Plane Cover Multiple Access: A Media Access Control Strategy for Wireless Environments*

Paul M. Blair†, **George C. Polyzos**†, and **Michele Zorzi**‡

†Department of Computer Science and Engineering, University of California San Diego, La Jolla, CA 92093-0407, U.S.A.
fax: +1-619-534-7029 – e-mail: {pblair,polyzos}@cs.ucsd.edu

‡Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA 92093-0407, U.S.A.
fax: +1-619-534-2486 – e-mail: zorzi@ece.ucsd.edu

Abstract— **In order to support future demand for mobile, wireless data communications, a packet based transmission scheme needs to be adopted. Such a paradigm could support data as well as voice and multimedia transmissions and would use the scarce wireless spectrum in a more efficient way than the conventional circuit switched model. In this paper we present a Media Access Control strategy, called Plane Cover Multiple Access (PCMA), which provides a means of allocating wireless bandwidth in a packet based system. Mobiles in all cells are assigned various reuse factors and bandwidth is allocated in the form of time slots to users such that their reuse factors are preserved on a system wide basis. As such, we are proposing a simplified version of the real-time packet scheduling problem. In its simplest form, reuse factors are determined by position within a cell in a way that essentially remains fixed after being established as part of the system design, assuming a fixed traffic distribution. However, the reuse factor determination may change in real-time to account for changing system traffic distributions. Through a simple example of PCMA we show that considerable increases in system capacity can be gained by careful selection of the reuse factors. Our example achieves a system utilization factor of almost 2 times that of the standard one-seventh frequency reuse plan. Finally, PCMA is ideal for supporting non-uniform user traffic and can easily adjust to time varying user traffic patterns so as to maximize system throughput.**

I. Introduction

Given the foreseeable demand for next-generation wireless communication services, the circuit switched transmission mode and resource allocation methods used in current cellular systems will be inadequate. Consumers want more than voice communications. They want data transmission capabilities to permit sending files, electronic mail, and multimedia traffic, including real-time video. The popularity of the Internet is such that consumers want mobile access to it. In order to provide the bandwidth to support such an array of services, some major changes need to take place.

The first is that a packet switching based paradigm must be invoked. Circuit switching has served wireless voice communications well over the past few decades. However, the bursty nature of data communications demands packet switching. Even with apparently continuous media like video, packet switching is preferable. Video will need to be compressed to fit bandwidth constraints in a wireless environment and the compression algorithms used create variable data rates for real-time video. It is well known that voice, too, can be supported by packet switching [1]. The move toward ATM in wired networks is further evidence of the flexibility of packet based transport.

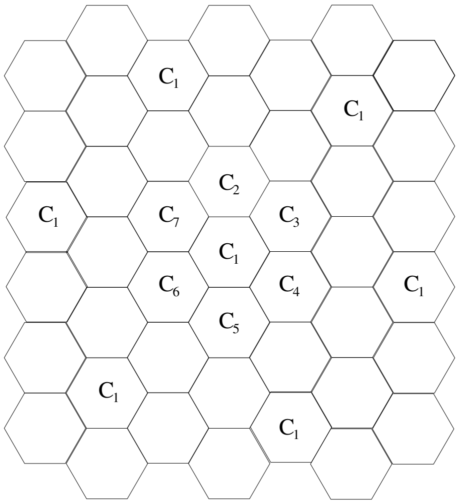

Fig. 1. The one-seventh frequency reuse plan

The second change is that new resource allocation schemes need to be implemented. The current common scheme (as used in AMPS and GSM, among others) is referred to as the one-seventh frequency reuse plan. It takes all the available channels in the system and splits them into seven (equal) sets, $S_1, S_2, ..., S_7$. Cells are also split into seven sets, $C_1, C_2, ..., C_7$, based on their location, resulting in a pattern as in fig. 1.

Then, users in a cell belonging to C_i may only use channels from set S_i, $i \in \{1, ..., 7\}$. Although the one-seventh reuse plan avoids interference through strict channel reuse rules, it performs poorly as far as system capacity. Since each cell can only use $1/7$ of the total resources, the maximum system utilization can never exceed $1/7$. In order to achieve greater system utilization, more intense reuse must be invoked. One possibility is to permit more intense reuse to all users in the system, while also providing a means of tolerating the increased interference level, by either using an interference-robust modulation (e.g., spread-spectrum as in

* Work supported by the Center for Wireless Communications at the University of California San Diego.

