Chapter 7

Satellite Channel Assignment

FPLMTS will need to use spectrally efficient channel assignment processes to meet anticipated demands for capacity and global coverage at reasonable cost. This chapter introduces channel assignment techniques [KATZELA] and terminology and the unique difficulties of channel assignment in satellite FPLMTS. The merits of FDMA, TDMA and CDMA are highlighted. The remainder of the chapter investigates techniques to make FDMA and TDMA channel assignment process more efficient.

7.1. Satellite FPLMTS and Cellular Radio

Satellite FPLMTS networks will use a technique known as "frequency re-use" that has been developed in terrestrial cellular telephone networks. Whilst in use, a customer's mobile telephone occupies a radio channel that must not be used by any other customers' terminals within a certain range otherwise the telephones would interfere with each other. In cellular radio networks low-powered base station transmitters are sited in a cellular structure and each telephone communicates only with the closest base station at the minimum possible power level. As the radio signal travels outwards from the transmitter, its power is dispersed and attenuated by a factor of roughly $r^{-3.5}$, where r is the distance from the transmitter. For a large enough r, the power flux density becomes sufficiently low that it does not cause interference with another customer's telephone even if both telephones use the same radio channel. This distance is known as the minimum frequency re-use distance.

A range of channels is assigned to each cell from the pool of all possible radio channels in the licensed frequency band so cells that are further apart than the minimum frequency re-use distance can use the same channels, whilst cells closer together do not. Over the entire cellular network, channels will simultaneously be re-used many times. The penalty for this spectral efficiency is that any mobile customer's terminal must only use channels that have been assigned to its current cell and it must therefore be able to hand the call over from one cell's channel to another's each time it crosses the boundary between cells. This handover process involves the automatic re-routing of the call to the new cell site, the assignment of one of the new cell-site's channels to the mobile and the re-tuning of the mobile telephone to that new channel, all without the customer's call being disrupted.

Mobile satellite systems use a very much longer radio path that has different characteristics to terrestrial cellular radio but certain similarities can be noted. A combination of low altitude satellites and large, multi-beam antennas can mimic cellular coverage. Handover can be introduced to continue calls as customers move through the spot beam pattern or, considering the high speeds at which the satellites move, the satellites move over the customers. The relevant distinguishing feature of satellites is that the minimum frequency re-use distance is not determined by an empirical $r^{-3.5}$ power law but by the position of the mobile terminal in a satellite antenna gain pattern. With antennas of a size practical for satellite service within the next decade, power 'roll-off' moving away from the centre of a satellite spot beam is much less than a terrestrial cellular radio cell's $r^{-3.5}$. The a reflector antenna spot beam's gain pattern always has parabolic roll-off at θ^{-2} from the centre of each antenna feed in the reflector system. Using phased arrays of feeds can produce more directive beams with better out-of-beam attenuation but satellite spot beam areas and frequency re-use distances will be larger than their terrestrial counterparts, even with the very large phased array antenna systems currently proposed.

7.2. FDMA, TDMA and CDMA

The principles of CDMA rest on the work of Claude Shannon (1948) and V. Kotelnikov (1947). Shannon showed that performance was related to the time × bandwidth × power product of the signal. Kotelnikov showed that in white noise the detection performance depends only on received signal energy. Radio communications systems can be considered to use three finite resources: frequency, time and power. A service will be licensed to use a range of frequencies within which it must be confined to avoid interference to other services occupying frequency bands around it. The usable power domain is delimited by the minimum signal power required above thermal noise to allow signal recovery and the maximum transmitted power from the transmitter, which is limited by power supply, equipment and safety considerations. Usually the whole of the time domain can be used, with sharing planned to avoid any unacceptable transmission delays. Figure 33 shows these three radio resources. TDMA would slice up this cube horizontally, FDMA vertically through the frequency axis. CDMA would in effect slice the cube vertically through the power axis, but the power usage of each channel might change rapidly with time.

Since techniques exist that make our use of frequency, time and power interchangeable, looking at figure 33 it seems that TDMA, FDMA and CDMA should all have the same maximum capacity. What follows is not a comparison of TDMA, FDMA and CDMA to determine which offers the most benefit to satellite FPLMTS but is an introduction to the factors that effect cellular network capacity and their relevance to a satellite channel assignment scheme.

