definition of a unit slice (of the wireless resource) is more complex since the spread signals contribute different amounts of noise density, depending on the rate of the spreading code. First, a minimum quanta of noise density is defined, referred to as the code slot, corresponding to the noise power level (per hertz) at the receiver of a minimal-power $(P_{\min}^{\text{signal}})$ narrowband signal spread with the maximum-rate code. The larger the rate of the spreading code, the lower the level of noise per hertz. Then, the *code space* (i.e., the total number of code slots) is the maximal allowable total (aggregated over all the users) density of noise, based on some quality-of-service consideration, such as, for example, bit error rate (BER) or subjective voice-intelligibility tests.

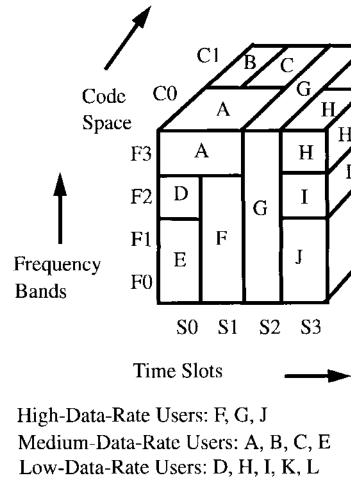

Fig. 3. An example of a time-frequency-code-sliced system.

From (1), the noise level (per hertz) of different spreading codes will always be an integer numbers of code slots. This allows us now to define a *unit slice* as a fraction of the total wireless resource that occupies one frequency band during one time slot and contributes at the receiver a noise power level equal to the power of the minimal power signal $P_{\min}^{\text{signal}}$ spread with a code of maximum rate $R_{\max}^{\text{code}}$. This implies a time-frequency-code-sliced system (Fig. 3) in which high-bandwidth users modulate their signal to occupy more than one frequency band. A user would have to be assigned enough frequency bands to accommodate the spreading associated with the chip rate R_i^{code}. Other users may share the same bandwidth at the same time using different codes.

The idea in allocating users a portion of the code space is in each time/frequency slot to allow noise to accumulate up to some maximum noise level. Each call is identified by an incremental noise level, measured at the receiver, and contributed in a specific frequency slot during a specific time slot. The resulting total noise level allows the participating transmissions to ensure the required quality-of-service although the allocation procedure may not necessarily lead to maximal packing of users in the code domain.

A particular instance of the time-frequency-code-sliced system is the frequency-code-sliced system (Fig. 4). In this case, users are continuously granted some portion of the frequency spectrum to be used with different rate spreading codes. Different users then continuously contribute a different amount of noise (i.e., occupy different portions of the code space).

Although the chip rate does not need to be fixed (among users and codes), in the rest of this section we assume, for exemplary purposes only, that the chip rate of the spreading codes is of constant rate. Also, we assume a single (and fixed)

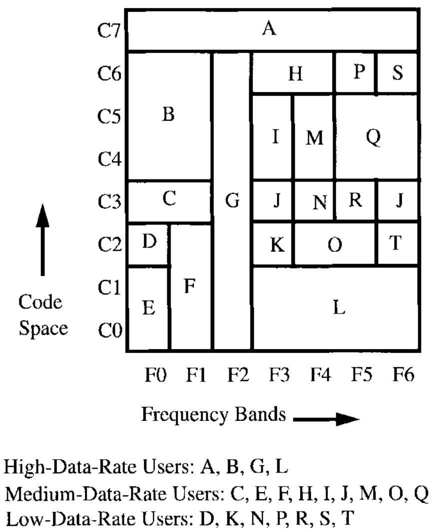

Fig. 4. An example of a frequency-code-sliced system with variable chip rates.

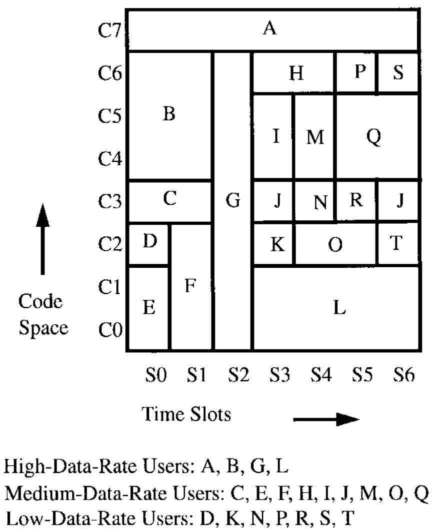

Fig. 5. An example of a time-code-sliced system, where a subset of the codes is allocated to each user.