CDMA [2]) or providing error recovery mechanisms (e.g., by retransmission as in CDPA [3]). A second approach is to permit more intense reuse only to those users for which transmission quality will not be compromised. An example of such an approach is reuse partitioning [4] in which users closer to base stations are granted more intense reuse since their signals are still received strongly enough. In this paper we take this second approach and propose a new multiaccess protocol that avoids the high price of splitting resources ahead of time and permits statistical multiplexing. At present we do not consider using error recovery techniques to permit even greater frequency reuse, however such methods are currently being investigated.

A basic difference between our proposed scheme and classic reuse partitioning is essentially in its being packet-oriented, which provides a greater degree of flexibility in the management of traffic. Also, the use of transmission resources in a dynamic TDMA fashion allows statistical multiplexing, soft handoff and macrodiversity to be supported. Finally, as it will be explained later in the paper, our effort to maximize the reuse at any given point leads to topologies which are different from those considered in reuse partitioning.

The remainder of this paper is structured as follows. Section 2 indicates the basic system assumptions made throughout the paper. Section 3 formalizes the media access control problem in terms of system utilization. Section 4 presents an example of the PCMA protocol and analyzes its performance. Section 5 contains a discussion including future directions for this research. Section 6 is the conclusion.

II. System Assumptions

For simplicity of exposition we assume that separate frequency bands are used for the uplink (mobile to base) and downlink (base to mobile) directions. Furthermore, we assume that there is no splitting of the spectrum into "channels" in the frequency domain. Rather, we view all uplink (and similarly downlink) spectrum as one wide channel, or bit pipe. Resource sharing is accomplished by splitting this channel in the time domain. Thus, we form a series of equally sized time slots such that each slot can transmit exactly one packet, plus any system overhead. For simplicity of analysis, the overhead will be ignored.

Users generate packets at varying rates and transmit them in the available time slots. For real-time traffic there are deadlines by when successful transmission must take place, while for best-effort service there are no such deadlines. In the latter case, however, we would like to transmit as soon as possible. We assume that a base station can successfully receive at most one packet per time slot. Capture is not considered in this paper. Rather, we follow the assumption that the one-seventh reuse plan provides adequate interference protection. That is, in fig. 1, simultaneous transmissions in all cells labeled C_1 can be successfully performed if no other cells transmit at the same time. We will see that for mobiles located close to base stations, there are further possibilities.

We assume that each base station uses polling to assign slots to mobiles in its cell. The overhead for the polling is minimal because the cells are small in size. For example, the propagation delay incurred over 1 km at 1 Mb/s (or over 100m at 10 Mb/s) represents only 3 bits, which is negligible even for the small 53-byte packets used in ATM. An alternative to polling would be to use a random access channel to service the mobile's initial transmission requests, piggybacking later transmission requests on current packet transmissions as in the DQRUMA [5] or C-PRMA [6] protocols. Regardless of how the mobile's transmission requests are known, the base stations do scheduling of mobile transmissions using whatever discipline they choose. We assume that at the start of each time slot the base sends a mobile ID number indicating which mobile in the cell is permitted to transmit in that slot. Since only one mobile per time slot is given permission there are no intra-cell collisions. The idea is for the base station to optimally schedule packet transmissions so that all packets meet their deadlines and the interference generated is kept to a minimum. Thus, the highest number of packets can be successfully transmitted, maximizing system throughput.

We make no specific assumptions about the probability of interference when using the one-seventh reuse plan. It was designed to achieve the low interference probability required to support voice transmission. Whether or not it reaches this design goal in practice is not central to our study, since we will be comparing our protocol directly with the one-seventh plan and using its assumptions regarding minimum distances required for interference free communication. Thus, only the relative performance of the two schemes is of interest here, not the actual absolute performance of the one-seventh reuse plan.

Before proceeding to introduce the protocol, a few words are in order about the advantages of this setup relative to other systems. First of all, we are using packet switching which is inherently more flexible than circuit switching. Furthermore, since we are allocating time slots from a single channel used in all cells, soft handoff and macrodiversity can be supported. Also, increased traffic in a given cell can be handled by allocating a greater proportion of time slots to users in that cell. Thus, Dynamic Channel Allocation and bandwidth on demand are readily supported. Another advantage of invoking packet switching is that user calls do not necessarily need to be dropped due to handoff failure. Take, for example, the case in conventional systems where one runs out of channels in a cell. Then, any call in progress will have to be terminated if it roams into this congested cell. Now, even if none of the calls in progress in the congested cell is transmitting at that instant, the handed off call is still terminated. In our packet model, however, there would be unused slots available to service the new user at that instant. Even if there were no slots available at the instant of handoff, the call could still remain connected in our model at the cost of a degradation in service quality. A few packets might miss their deadline and need to be dropped, but as soon as some of the other calls left slots unused, the handed off call would be able to continue transmission. Such an approach is preferable to the alternative of dropping the call outright.