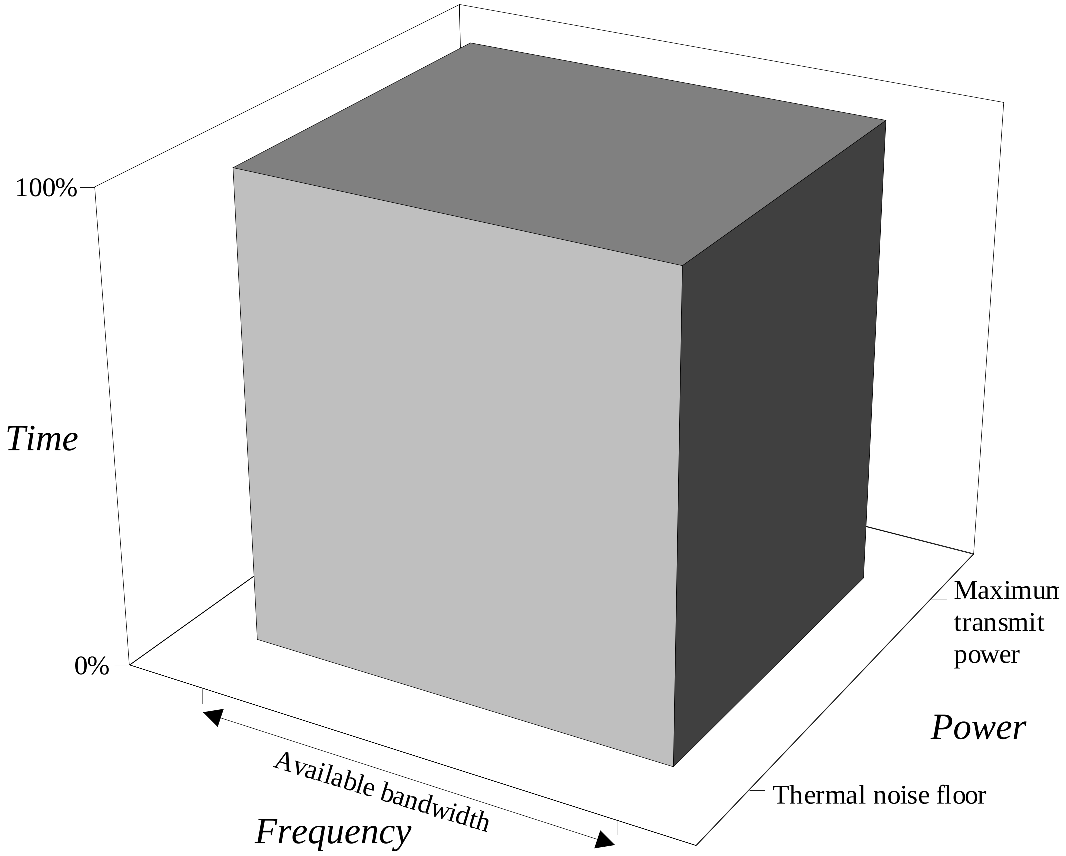

Figure 33 Time, frequency and power radio resource

7.2.1. The Importance of Power

A cellular radio system is designed to re-use its time, frequency and power radio resources across a spatial domain. Whilst a FDMA system appears to re-use frequencies and a TDMA system appears to re-use time slots neither the frequency nor timing of a signal change significantly as it travels through space. What does change is the signal's power, so it is the power resource that is really being re-used. CDMA re-uses the power resource directly without having to slice up the frequency or time domains in such a regimented fashion as TDMA or FDMA. Propagation measurements show that the general trend for signal strength roll-off is proportional to the distance from the cell centre but that there are local variations due to shadowing from obstructions and multi-path fading. In a mobile radio environment the depth of these signal strength fades also varies with time, further complicating the challenge of making maximum use of the radio resources. It might be helpful to think of the resource cube in figure 33 as containing moving "bubbles" where the resources are being under used. Efficient channel assignment would fill these "bubbles" in with more traffic, increasing the total capacity of the system.

7.2.2. CDMA

CDMA uses the carrier's modulation to distinguish between channels. FDMA transmitters modulate data onto very narrow band (sinusoidal) carriers at different frequencies so receivers can identify their intended data from its carrier frequency. CDMA transmitters modulate data onto wide band carriers that are distinguishable from each other by different preset PN (pseudo noise) sequences that they follow in time. All of these different carriers share the same band of frequencies and receivers can identify their intended data by searching for their PN sequence.

If we wanted to avoid self-interference completely, so that one traffic channel does not interfere with any other traffic channel, all traffic channel carriers must be spread with synchronized, orthogonal PN sequences. The number of orthogonal codes available happens to be equal to the spreading factor, which is the spread in signal bandwidth on the wide band carrier compared with a narrow band FDMA carrier. If the spread bandwidth of the carrier was 128 times the baseband bandwidth then there would be 128 orthogonal PN sequences that could be used to support 128 traffic channels. The capacity of a single-cell CDMA system is therefore exactly the same as the theoretical capacity of a single-cell FDMA or TDMA system. This is no surprise because a synchronous orthogonal CDMA system operates in a fashion analogous to a very fast TDMA system (direct sequence CDMA) or a rapidly changing dynamic FDMA plan (frequency hopping CDMA).

CDMA carriers used in the cellular radio environment do not use orthogonal PN sequences, so they cause self-interference to other carriers using the system. To overcome this, slightly more power than the minimum must be used to compensate for the parts of the signal that are being interfered with and corrupted, reducing the total number of carriers that can be used. CDMA’s de-spreading process averages the signal out to recover the intended signal and rejects the interference as long as the intended signal is sufficiently more powerful than the interference. Although the use of non-orthogonal PN sequences causes self-interference and reduces system capacity by requiring the use of extra power to overcome its own interference, it does provide some advantages.

If orthogonality is not required for the chosen PN sequences, then the number of PN sequences that can be chosen from is vast. Sets of low cross-correlation codes such as the Gold, Kasami and Bent sequences have been developed and their lengths can be longer than the spreading factor allowing the number of codes in a set to be so huge that an unused sequence can be picked very simply by seeding the sequence generator with the time and date, for example. Furthermore, when orthogonality is not required between PN sequences then it is no longer necessary to synchronize the sequences. In cellular systems where cell antenna sites are spread over tens of kilometres synchronization of these very fast PN sequences would complicate the system.