BER threshold for all the users above which the quality of service becomes unacceptable to all the users in the system. In general, users with higher bit rates transmit at higher power levels (i.e., to keep the energy per bit constant). Thus, because of the constant spreading-sequence chip rate, higher data-rate users contribute more interference (perceived as noise) than lower bit-rate users. Consequently, the scheduling process consists of granting users a number of codes (e.g., as in a multicode CDMA system [8], [9]) so that the BER caused by the total level of interference from all the transmissions remains below the threshold, i.e., so that the "code space" is not exceeded.

In the time-code slicing approach, codes can be reused in different time periods, thereby supporting a large user population with a relatively small number of codes. Also, scheduling attempts to pack each time slot while maintaining acceptable bit error rates. An example of such a schedule is shown in Fig. 5, where the amount of code space a user occupies is related to its transmission rate.

In unslotted CDMA systems, mechanisms are sometimes required to limit the maximum number of users accessing the system so that a minimum quality can be guaranteed for each user. In this time-code-sliced system, the scheduler and time slots provide direct control on the number of users accessing a time slot (i.e., exercising an aspect of congestion control), thereby guaranteeing a particular quality of service for large user populations. Alternatively, higher data-rate users can use more power or a lower chip rate, corresponding to higher levels of interference and a smaller number of users sharing the spectrum.

IV. Performance of Time-Frequency-Sliced Systems

Users can demand a number of slices, based on their traffic requirements. For example, a large number of slices may be necessary to support video traffic, while for a telephone call, a small number of slices suffices. For multimedia traffic, different allocations can be given to different traffic types, based on the particular requirement of a specific traffic type. Moreover, when there is a large demand for the wireless resource, the actual allocation granted to the users may be lower than requested, thus allowing an allocation of resources with flexible quality of service, compared with the current "all-or-nothing" traditional channel-assignment schemes.

The actual allocation of slices in the time, frequency, and code domains depends on the availability of the "slots" along each one of these dimensions and the capabilities of the transmitting equipment. For example, if the transmitter is capable of fast frequency tuning, different frequency slots (i.e., bands) can be allocated in adjacent time slots. If the transmitter is capable of spreading with a maximum-rate spreading code, a larger number of frequency bands can be allocated. On the other hand, if the transmitter bandwidth is limited to the bandwidth of the narrowband signal itself, there can be no spreading and the traffic will be allocated all the code slots in a particular frequency band during a particular time slot.

Some multimedia traffic is continuous in nature, whereas other traffic may need allocation in an on-demand basis. For continuous-type traffic, slices can be periodically allocated in every frame. For randomly arriving traffic, such as data or images, a request for an allocation is made, which carries with it an indication of the amount of data to be transmitted in addition to the specifications of the transmitter (such as the maximum agility rate or the maximum spreading code rate). The assignments of slices are then performed by the scheduler. The function of the scheduler is to allocate slices to new requests of the continuous and on-demand traffic types. The slicing technique facilitates the operation of the scheduler as to maximally pack the available wireless resource.

To illustrate the performance advantages of time-frequency-code slicing, we now show the dramatic reduction in blocking probability attainable with TFS systems. Consider a system where each user has a single tunable variable-bit-rate transmitter/receiver. In contrast with today's systems (e.g., GSM), we assume the bandwidth and center frequency of the transmitter/receiver can be adjusted on a slot-by-slot basis.

The main focus of Sections IV–VI is to determine how much performance advantage is gained by allowing users to modulate more than just one unit (frequency) slice per time slot.

There are F frequency bands and N time slots per (periodic) frame, and the traffic load is ρ. For simplicity, we assume each time-frequency unit slice is (independently) busy with probability ρ. Blocking probabilities may actually be lower than the results obtained using these independence assumptions since the idle frequencies will tend to be contiguous in each time slot (recall that the frequency assignments are assumed to be contiguous). Under these assumptions, we compute the blocking probability as a function of the "bandwidth request" (i.e., the requested number of time-frequency unit slices). We compare the blocking probability of "spectrally efficient assignment" (i.e., TFS) to the blocking probability of a "traditional allocation" scheme. In the traditional allocation scheme, because of the fixed division of the frequency spectrum into equal-size frequency bands and the single-transmitter-per-user constraint, at most one frequency band per time slot can be assigned to a user. However, different frequencies can be assigned to the same user in different time slots. Examples of systems that can, in principle, support the traditional allocation include the digital IS-54 and GSM standards.