Finally, the packet switching model provides support for multiple levels of Quality of Service (QoS) guarantees. Assume that we have a number of service classes defined with different pricing levels. Then take the case where a call is handed off to a congested cell, but this call has a higher service class than a certain set, S, of calls in its new cell. In such a case the base station could actually take time slots from calls in the set S and give them to the handed off call so that it perceived no disruption in its more expensive service quality. In the circuit switched multi-channel model such an occurrence is not possible, since channels require lengthy set up and tear down procedures. Therefore, a call cannot simply borrow one or more of another call's channels for a few milliseconds.

III. Improved Reuse

In order to maximize throughput, we would ideally have a central scheduler knowing which mobiles have packets awaiting transmission, their deadlines, and the current physical environment conditions. Packets would then be optimally scheduled in all cells such that the maximum system throughput could be achieved with the fewest possible packets missing their deadlines. Due to the complexity of such an approach, traditional systems have adopted more conservative schemes. Under a system such as the one-seventh frequency reuse plan we would schedule one packet in every seventh cell per time slot. This would avoid missed packet deadlines due to inter-cell interference, but would achieve a poor system throughput of at most $1/7$.

We are faced with two extremes: one too ambitious and one far too conservative. Therefore, we seek a means of designing a Media Access Control protocol that finds a compromise between simple implementation and improved performance. We choose to further develop the reuse factor concept for simplicity. However, instead of limiting ourselves to the case where all users have the same reuse factor, we introduce a new model in which users at different locations in the system would be assigned different reuse factors depending upon their position within a cell. This approach is similar to reuse partitioning [4]. However, the assignment should be done based on current traffic conditions so as to maximize system throughput by taking advantage of parallel transmissions whenever possible. Packet switching and the TDMA structure provide the benefit of additional scheduling flexibility. Formally, we can state the problem as follows:

For each point, x, in a cell, assign a reuse factor $r(x)$. The reuse factor is such that users at all points in all cells that have the same position as x can be placed into clusters of size $r(x)$ and achieve the same quality as the one-seventh reuse plan[1]. Note that users within the same cluster cannot transmit simultaneously, but rather must alternate transmissions according to their reuse factor. Users in the best conditions, i.e., those very close to their base station, can all be given permission to transmit at the same time (one per cell), so that in some area around the base stations $r(x) = 1$ (note in fact that the one-seventh reuse scheme is designed to guarantee sufficient quality at the cell border, and therefore users close to the base station will enjoy exceedingly high quality). On the other hand, users farther away from the base stations would suffer if the same channel (slot) were assigned to users at the same location in neighboring cells, so that the system should not allow them to simultaneously transmit but rather to access the channel at different times, i.e., $r(x) > 1$ in this case. As an example (see fig. 2), if $r(x) = 3$ and mobile i is in cell C_m at point x, then during each time slot when i is transmitting, no cells in the ring of neighboring cells of C_m may permit users to transmit, whereas all shaded cells may allow users at location x to transmit simultaneously[2]. In the worst case, we would assign $r(x) = 7$ which is known (from our initial assumption) to provide good quality regardless of the location.

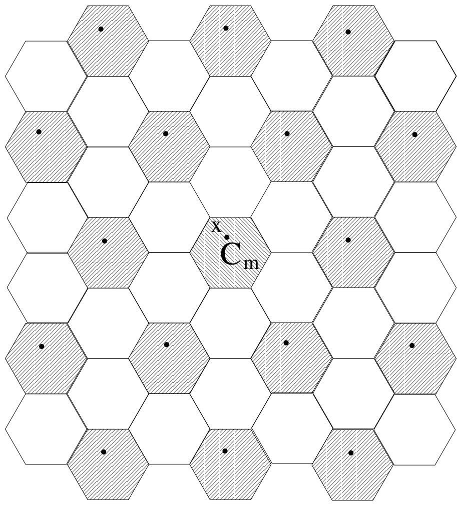

Fig. 2. One user with reuse factor 3 in each patterned cell can simultaneously transmit

In order to account for possibly non-uniform traffic distributions, we let $\lambda(x)$ be the traffic intensity at point x. The system throughput is then computed as

$$\frac{\int_{area} \lambda(x)\,dx}{\int_{area} \lambda(x) r(x)\,dx} \qquad (1)$$