Use of the wider set of all PN sequences allows unique PN codes to be allowed for many more Traffic Channels than orthogonal PN sequences, TDMA or FDMA would allow. Under ideal propagation conditions self-interference would cause inefficient use of the power resource, reducing the maximum number of Traffic Channels that could be handled before interference would become intolerable. Fortunately the multi-cell mobile radio environment is not ideal and shadowing and fading increase the attenuation of interfering signals compared to ideal propagation. Muting of carriers during the pauses in customer’s speech and variable bit-rate data channels also reduce interfering signals, as does the use of tight closed loop power control where the minimum power is transmitted to enable the intended receiver to just decode the signal. The combined use of all of these techniques dramatically reduces the self interference compared to what would be expected based on a simple, ideal propagation model. In practice it is found that in a cellular system, a larger number of Traffic Channels can be handled than would be available using orthogonal PN sequences, TDMA or FDMA. The number of non-orthogonal PN sequences available is more than enough to allow this large number of Traffic Channels to be individually identified and decoded, allowing their simultaneous use.

7.2.3. Advantage of CDMA

It is recognized that CDMA's capacity gains over TDMA and FDMA are entirely due to its tighter, dynamic control over the use of the power domain. Because of the ease of choosing a new non-orthogonal PN sequence a CDMA system does not encounter the difficulties of choosing a spare carrier frequency or time slot to carry a Traffic Channel, it simply needs to ensure that interference will not be too great if it begins to transmit - that there is still enough space left in the power domain for it to use.

Because its PN sequence is unique and not being re-used anywhere else in the system the mobile can move into other cells, taking its PN sequence with it without fear of meeting interference from another mobile using it. The system requires only that the total interference power from other mobiles is not too much to swamp detection of the intended data. This allows simpler radio resource sharing than FDMA where carrier frequencies cannot be exported with a mobile into another cell in case it interfered with a mobile re-using the frequency in a neighbouring cell. Freedom from tight frequency co-ordination also eases resource sharing in the presence of large Doppler frequency shifts.

7.2.4. Disadvantages of CDMA in the Satellite Environment

In the satellite radio environment CDMA is less attractive than it is in the terrestrial cellular radio environment for the following reasons.

- Power control cannot be as tight as it is in a terrestrial system because of long round-trip delay. Instead a fade margin must be included to fill in Rayleigh and log-normal fades, increasing self-interference (see section 2.5.5).

- Satellite transponders are channelized too narrowly for Broadband CDMA, which is the most attractive form of CDMA. Broadband CDMA reduces the depth of fades and so reduces the reliance of CDMA on power control. Fading is frequency selective - a narrow band signal will appear to suffer from temporal Rayleigh fading but if a signal is spread over a sufficiently large portion of the spectrum then, whilst parts of the signal will suffer Rayleigh fading, the average depth of fade over the whole bandwidth is much reduced. In fact, broadband fading is better described by the Rice distribution than the Rayleigh distribution. CDMA's de-spreading process's averaging effect can therefore be used to reduce the depth of temporal fades and reduce the fade margin that would be required for satellite communications. To do this successfully the spread bandwidth needs to be an order of magnitude greater than the radio channel's coherence bandwidth - it must be more than wide enough to encompass even the widest of specular fades. Terrestrial 2GHz FPLMTS band micro-cellular system measurements in urban areas indicate that for terrestrial use the coherence bandwidth is about 200kHz - 300kHz but measurements on satellite channels indicate that for satellite use the coherence bandwidth is closer to 10MHz.

This requires spread bandwidths of approximately 100MHz to benefit from the elimination of Rayleigh fading. Unfortunately, channel bandwidths of this size are not possible on satellites for technical and regulatory reasons. Although 72MHz transparent transponders are commonplace, the digital spot-beam forming networks that are anticipated for FPLMTS are currently limited to processing bandwidths of about 30MHz and, more importantly, 30MHz is the bandwidth of the entire FPLMTS Mobile Satellite Service allocation. Perhaps spreading a signal to 30MHz is feasible, which would reduce temporal fade depths significantly, though not by as much as a 100MHz spread bandwidth would. CDMA satellite systems such as Globalstar and Odyssey will use narrowband CDMA with signals spread to only 1.25MHz and 4.83MHz, respectively.

- Satellite transponder power is inefficiently used by relaying thermal noise in such large bandwidths.

7.2.5. Improving FDMA and TDMA

CDMA shows that there are significant capacity gains to be made by dynamically sharing out the power/frequency/time resource finely along the power axis. TDMA and FDMA cannot achieve this resource sharing as well as CDMA can, but it is possible to dynamically assign TDMA time slots and FDMA frequencies to make better use of the power resource in poor propagation conditions than fixed channel assignments would make. In doing this, one would expect to benefit from some of the potential capacity gains that CDMA shows to exist. In terrestrial cellular networks, this technique is known as DCA and has been studied for many years [NETTLETON, BECK].