In a time-frequency-sliced system, any number of idle frequencies per time slot can be assigned to satisfy the bandwidth request, provided the frequencies are contiguous (recall, again, the assumption of a single tunable transmitter). Consequently, the bandwidth request can potentially be for as many as FN unit slices. In contrast, in the traditional allocation scheme, the bandwidth request can be for at most N unit slices.

V. Analytic Computation of Blocking Probabilities

In this section, we will compute the probability that a new connection request is blocked as a function of the traffic loading ρ and the "bandwidth" requested (i.e., the number of unit slices). We first focus our attention on one particular time slot. There are F frequencies in the time slot, each (independently) busy with probability ρ. Suppose we let $P_F(m)$ denote the probability that the maximum number of contiguous idle frequencies is greater than or equal to m. Then

$$P_F(1) = 1 - \rho^F \tag{2}$$

and

$$P_F(F) = (1-\rho)^F \tag{3}$$

since $1 - \rho^F$ is the probability that there is at least one idle frequency in the time slot and $(1-\rho)^F$ is the probability that all F frequencies are idle. Naturally, $P_F(m) = 0$ if $m > F$. For each m in the range $2 \le m \le F - 1$, we first express $P_F(m)$ as a function of $P_1(m), P_2(m), \cdots$ and $P_{F-m}(m)$

$$P_F(m) = (1+\rho)(1-\rho)^m + \sum_{f=2}^{F-m} [1-P_{f-1}(m)]\rho(1-\rho)^m. \tag{4}$$

In (4), $[1 - P_{f-1}(m)]$ is the probability that there are fewer than m contiguous idle frequencies in the "first" $f - 1$ frequencies. This term is multiplied by ρ, which is the probability that the "next" frequency (i.e., the "fth" frequency) is busy and multiplied by the probability $(1-\rho)^m$ that the "following" m frequencies are idle (thereby yielding at least m contiguous idle frequencies in the time slot). Furthermore, because the summation in (4) starts at $f = 2$, to obtain $P_F(m)$ we must also add $(1-\rho)^m$ and $\rho(1-\rho)^m$, which represent, respectively, the probability that the "first" m frequencies are idle and the probability that the "first" frequency is busy, but the "next" m frequencies are idle.

Using (4), we can compute $P_F(m)$ by recursion on the number of frequencies. Specifically, for $(m+1)$ frequencies, we first compute $P_{m+1}(m) = (1+\rho)(1-\rho)^m$. We then compute $P_{m+2}(m)$, $P_{m+3}(m), \cdots$, and finally $P_F(m)$.

After computing the probability $P_F(m)$ that the maximum number of contiguous idle frequencies is greater than or equal to m (for all m in the range $1 \le m \le F$), it is easy to determine the probability $Q_1(m)$ that the maximum number of contiguous idle frequencies (out of F total) in one time slot is exactly equal to m

$$Q_1(m) = P_F(m) - P_F(m+1). \tag{5}$$

The subscript "1" in $Q_1(m)$ corresponds to the one *time slot*. The dependence on the number of frequencies (F) is not listed for simplicity.

The last step in the analysis is to sum the "bandwidth availability" (i.e., the maximum number of contiguous idle frequencies) across all time slots[2] and check whether or not the total bandwidth availability is greater than (or equal to) the "bandwidth request." Note that the "available bandwidth" over n time slots equals the bandwidth over $(n-1)$ time slots plus the bandwidth available in one more slot. Therefore, we can compute the bandwidth availability by recursion on the number of time slots n and by using the following convolution:

$$Q_n(k) = Q_{n-1}(k) * Q_1(k) = \sum_{r=0}^{k} Q_{n-1}(r) \cdot Q_1(k-r). \tag{6}$$

In (6), $Q_n(k)$ represents the probability that the total "bandwidth availability" is k across the complete set of n time slots.

Finally, for a time-frequency-sliced system with F frequencies and N time slots (per periodic frame), the blocking probability is

$$\text{Blocking Prob.} = \sum_{k=0}^{B-1} Q_N(k) \tag{7}$$

where B denotes the bandwidth request (in unit slices) and can be any value in the range $1 \le B \le FN$.