Here the numerator represents the minimal possible bandwidth required to service the system, which would be the case if all users in a typical cell had reuse factor 1. The denominator represents the actual amount of bandwidth that is required in the typical system cell given that users have varying reuse factors, $r(x)$, many of which will be greater than one. For example, if $r(x) = 3$, in order to serve users at location x in all cells we need to allocate three time slots which will be alternately accessed, and therefore three times as much bandwidth as needed for a single user must be reserved. The expression in (1) determines the system throughput, which would achieve a maximum of one if $r(x) = 1$ for all x. Note that in the one-seventh reuse plan, $r(x) = 7$ for all x since the reuse factor is always 7, regardless of a mobile's location. Therefore, the system throughput is $1/7$. The traffic intensity, $\lambda(x)$, allows the flexibility of analyzing real systems with non-uniform traffic distributions. However, in the simplified case of uniform traffic distributions, $\lambda(x) = C$ for all x. Therefore, C factors out of the integrals in both the numerator and denominator and cancels.

The motivation for defining the problem in these terms is similar to the motivation behind reuse partitioning [4]. We are trying to use the minimum reuse factor possible in order to maximize system throughput. There are, however, several differences between our approach and reuse partition-

[1] Since we are using packet switching, we are not limited to integral values for $r(x)$. By varying the transmission probability for mobiles at various points we can achieve any expected cluster size $r(x)$, $r(x) \in [1, 7]$

[2] Note, on the other hand, that users at different locations (e.g., at the cell border) may not have the right to transmit, as they might interfere with transmission in C_m.

ing. For one thing, we are looking at the reuse factors necessary for packet transmissions, not at the reuse factors necessary to support long term connections, or channels. The interference characteristics are vastly different in the two cases. In our case, because of the statistical nature of packet generation, we can take advantage of statistical multiplexing and time diversity to attain even smaller reuse factors than possible with circuit switching. For example, we might be able to assign a certain group of points reuse factor 1, since users at these points generate packets in a round-robin fashion where no two require simultaneous transmission. In the classic reuse partitioning using circuit switching, such a low reuse factor would not be possible, since dedicated channels would have to be established with each user such that all could be simultaneously active. It is possible that all users have simultaneous packet transmissions at all times. In this case a reuse factor greater than one would have to be employed. However, our scheme is designed to take advantage of lower reuse factors whenever possible, in ways that circuit switching schemes cannot. Another means of achieving lower reuse factors is to take advantage of packet capture. When using capture, a packet is correctly received if the ratio of its SIR to the sum of all interferers is greater than a given threshold. Therefore, capture also permits lower reuse distances to be used. Even though the introduction of the analytical models necessary to work with capture is out of the scope of the present paper, such extensions are currently being explored. (A description of these models can be found in [3].)

Another way in which our approach is unique is that it does not limit users with a given reuse factor to all lie in a circle centered about the base station, as is the case with reuse partitioning. The example in the next section will illustrate this characteristic. In general there are no constraints on the location of users with a given reuse factor. They do not have to lie within any particular region. However, in order to achieve better throughput, users with a given reuse factor will typically be grouped by area, although it is also possible to take into account a user's rate and nature (bursty or not) of packet transmission as well as the Quality of Service it may have been guaranteed.

At the present time, we do not try to optimally assign reuse distances. This problem is essentially equivalent to optimally scheduling all packet transmissions, which is hard to implement in real-time. Consequently, the pursuit of good heuristic algorithms will be of paramount importance. In the next section we present one such algorithm. It assigns reuse factors to mobiles based upon their location in the plane. Regions of the plane, which we call "virtual cells," will be defined such that all users in a given virtual cell will have the same reuse factor. The system throughput achieved is considerably greater than that achieved with the one-seventh frequency reuse plan.

IV. The Protocol

For simplicity of analysis, we will assume throughout that all system cells are hexagonal and have unit side length. In order to facilitate description of the protocol, we define a set of virtual cells as a number of equally sized areas in the system such that one user in each virtual cell can transmit simultaneously and all can be correctly received. Our first example

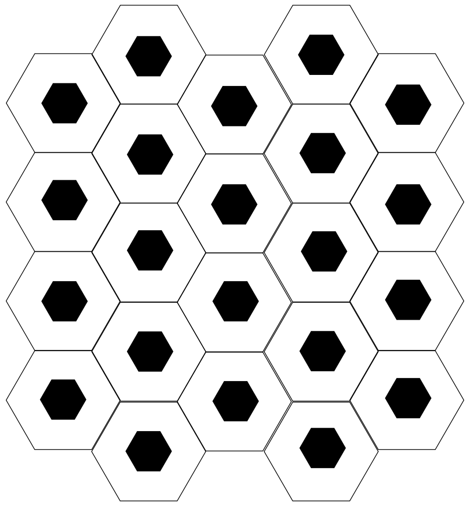

Fig. 3. Illustration of the first level of virtual cells (dark regions). One user from each may simultaneously transmit without interference.