7.3. Satellite Channel Assignment

CDMA will be used in some satellite FPLMTS networks because of its simplicity of radio resource sharing. In LEO and MEO constellations the attraction of a multiple access scheme with no need to manage the sharing of a very small number of unique channel identifiers is considerable, especially with the rapid motion of satellites and the constantly changing pattern of interference. Globalstar and Odyssey will both use CDMA. Globalstar's system is very similar to the US IS95 CDMA standard, which may allow the same terminal CDMA functions to be used in both IS95 and Globalstar modes of dual-mode terminals.

Other satellite FPLMTS networks will choose TD/FDMA because of its efficiency on satellite links and because for LEO systems with ISLs like Iridium, the simplicity of multiplexing and switching TD/FDMA carriers on board satellites will be useful.

This chapter describes mechanisms that could be used to assign TD/FDMA Traffic Channels to mobile terminals communicating with FESs through non-GEO satellites. It is supported by simulation work for the baseline 769km LEO scheme used in the link budgets in section 2.5. The simulation is documented in appendices A, B and C and a summary is reported in an IEE colloquium paper [FINEAN]. The mechanisms could also be used with GEO satellite systems as well as non-GEO systems at any altitude. The main feature is providing solutions to the unique non-GEO problems of frequent satellite changes and highly dynamic channel re-use co-ordination. Independently of this work, the University of Firenze, Italy, has been studying a number of dynamic channel allocation schemes for use with LEO satellites, including handover queuing techniques [DEL RE].

7.4. Need for Frequency Guard Bands

One of the biggest inefficiencies in an FDMA system is the guard band required around each channel's occupied bandwidth. This band is to allow for realistic implementations of the filters required to select the intended carrier without inadvertently selecting parts of other carriers. They also allow carriers to drift slightly from their assigned frequency without causing interference with adjacent carriers. Because of the speed of satellite motion, communications to LEOs will inevitably be affected by very high Doppler frequency shifts. Guard bands could be made wide enough to allow such shifts without interference to neighbouring carriers but the inefficiency of spectrum use would be considerable. Fortunately it is possible to track the frequency shifts at the receiver and/or to pre-shift the transmitter carrier frequency to compensate for the Doppler effect. This allows a near constant carrier frequency to be observed at the receiver.

Pre-shifting the transmitter carrier frequency would appear to obviate the need for guard bands but a closer look at the interference that this causes to other receivers indicates that this is not always the case.

7.4.1. Pre-Shifting Downlink Carrier Frequencies

Figure 34 shows a satellite transmitting to the two extremities of a single-beam satellite coverage footprint, with the transmitter carrier frequencies pre-shifted by the frequencies shown.

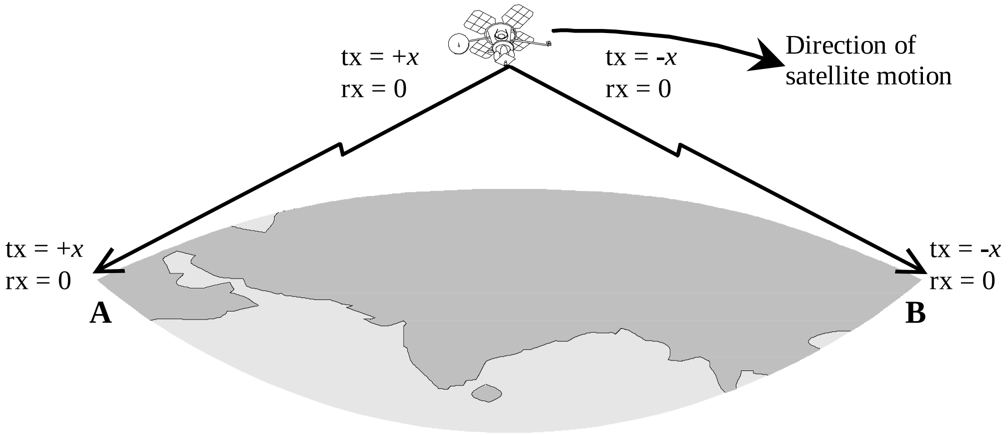

Figure 34 Transmit and receive frequency offsets within a satellite coverage footprint using pre-shifted carriers

All the footprint's downlinks are from one source, the satellite transponder, to many mobile terminals which can be assumed to be stationary on the Earth's surface. Each carrier is pre-shifted individually by the FES to allow the intended recipient to receive the signal with no Doppler frequency shift. This can be achieved by transmitting a constant transmit frequency pilot carrier from the satellite so that each mobile terminal can calculate the frequency pre-shift required for its own position, which can then be used by both the FES and the mobile terminal for their transmissions. These pre-shifts range from -*x* to +*x* kHz, where *x* is the worst case Doppler shift at the edge of the satellite footprint. The drawback is that all mobiles within the beam would receive all the satellite's transmissions, with Doppler shifts appropriate to their own positions, as potential interference. Thus the carrier intended for reception at **A** would also be received at **B** with +*x* kHz of Doppler shift in addition to the +*x* kHz already pre-shifted by the FES. Conversely the carrier intended for **B** would be received at **A** as interference with a frequency offset totalling -2*x* kHz. Fortunately the wanted carrier is always at the correct frequency so, to avoid interference between carriers, guard bands of 2*x* kHz would be needed between satellite downlink carriers. Evidently, since the carriers are being transmitted from the same transponder, shifting them in frequency relative to each other has introduced the need to leave bands into which carriers can be shifted without colliding with other carriers.