For comparison with our time-frequency-sliced system, we also compute the blocking probability associated with the traditional allocation scheme (in which users can only modulate a

[2] Suppose there is a constraint on the maximum bandwidth W (expressed as a number of contiguous frequencies) that a user can modulate per time slot. Then, this constraint can be incorporated in the performance analysis by replacing Q with Q^T in (5), where Q^T is a truncated version of Q: $Q_1^T(m) = Q_1(m)$ if $m < W$, $Q_1^T(W) = P_F(W)$ and $Q_1^T(m) = 0$ if $m > W$.

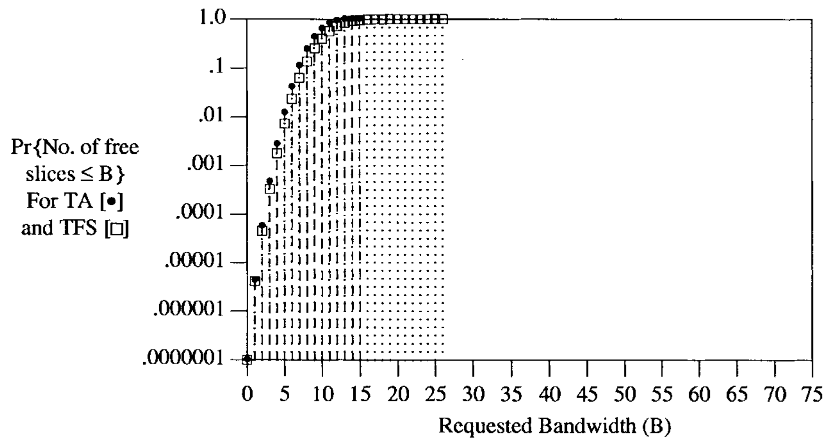

Fig. 6. Blocking probability as a function of bandwidth request for $F = 10$, $N = 15$, and $\rho = 0.9$.

single frequency band per time slot). The maximum allowable bandwidth request is $B = N$, and the blocking probability is simply the probability that there are fewer than B time slots with an idle frequency band. This is given by

$$\text{Blocking Prob.} = \sum_{k=0}^{B-1} \binom{N}{k} [P_F(1)]^k [1 - P_F(1)]^{N-k}. \quad (8)$$

In (8), $P_F(1)$ represents the probability that a time slot has (at least) one idle frequency band, and $[1 - P_F(1)]$ is the probability that there are no idle frequencies in a time slot, which is equivalent to $Q_1(0)$.

VI. Numerical Results

The following graphs present, on a logarithmic scale, some numerical results obtained using (6)–(8).

Figs. 6 and 7 compare the blocking probability for the TA [·] and the TFS technique [□]—for $\rho = 0.9$ and 0.7, respectively (with $F = 10$ and $N = 15$). The decrease in the blocking probability attained with a spectrally efficient assignment (i.e., TFS) is clearly demonstrated in these graphs. For example, at $\rho = 0.9$ with 15 time slots per frame, there is still a 1% chance of providing a 17-slice allocation with the TFS technique, whereas there is no chance at all to allocate more than 15 unit slices in the TA scheme. These comparisons are even more profound as ρ decreases. For example, at $\rho = 0.7$ with 15-slot frames, the TA scheme can again accommodate requests for at most 15 unit slices, while the TFS technique still has a 1% chance of satisfying a request for even 34 unit slices. This increased gain of the TFS technique at lower utilization, as compared with the TA scheme, is expected; as the utilization decreases, there are more and more contiguous idle frequency bands and more chances of accommodating larger requests.

The next two graphs (Figs. 8 and 9) show the same types of results as in Figs. 6 and 7, but for 30 frequency bands (i.e., $F = 30$). With more frequency bands, the TFS technique advantage over the TA scheme is more pronounced. For example, at $\rho = 0.9$ with 15 time slots per frame, there is still a 1% chance of providing a 24-slice allocation with the TFS technique, whereas there is no chance at all to allocate more than 15 unit slices in the TA scheme. Also, as before, as ρ decreases, the TFS technique performs considerably better than the TA scheme. For example, at $\rho = 0.7$ with 15-slot frames, the TA scheme can again accommodate requests for at most 15 unit slices, while the TFS technique still has a 1% chance of satisfying a request for even 47 unit slices.