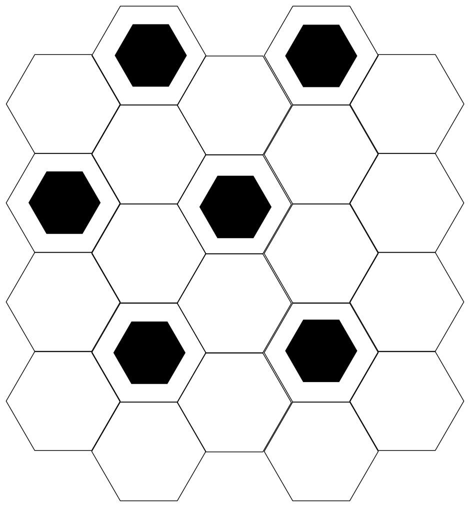

Fig. 4. Second level virtual cells. Users in each have a reuse factor of 3.

of a set of virtual cells will be constructed by shrinking all system cells down to a size such that one user per (smaller) virtual cell can transmit simultaneously in all system cells. See fig. 3 for a picture of the first level of virtual cells. When users inside the first level virtual cells are transmitting, their base stations are at a distance $\sqrt{3}$. Under the one-seventh frequency reuse plan, base stations using the same set of frequencies are at distance $\sqrt{21}$, since we have assumed unit side length. In order to achieve the same performance as the one-seventh frequency reuse plan, we need the radius,r_1, of the first level virtual cells to be such that:

$$\frac{1}{\sqrt{21}} = \frac{r_1}{\sqrt{3}}.$$

Therefore, $r_1 = \sqrt{1/7}$.

This process is then repeated to construct a second level of virtual cells with a reuse factor of 3. That is, every third system cell contains a second level virtual cell such that one user in every second level virtual cell can transmit simultaneously (see fig. 4).

It is well known that to achieve a reuse factor of N with

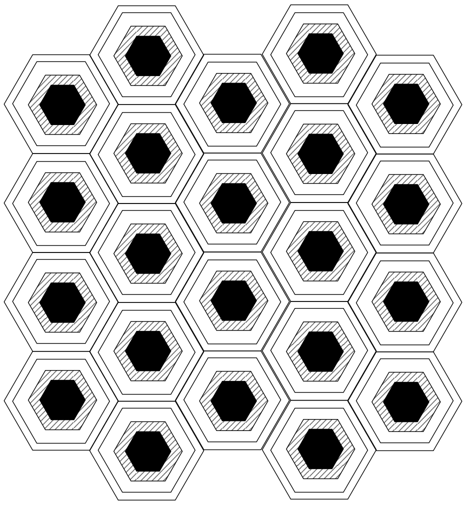

Fig. 5. First three levels of virtual cells.

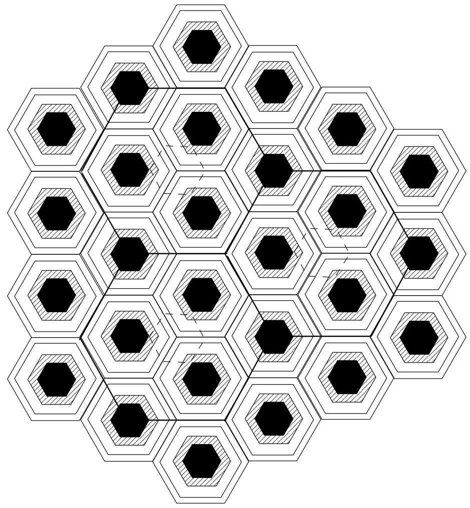

Fig. 6. Construction of second level virtual cells. Each second level cell sits inside an area equal to 4 original system cells.

base stations with unit side length, base stations using the same set of resources are located at distance $\sqrt{3N}$. For reuse factor 3 we have $\sqrt{3*3} = 3$. In order to preserve the same performance as the one-seventh frequency reuse plan, we need the radius, r_2, of the second level virtual cells to be such that:

$$\frac{1}{\sqrt{21}} = \frac{r_2}{3}.$$

Therefore, $r_2 = \sqrt{3/7}$. The same process is repeated for reuse factor 4 which creates a sequence of concentric hexagons as seen in fig. 5.