If downlink carriers are not pre-shifted then 2*x* kHz of frequency guard band is still required to avoid interference between transmissions from different satellites (the worst case being +*x* kHz Doppler from one and -*x* kHz Doppler from another). However, at any given mobile terminal all the carriers transmitted from any one satellite share the same radio path and therefore have the same Doppler frequency shift. These carriers cannot interfere with each other even if no guard bands are left between them. Therefore if carriers from one satellite could be grouped together to be adjacent in the satellite frequency band, then no frequency guard bands would be required between carriers within the group. Guard bands of 2*x* kHz would be required only between groups of carriers that could be used by different satellites.

7.4.2. Pre-Shifting Uplink Carrier Frequencies

For uplinks to the satellite the pre-shifting technique works better, especially where only one satellite is considered. This is because all transmissions are being pre-shifted in frequency to ensure that there is no residual Doppler offset at one common point, the satellite transponder. In general, this minimizes the need for frequency guard bands on the uplink to the negligible residual Doppler shift that cannot be removed by open or closed loop frequency control. There is a spread of Doppler shifts at any other point in the sky and on the ground but this does not matter as long as there are no receivers elsewhere to interfere with.

Of course in any satellite constellation there are other satellites within range, especially at footprint peripheries. Hence in practice similar problems to those of the downlink are met, with interference between carriers used by different satellites. The worst case parameters would also be the same - a mobile on the edge of two satellites' coverages could experience Doppler frequency shift of -*x* kHz from one satellite and +*x* kHz from the satellite following. But if the mobile was communicating with the former satellite using +*x* kHz of pre-shift then the latter would experience interference at +2*x* kHz, so a guard band of 2*x* kHz is required to avoid conflict with the latter satellite's carrier (which should always be received on target with 0 kHz aggregate Doppler shift). Like the downlink, these guard bands are only required where adjacent carriers are used by different satellites, so grouping carriers together for use by one satellite at a time eliminates the need for guard bands between carriers within the group.

Thus for single-beam satellites we must have no pre-shifting of downlink carrier frequencies and carriers transmitted by different satellites must be separated by guard bands of 2*x* kHz. For uplink carriers, mobile terminals must pre-shift their transmission frequency in response to the movement of the satellite's pilot carrier to compensate for Doppler shifts. Then no guard band will be required between carriers destined for the same satellite but guard bands of 2*x* kHz will be required between carriers destined for different satellites. It is therefore advantageous to co-ordinate the satellite frequency spectrum into groups of carriers, each group used by only one satellite at any given time. Figure 35 illustrates this arrangement and how it reduces the need for guard bands. These provisions will be sufficient for any number of satellites operating over the same geographic area, subject to their being sufficient numbers of groups of satellite channels.