VII. Summary

We have proposed time-frequency-code slicing techniques for providing access to a communications resource for multiple users at a variety of rates while maintaining both low-cost access for low-data-rate users and good spectral efficiency. For the specific case of TFS, we showed that it offers higher bandwidth allocations and reduced blocking probabilities, especially for smaller utilizations and a larger number of frequency bands. And because current cellular systems rely on large number of channels (e.g., 124 channels in GSM), mostly operated at low utilization, the proposed scheme will allow use of these channels to accommodate bursty traffic with low access latency. As the number of data applications increases on wireless networks, access for bursty traffic (which is characteristic of computer communications) will become of prime interest and importance in such networks. Finally, although we focused on just one variation of the general time-frequency-code slicing technique, similar performance advantages are attainable with the other variations. Naturally, there are numerous implementation tradeoffs that need to be considered along with the performance comparisons.

Appendix A
Performance with Noncontiguous Frequency Assignments

As mentioned in Section II, there are scheduling advantages associated with noncontiguous frequency assignments. In this Appendix, we provide a brief outline of the performance analysis associated with time-frequency-sliced systems that schedule the transmissions with noncontiguous frequency assignments. The analysis is actually simpler than the analysis in Section V for the assignment with contiguous frequencies.

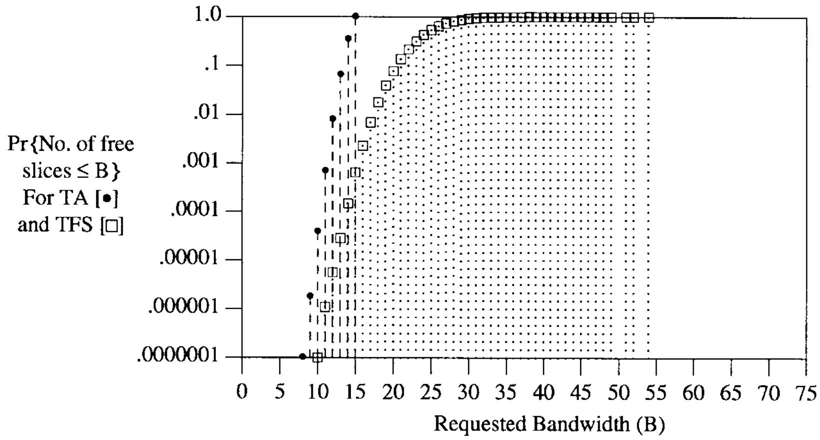

Fig. 7. Blocking probability as a function of bandwidth request for $F = 10$, $N = 15$, and $\rho = 0.7$.

As in Section V, we will compute the probability that a new connection request is blocked as a function of the traffic loading ρ and the "bandwidth" requested (i.e., the number of unit slices). We first focus our attention on one particular time slot. There are F frequencies in the time slot, each (independently) busy with probability ρ. Suppose we let $V_1(m)$ denote the probability that there are exactly m idle frequencies (out of F total) in one time slot. Then

$$V_1(m) = \binom{F}{m}(1-\rho)^m \rho^{F-m}. \tag{9}$$

Now, let $W_{\min}$ and $W_{\max}$ represent the minimum bandwidth and the maximum bandwidth, respectively, that a user can modulate per time slot (expressed as a number of contiguous frequencies). The special case $W_{\min} = W_{\max}$ corresponds to a constant modulation bandwidth. We define a truncated version of V_1 as follows: $V_1^T(m) = 0$ if $m < W_{\min}$ or $m > W_{\max}$, $V_1^T(m) = V_1(m)$ if $W_{\min} \leq m < W_{\max}$, and $V_1^T(m) = \Sigma_{r=m}^{F} V_1(r)$ if $m = W_{\max}$.

Summing the bandwidth availability (i.e., the number of idle frequencies that can be accessed) across all time slots, we obtain a convolution similar to (6)

$$V_n^T(k) = V_{n-1}^T(k) * V_1^T(k) = \sum_{r=0}^{k} V_{n-1}^T(r) \cdot V_1^T(k-r). \tag{10}$$

As in Section V, we compute $V_n^T(k)$ from (10) by recursion on the number of time slots n.