If at this point we assigned all points not in any virtual cell a reuse factor of 7 as in the conventional one-seventh frequency reuse plan, we would have a classic instance of reuse partitioning, yet in the time domain using packets rather than in the frequency domain using radio channels. Such an approach, in the case of circuit switching, can be found in the literature on reuse partitioning [4, 7]. In our case, using packets would be an improvement due to statistical multiplexing. However, we can do even better by forming additional sets of virtual cells. If we were to consider capture there would be even more flexibility in using additional areas with smaller reuse distances than 7. At present we use only the interference assumptions of the one-seventh reuse plan when assigning reuse distances to virtual cells. This conservative approach illustrates the power of PCMA without introducing analytical complexities. Since we are not yet exploring capture effects, we limit the current example of PCMA to having one additional set of virtual cells that we will call fourth level virtual cells. We want this new set of virtual cells to cover as many new users as possible (those that would have had reuse factor 7 in classic reuse partitioning), and we want the new set to have the minimum reuse factor possible.

Looking at fig. 5 one can see that in order to maximize the number of new users covered, a new set of virtual cells should be interwoven with the first sets, allowing the new virtual cells to cover as much uncovered area as possible. The new set of virtual cells is indicated in fig. 6.

They can be viewed as sub cells of larger hexagonal areas (those indicated by the thicker black lines) whose side

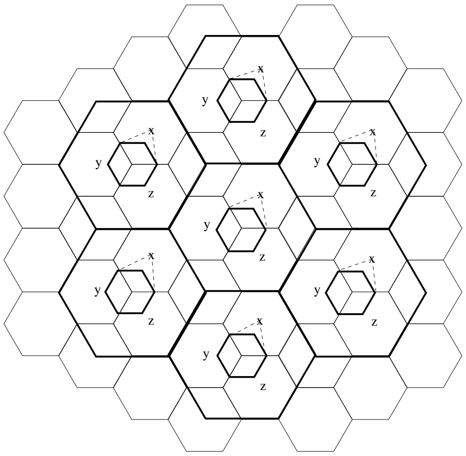

Fig. 7. 120 degree sectorized antennas in each cell allow the second level virtual cells to be induced without additional base station placement.

lengths are twice that of an original system cell, namely 2. From the earlier calculation for first level virtual cells, we know that we can form a set of virtual cells by taking a sub-cell of radius $1/\sqrt{7}$ that of the larger cell. In this case we have $2 * 1/\sqrt{7} = 2/\sqrt{7}$. As in the first level virtual cells, one user in each of the new fourth level virtual cells can transmit simultaneously without interference. The fact that there is no base station at the center of the new virtual cells is not a problem. The three base stations surrounding a given virtual cell "induce" the new virtual cell by using sectorized antennas. Under this arrangement, each surrounding base station must use one of its sectorized antennas to form the new virtual cell. Using antennas with just 3 sectors allows the base stations to service the new virtual cells such that the interference between users in the virtual cells is below the threshold required under the one-seventh frequency reuse plan. This can be seen in fig. 7.

To remain below the interference threshold, it is necessary to coordinate the transmissions among the sectorized antennas, but this is a simple matter. In the case of fig. 7, all sectors labeled "x" transmit together, all "y" sectors transmit together in different slots, and similarly for "z" sectors.

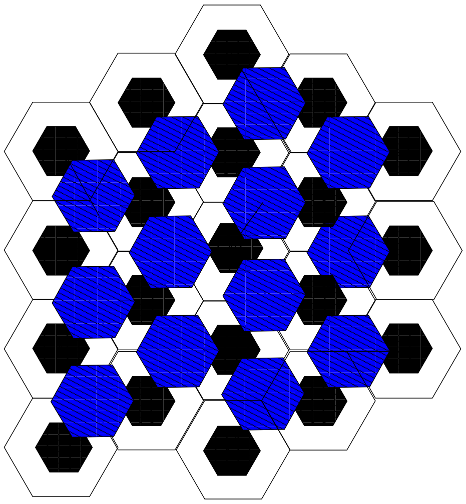

Fig. 8. The result of repeating the coverage pattern of the fourth level virtual cells a total of 4 times. The first three levels of virtual cells are contained inside the smaller hexagons.

The fact that each fourth level virtual cell sits inside an area equal to 4 system cells implies that users in fourth level virtual cells have reuse factor 4. Looking at fig. 6 one can see that 4 such sets of fourth level virtual cells are necessary to create the coverage pattern in fig. 8. This clearly implies a reuse factor of 4.

Finally, we form 5th level virtual cells according to the one-seventh frequency reuse plan. In every 7th system cell, all users not covered by any of the first four levels of virtual cells make up a set of 5th level virtual cells, which have reuse factor equal to 7.