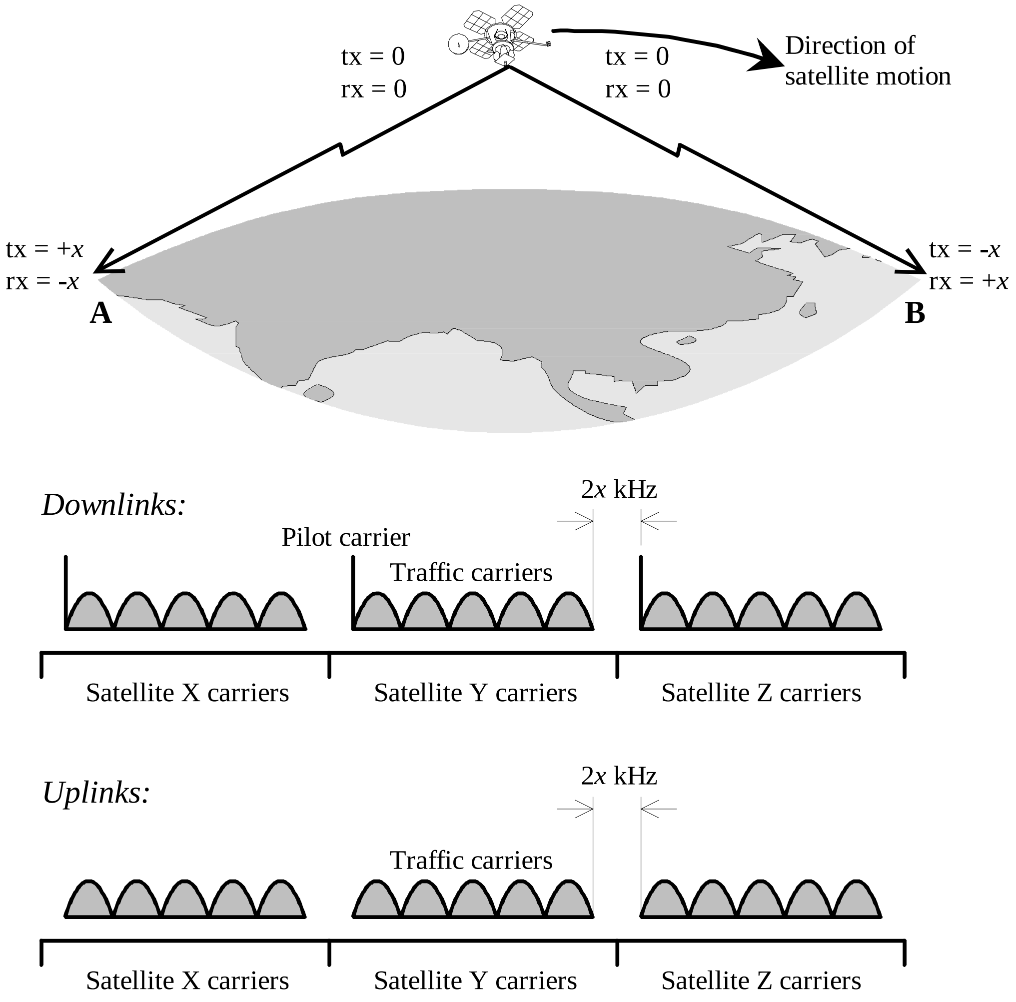

Figure 35   Transmit and receive frequency offsets within a satellite coverage footprint and frequency spectrum usage at the satellite transponder using individually pre-shifted uplink carriers but no pre-shifting on downlink carriers

7.4.3. Satellite Diversity

To use satellite diversity the mobile's uplink is by definition destined for two or more satellites, to be combined at the FES. This prevents us from eliminating the uplink guard bands between carriers destined for the same satellite because depending on the signal that the FES chooses, the carriers may effectively be routed through different satellites. Therefore when satellite diversity is in use uplink Doppler compensation becomes redundant for uplink carriers and the full $2x$ kHz guard band must be left between all carriers.

In TDMA and FDMA systems using satellite diversity, only one of the possible downlink routes can be used at a time. Which satellite to use is determined by the FES's measurements of which satellite relayed the best uplink in the most recent uplink burst. The FES transmits the downlink burst only through this satellite, avoiding potentially destructive interference to the strongest signal from weaker ones arriving out of phase because of different path delays and Doppler effects. Because only one downlink carrier is active at a time, downlink guard bands are not required between carriers on the same satellite, as discussed in 7.4.1. Note that the downlinks relayed by different satellites cannot share the same frequency because their frequencies are not co-ordinated to remove Doppler shifts. Also, burst timing is not co-ordinated to accurately avoid burst collisions when bursts are relayed through different satellites with different radio path lengths.

7.4.4. Spot Beams

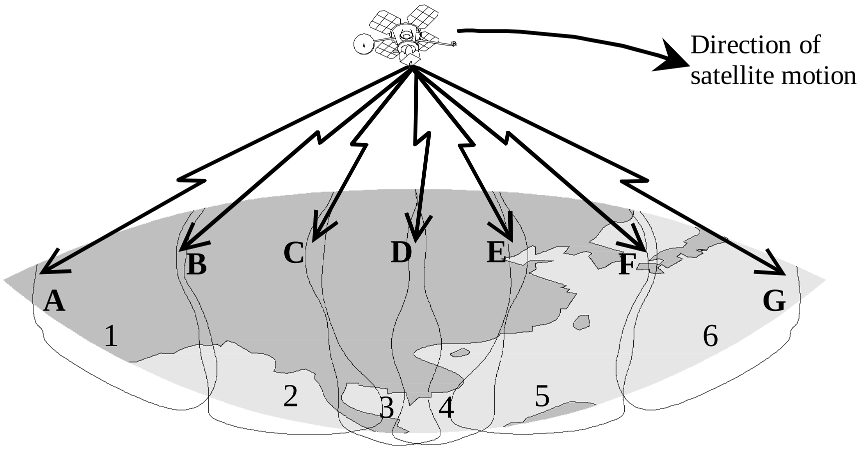

Figure 36 Satellite footprint split into spot beams with downlink carrier pre-shifting constant within each spot beam

Downlinks:		**Residual Doppler** at receiving mobile terminal:						
Transmit beam	**Pre-shift**	**A**	**B**	**C**	**D**	**E**	**F**	**G**
1	$+2\frac{1}{2}y$	$-\frac{1}{2}y$	$+\frac{1}{2}y$					$-\frac{1}{2}y$ *
2	$+1\frac{1}{2}y$		$-\frac{1}{2}y$	$+\frac{1}{2}y$				
3	$+\frac{1}{2}y$			$-\frac{1}{2}y$	$+\frac{1}{2}y$			
4	$-\frac{1}{2}y$				$-\frac{1}{2}y$	$+\frac{1}{2}y$		
5	$-1\frac{1}{2}y$					$-\frac{1}{2}y$	$+\frac{1}{2}y$	
6	$-2\frac{1}{2}y$	$+\frac{1}{2}y$ *					$-\frac{1}{2}y$	$+\frac{1}{2}y$

* Doppler on interference observed by mobile using a different satellite

Table 5 Residual downlink Doppler shifts at mobiles in Figure 36 using constant pre-shifted downlink carriers

Using multiple spot beams in the satellite footprint, as recommended in section 4.4, can help to reduce the Doppler spread that needs to be taken into consideration. Figure 36 depicts a satellite footprint split into six roughly equal-area spot beams where the maximum difference in Doppler shift from one side of a spot beam to another (e.g. **A** to **B** or **B** to **C**) is *y* kHz. Downlink carriers within a spot beam are all pre-shifted in frequency by a constant shift that compensates for the Doppler shift at the centre of the spot beam. Table 5 shows the Doppler shifts of the carriers received by the mobile terminals.