Finally, for a time-frequency-sliced system with F frequencies and N time slots (per periodic frame), the blocking probability is

$$\text{Blocking Prob.} = \sum_{k=0}^{B-1} V_N^T(k) \tag{11}$$

where B denotes the bandwidth request (in unit slices) and can be any value in the range $1 \leq B \leq FN$.

Finally, note that the above analysis can be greatly simplified in the special case $W_{\min} = 1$ and $W_{\max} = F$. Then, the blocking probability is simply the probability that there are fewer than B idle unit slices out of the total FN

$$\text{Blocking Prob.} = \sum_{k=0}^{B-1} \binom{FN}{k}(1-\rho)^k \rho^{FN-k}. \tag{12}$$

Comparing the results obtained with (12) with those in Section VI, we see that for $\rho = 0.9$ (with $F = 10$ and $N = 15$) and a requested bandwidth of $B = 17$, the blocking probability decreases from 99% (with contiguous frequency assignments) to 67% (with noncontiguous frequency assignments). Also, for $\rho = 0.7$ (with $F = 10$ and $N = 15$) and a requested bandwidth of $B = 34$, the blocking probability decreases from 99% (with contiguous frequency assignments) to 2% (with noncontiguous frequency assignments). The reductions in blocking probabilities are even more dramatic when the number of frequencies is increased to $F = 30$. However, one needs to remember that these comparisons are based on the assumption that noncontiguous frequency assignments are possible with $W_{\min} = 1$ and $W_{\max} = F$. Further investigations are necessary to examine the performance and implementation tradeoffs of contiguous and noncontiguous frequency assignments in time-frequency-sliced systems.

Appendix B
A Short Overview of CDMA

In a CDMA system based on spread spectrum, each user is given a unique and orthogonal spreading code from a set of signals. This set, large but finite, can be composed of different phases of a long pseudo-noise (PN) sequence. When users access the channel, they multiply their data stream by their assigned code. The code rate is considerably higher than the data bit rate and is referred to as the *chip rate*. Thus, the spreading codes are used to spread the spectrum of the transmitted narrowband signals. At the receiving end, the destination correlates the received signal by a synchronized replica of the source code to recover the original signal.

CDMA supports multiple simultaneous transmissions because the cross correlation between two different codes is small. Thus, if a signal encoded with one code C_1 is decoded

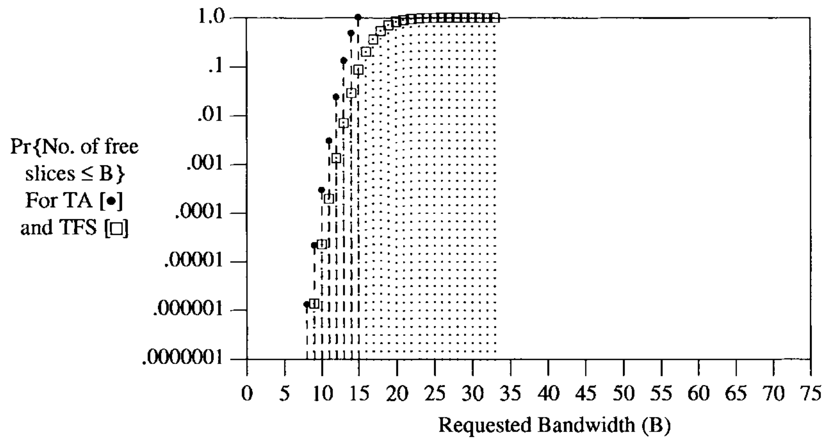

Fig. 8. Blocking probability as a function of bandwidth request for $F = 30$, $N = 15$, and $\rho = 0.9$.

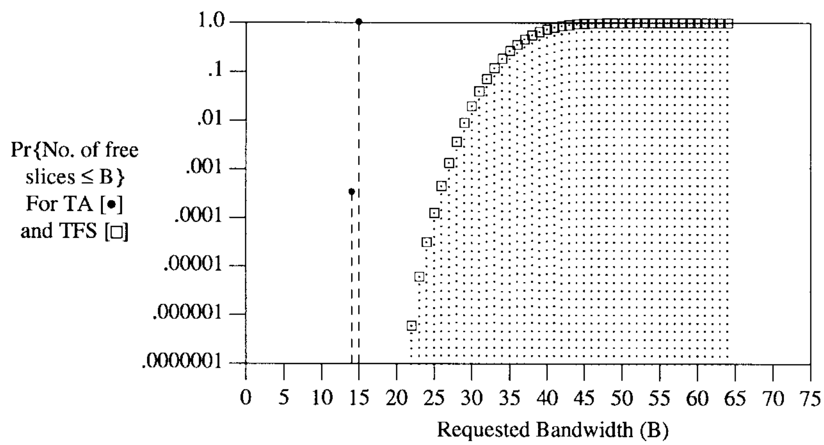

Fig. 9. Blocking probability as a function of bandwidth request for $F = 30$, $N = 15$, and $\rho = 0.7$.