In order to determine system throughput under this arrangement, we need to compute expression (1). Assuming uniform traffic intensity, $\lambda(x)$ is constant in both numerator and denominator and can be factored out. Then the numerator is simply the area of a system cell, namely $A = 3\sqrt{3}/2$. The integral in the denominator is easy to compute, since points with the same reuse factor are grouped into regions in the plane. In this case the denominator simplifies to $\sum_{i=1}^{5} a_i r_i$, where a_i is the total area of all ith level virtual cells in a system cell and r_i is the reuse factor for virtual cells of level i. Thus, to find throughput, the sizes of the different virtual cells must be determined.

The first level virtual cells, those with a reuse factor of 1, were found to be hexagons of side length $\sqrt{1/7}$. Thus, by simple geometry they have area $A/7$. The next virtual cells, those with reuse factor 3, have side length $\sqrt{3/7}$. Thus, they would have area $(3/7) * A$. However, each contains an entire first level virtual cell, and that area is already being serviced. Consequently, the second level virtual cells are actually hexagonal rings of size $(3/7)*A-(1/7)*A = (2/7)*A$. The virtual cells with reuse factor 4 have side length $\sqrt{4/7}$, so they would have size $(4/7)*A$. Again, however, they contain first and second level virtual cells, so they actually are hexagonal rings of area $(4/7) * A - (3/7) * A = (1/7) * A$.

The final computation is to determine the size of the fourth level virtual cells, those that are interwoven with the first 3 levels of virtual cells. Each of these new cells sits inside a hexagonal area with side length two times that of the system cells. Therefore, we can use the earlier calculation for the size of the first level of virtual cells if we simply double the side length ($s = 2$). We find that the radius should be $2/\sqrt{7}$. Consequently, the area is $(2/\sqrt{7})^2 * A = (4/7) * A$. Notice, however, that the new virtual cells cover some of the area already covered by the first few levels of virtual cells. After subtracting this extra coverage, each new virtual cell still covers a previously uncovered area equal to 0.31*A. In a given system cell we have 3 fourth level virtual cells that overlap with it, but only 1/3 of the area of each is actually contained in the system cell. Therefore, we have an area covered equivalent to one whole fourth level virtual cell, or 0.31*A. The remaining uncovered area is the fifth level virtual cell, which is computed by taking the area of an original system cell and subtracting the sum of the areas of the first three levels of virtual cells, $(4/7) * A$, and the area of the fourth level virtual cells, 0.31*A. Thus, we have $A - (4/7) * A - 031 * A$.

Expression (1) thus becomes:

$$\frac{A}{\frac{A}{7} + \frac{2A}{7}(3) + \frac{A}{7}(4) + 0.31 * A(4) + (A - \frac{4A}{7} - 0.31 * A)(7)}$$

which simplifies to 0.2745 In other words, on a per cell basis, we are making productive use of over 27% of the total available time slots in the system.

In order to facilitate comparison, we note that on a per cell basis, the one-seventh frequency reuse plan only makes available $1/7 = 14.29\%$ of the total available time lots in the system. After all, each cell is only permitted to transmit during the $1/7$ of the time slots reserved for its use. Therefore, the throughput of even this early version of our proposed technique is $27.45/14.29 \simeq 1.92$ times greater than the one-seventh reuse plan.

A. Non-uniform Traffic

In most papers on wireless multi-access protocols, uniform traffic distributions are assumed (often in order to simplify analysis). However, in real systems such an assumption is not likely to be realized. We have made a uniform traffic assumption above, but we would like to indicate here how our proposed protocol might operate under non-uniform traffic conditions.

Specifically, we have in mind the system proposed by the Center for Wireless Communications at the University of California, San Diego for Universal Broadband Wireless coverage to the home. The proposal includes base stations mounted on top of homes that would provide Mb/s speeds through wireless radio communications with users in the home or outside, including passing pedestrian and vehicular traffic. The base stations communicate with each other through optical links, thus forming a network mesh that is connected directly to the wireline network at several points. We will not discuss the details of the above proposal any further, but we do note that in such a system, most of the traffic would be generated from inside the home, for example from multimedia applications running on a desktop PC. PCMA can take great advantage of such non-uniformity in traffic. As a concrete example, let us assume that 50% of all traffic is generated from within homes and that all homes fall within first level virtual cells. We further assume that all other traffic is generated from outside the first level virtual cells, but is uniformly distributed outside the first level virtual cells. Note that this is being quite generous since it is likely that even more traffic would be generated from within the home. Also, some traffic generated outside the home might actually fall within the virtual cells (like pedestrian traffic on a sidewalk just outside a home).