This constant downlink pre-shifting technique works very well, even regarding interference from other satellites, because interference is physically limited by the highly directional antennas of the satellite to a spot beam on the Earth's surface where Doppler shift will be approximately the same. Furthermore, the pre-shifting of carriers is simple to implement at the FES since all the mobile terminals in a spot beam share one constant carrier frequency offset which does not require open or closed-loop control.

Using this technique the frequency guard bands required between downlink carriers from different satellites are reduced to only *y* kHz.

Uplinks:		**Residual Doppler shift** at satellite transponder:					
Transmit mobile	**Pre-shift**	1	2	3	4	5	6
A	+3*y*	0					+6*y* *
B	+2*y*	0	0				
C	+*y*		0	0			
D	0			0	0		
E	-*y*				0	0	
F	-2*y*					0	0
G	-3*y*	-6*y* *					0

* Doppler on interference observed by transponder on another satellite

Table 6 Residual uplink Doppler shifts at satellite transponders in Figure 36 with individually pre-shifted uplink carrier frequencies

The spot beams do not help or hinder the individual pre-shifting of uplink carriers as previously discussed. They cannot reduce the size of the guard bands required between groups of satellite carriers, fundamentally because the mobile terminals' antennas are omni-directional and relatively stationary. Therefore their transmissions, which are pre-shifted for a satellite moving in a particular way relative to the mobile, can be received as interference by all other satellites in the sky, which will be moving differently to the intended receiving satellite and experiencing different Doppler effects. The result is that for all but the intended satellite the pre-shift and the real Doppler shift do not cancel out but could add, demanding the use of frequency guard bands that allow for worst case Doppler shifts across the whole satellite coverage footprint. This is verified in table 6, which shows the need for uplink guard bands between carrier groups of 6*y* kHz (= 2*x* kHz, the same as for the no spot beam case) but only negligible guard bands between carriers in the same group. Note that a group of carriers can be used across the whole satellite coverage, in all spot beams, since they all share the same antenna platform.

7.4.5. Maximum Doppler Spread

To see how big an overhead guard bands of *x* and *y* kHz are on a real system it is necessary to find the maximum Doppler frequency spread across any spot beam, limited by the maximum width of beam in direction of satellite motion.

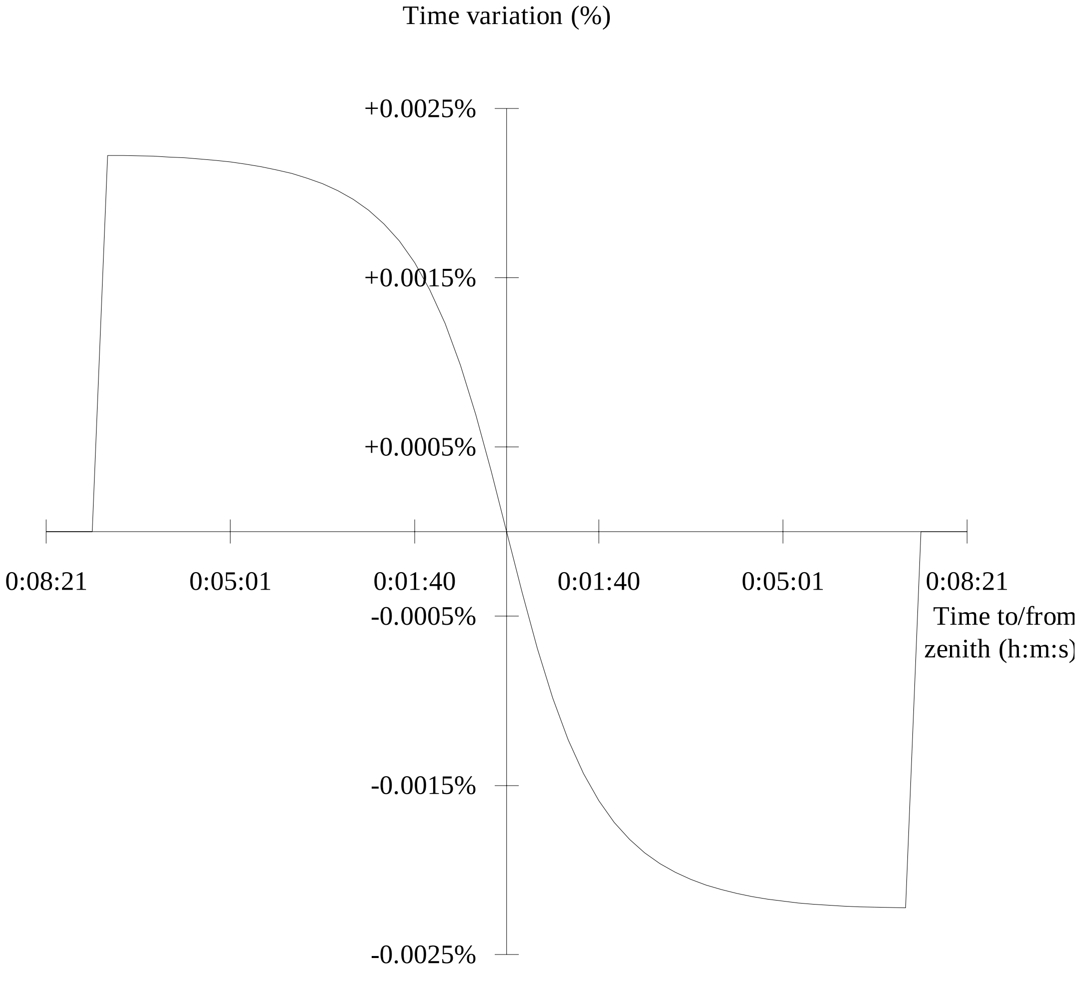

Figure 37 Variation of time at receiver relative to transmitter during pass of LEO satellite at 769km altitude