with a different code C_2, the result appears to the receiver as noise. The limitation of the scheme (i.e., the maximum number of users that can utilize the multiaccess channel) depends on the total amount of "noise" contributed by "interfering" users to the detected signal. In other words, the more users simultaneously transmitting on the channel, the more noise is detected by the receivers. The signal-to-noise ratio (SNR) determines the BER performance of the system.

By spreading a signal over a larger bandwidth (i.e., with a faster chip rate), more independent transmissions can be scheduled in the same bandwidth. Scheduling of the independent transmissions depends on the interference level contributed from each transmission so that the BER of the scheduled transmissions is kept below some predetermined threshold.

In the spectral domain, the multiplication of the data by the fast-bit-rate code corresponds to spreading the data spectral components over a broader spectrum. Thus, a larger spectrum is required to convey the transmission. However, because of the multiple-access feature, a number of users may coexist at any time on the channel. The bandwidth ratio of the spread to the unspread signals is called the processing gain. The larger the processing gain, the less "noise" contribution a single user has on the other users' signals.

REFERENCES

[1] Z. Haas and R. D. Gitlin, "Optical distribution channel: An 'almost-all' optical lan based on the field coding technique," *J. High-Speed Network.*, vol. 1, no. 3, pp. 193–214, 1992.

[2] R. A. Thompson, J. J. Horenkamp, and G. D. Bergland, "The photonic switching of universal time slots," in *Proc. Int. Switching Symp.*, May 1990, pp. 165–170.

[3] C. B. Woodworth, M. J. Karol, Z. J. Haas, and R. D. Gitlin, "Spectrally efficient universal time slots using time-frequency-code slicing," in *Proc. WCN*, Sept. 1994, pp. 1009–1013.

[4] M. J. Karol, Z. Liu, and K. Y. Eng, "Distributed-queueing request update multiple access (DQRUMA) for wireless packet (ATM) networks," in *ICC'95 Conf. Rec.*, June 1995, pp. 1224–1231.

[5] J. A. C. Bingham, "Multicarrier modulation for data transmission: An idea whose time has come," *IEEE Commun. Mag.*, vol. 28, pp. 5–14, May 1990.

[6] R. C. Dixon, *Spread Spectrum Systems*. New York: Wiley-Interscience, 1976.

[7] R. L. Pickholtz, D. L. Schilling, and L. B. Milstein, "Theory of spread-spectrum communications: A tutorial," *IEEE Trans. Commun.*, vol. COM-30, pp. 855–884, May 1982.

[8] H. Azad and A. H. Aghvami, "Multirate spread spectrum direct sequence cdma techniques," in *IEE Colloquium on Spread Spectrum Techniques for Radio Communication Systems, Digest 1994/098*, Apr. 1994, pp. 4/1–4/5.

[9] C.-L. I and R. D. Gitlin, "Multi-code CDMA wireless personal communications networks," in *ICC'95 Conf. Rec.*, June 1995, pp. 1060–1064.

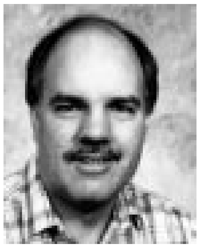

Mark J. Karol (S'79–M'85–SM'91–F'93) received the B.S. degree in mathematics in 1981 and the B.S.E.E. degree in 1981 from Case Western Reserve University, Cleveland, OH, and the M.S.E., M.A., and Ph.D. degrees in electrical engineering from Princeton University, Princeton, NJ, in 1982, 1984, and 1986, respectively.

Since 1985, he has been a Member of the Research Communications Sciences Division at Bell Laboratories (formerly part of AT&T, now part of Lucent Technologies), Holmdel, NJ. He is currently a Distinguished Member of Technical Staff, and his research interests include high-performance packet-switching architectures, local- and metropolitan-area network architectures, wireless communications networks, and multiuser lightwave communication networks.