Assume that $\lambda(x) = C$ for all points outside the first level virtual cells. Then $\lambda(x) = 6C$ for all points inside the first level virtual cells. This follows since the first level virtual cells are $1/6$ the size of the area not in first level virtual cells and we are trying to achieve an equal balance of traffic. We use equation (1) again, but this time plug in the new values for $\lambda(x)$. It follows that system throughput is over 39%, which is 2.75 times that of the one seventh frequency reuse plan. The advantage of our scheme comes from the fact that the relative distribution of time slots allocated to various virtual cells can be dynamically changed in real time in order to adjust to varying traffic distributions.

V. Discussion

Thus far we have presented PCMA in its simplest form. There is nothing complex or dynamic about the first example of the PCMA technique described in the previous section, yet it is almost two times more efficient than the one-seventh frequency reuse plan. At present we make no dynamic adjustments based on environment conditions or traffic distributions, yet the previous section also indicated how such real-time adjustments could dramatically enhance PCMA. Furthermore, we have not taken advantage of the possibility of packet capture and have simply used the inter-cell interference constraints of the one-seventh reuse plan in order to avoid inter-cell interference in the area outside virtual cells. In so doing, we have made it easier to evaluate the efficiency of PCMA as a strategy for developing MAC protocols for wireless environments. By using only a few levels of virtual cells we could achieve system throughput almost two times greater than the conventional one-seventh reuse plan. This figure represents a substantial capacity gain. Consequently, further research is needed into enhancements of PCMA. For one thing, a study needs to be done that makes provision for packet capture. Capture should allow the various levels of virtual cells to be slightly larger, which would improve throughput. Making virtual cells with low reuse factor larger would allow more users to be covered with a lower reuse factor, thus increasing system throughput. Further study is also needed of new means for creating fourth level and higher virtual cells. It is not necessary to form them inside of groups of 3 system cells as in fig. 6. Alternate coverage patterns need to be investigated. Finally, if several plane coverage patterns are found that achieve similar throughput, then different ones could be used at different times depending on the distribution of users throughout the system. This would allow users to be assigned as low a reuse factor as possible, regardless of the current traffic distribution. As a result, the greatest possible system throughput could always be realized, even under constantly varying traffic distributions.

VI. Conclusion

We have presented a new media access control strategy called Plane Cover Multiple Access. It is based on a packet-switched wireless cellular system, the style of system that will be needed to efficiently support the mix of data and voice transmission that will be required of the next generation of wireless networks. It operates by inducing various sets of "virtual cells" in the plane and then exploiting the possibility of simultaneous transmissions among these cells to enhance system capacity over that of the one-seventh frequency reuse plan that has been adopted for most existing wireless systems. By using several levels of virtual cells of different sizes, we can achieve system throughput almost 2 times greater than the one-seventh reuse plan. This increase is dramatic, considering that we are strictly adhering to the rules of inter-cell interference as dictated by the one-seventh plan. In particular, we are making no provision for packet capture.

References

[1] D.J. Goodman, R.A. Valenzuela, K.T. Gayliard, and B. Ramamurthi, "Packet Reservation Multiple Access for Local Wireless Communications," IEEE Transactions on Communications, Vol. 37, No. 8, August 1989, pp. 885-90.

[2] K. S. Gilhousen, I. M. Jacobs, R. Padovani, A. J. Viterbi, L. A. Weaver, Jr, and C. E. Wheatley III, "On the capacity of a cellular CDMA system," *IEEE Trans. on Vehic. Technol.*, vol. 40, no. 2, pp. 303-312, May 1991.

[3] F. Borgonovo, M. Zorzi, L. Fratta, V. Trecordi, and G. Bianchi, "Capture-Division Packet Access for Wireless Personal Communications," IEEE Journal on Selected Areas in Communications, Vol. 14, No. 4, May 1996, pp. 609-622.

[4] S. W. Halpern, "Reuse Partitioning in Cellular Systems," IEEE Trans. on Vehicular Tech., 1983, pp. 322-27.

[5] M. Karol, Z. Liu, and K. Eng, "Distributed-Queueing Request Update Multiple Access(DQRUMA) for Wireless Packet(ATM) Networks," ICC '95, June 1995, pp. 1224-1231.

[6] G. Bianchi, F. Borgonovo, L. Fratta, L. Musumeci, and M. Zorzi, "C-PRMA: a Centralized Packet Reservation Multiple Access for Local Wireless Communications," IEEE Transactions on Vehicular Technology, Vol. 46, No. 2, May 1997, pp. 422-36.

[7] S. Papavassiliou, L. Tassiulas, and P. Tandon, "Meeting QOS Requirements in a Cellular Network with Reuse Partitioning," IEEE INFOCOM '94, June 1994, pp. 4-12.