Dr. Karol is a Member of Eta Kappa Nu and Tau Beta Pi. He is currently on the Board of Governors of the IEEE Communications Society. He has served as an Associate Editor for the JOURNAL OF LIGHTWAVE TECHNOLOGY, as Chair of the IEEE Communications Society Technical Committee on Computer Communications, and as General Chair for the IEEE Infocom'94 Conference.

Zygmunt J. Haas (S'84–M'88–SM'90), for a photograph and biography, see this issue, p. 817.

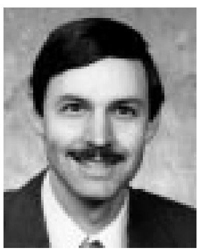

Clark B. Woodworth (S'76–M'77) received the B.S.E.E. degree in 1977 from Monmouth University, NJ, the M.S.E.E. degree in 1980 from Rutgers University, New Brunswick, NJ, and the Ph.D. degree in 1988 from Polytechnic University, Brooklyn, NY.

He worked in the Satellite Systems Research and the Radio Systems Research Departments at Bell Labs. His work in the Radio Systems Department included digital radio studies and multipath channel simulation. More recently, he worked in the Network Systems Research Department, where he developed a 150-Mb/s multicast ATM switch, a multimedia voice-video terminal, and a multimedia server. He is currently working on wireless ATM in the Broadband Systems Research Department, Bell Labs, Lucent Technologies, Holmdel, NJ, and his present interests include ATM local-area networking, B-ISDN, and multimedia.

Dr. Woodworth is a Member of the Eta Kappa Nu and Sigma Pi Sigma honor societies.

Richard D. Gitlin (S'67–M'69–SM'76–F'86) was born in Brooklyn, NY, on April 25, 1943. He received the D.Eng.Sc. degree from Columbia University, NY, in 1969.

He is the Research Vice President of the Communication Science Division of Bell Laboratories, the R&D unit of Lucent Technologies, Holmdel, NJ. This division conducts research in speech and image processing, wireless communications, access systems, photonic systems, and high-speed broadband networking. The goal of this division is to extend the limits of communications technologies to provide key innovations for Lucent Technologies. Throughout his career, he has both conducted and led pioneering research and development in digital communications that have led to many innovative products. Before undertaking his current assignment in 1995, he was Director of the Communications Systems Research Laboratory, where he was responsible for research in wireless systems, broadband networking, and local-access and switching systems. From 1987 to 1992, he was Head of the Network Systems Research Department, where he led research in broadband networking, ranging from gigabit-per-second packet switches and local-area networks to high-speed protocols, broadband applications, and the gigabit research network known as LuckyNet. From 1982 to 1987, he was Head of a department responsible for systems engineering, exploratory development, and final development of data communications equipment. He guided the pioneering efforts that led to the V.32/V.34 product family of high-speed modems, and he is a Coinventor of the high-speed (asymmetric) digital subscriber line HDSL/ADSL technology for multimegabit transmission on copper loops. During his career at Bell Labs, he did applied research and exploratory development in the field of high-speed voiceband modems and later supervised a group doing exploratory and advanced development in this area. He is the author of more than 70 technical papers and numerous conference papers, and he holds 39 patents in the areas of data communications, digital signal processing, wireless systems, and broadband networking. He is also coauthor of the book *Data Communication Principles* (New York: Plenum, 1992) and coauthor of a paper selected as the best paper in communications by the Bell System Technical Journal in 1982.

Dr. Gitlin received the IEEE Communications Society's Fred Ellersick Prize Paper Award for coauthoring the paper "Reducing the effects of transmission impairments in digital fiber-optic systems." He received the 1995 S. O. Rice Prize Paper Award given by the IEEE Transactions on Communications for the paper "Diversity Coding," which he coauthored. He has served as Chair of the Communication Theory Committee of the IEEE Communications Society and is a Member of the COMSOC Awards Board, an Editor for Communications Theory of the IEEE TRANSACTIONS ON COMMUNICATIONS, and a Member of the Board of Governors of the IEEE Communications Society and the Nominations and Elections Board. In 1985, he was elected a Fellow of the IEEE for his contributions to data communications technology. In 1987, he was named a Bell Laboratories Fellow.