Assuming that beams may be formed as in figure 19 of chapter 4, the limit on Doppler spread will usually be from the central beams since this is where the Doppler effect is changing most rapidly. Figure 37 shows the difference between the time-frame of the satellite and the time-frame of the mobile terminal as the coverage footprint moves across the mobile terminal. This time variation is defined as

$$\left(\frac{\text{bit rate received (measured at receiver)}}{\text{bit rate transmitted (measured at transmitter)}} - 1\right) \times 100\%.$$

Figure 37's time axis also shows how quickly the satellite passes from the horizon, overhead and to the opposite horizon. Doppler frequency spread is a function of both spot beam size and satellite velocity, both of which vary with orbital altitude. For a specified antenna size (for example one producing a 20° beam width), the combined effect is a monotonic reduction in maximum spread of relative velocity as orbital altitude increases, as shown in figure 38.

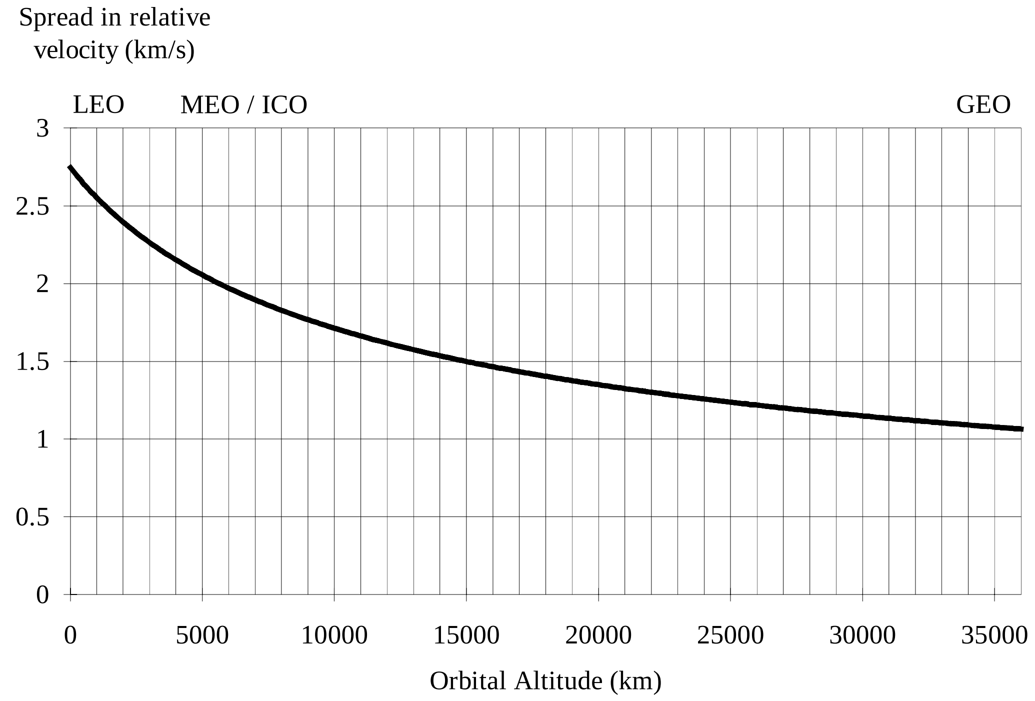

Figure 38 Spread of relative velocity as a function of orbital altitude

The maximum Doppler frequency spread will be for the lowest LEO satellites. Using the baseline 769km altitude LEO as an example, with 20° spot beam widths across their smallest axis, the satellite coverage footprint can be split into 6 spot beams across the footprint diameter in the direction of motion (see section 4.4). This results in a maximum difference of 2.6km/s between the relative velocity of the satellite and a mobile terminal at one edge of the beam and the relative velocity of the satellite and a mobile terminal at the opposite edge of the beam. This difference will be observed in the beams closest to the sub-satellite point.

7.4.6. Size of Guard Bands

The maximum relative velocity of a stationary mobile terminal at the edge of the coverage footprint of a 769km altitude LEO satellite is ±6.6km/s. In the 2GHz satellite FPLMTS frequency band this equates to ±44kHz of Doppler shift on a 2GHz uplink carrier and ±49kHz of Doppler on a 2.2GHz downlink carrier. These are the values of "*x*" in the discussion in sections 7.4.2 and 7.4.1 respectively, so an 88kHz guard band is required between uplink FD/TDMA carriers used by different satellites and 44kHz is required between satellite and terrestrial uplink frequency bands. If frequency re-use can be planned to ensure that large blocks of spectrum will be used exclusively by only one satellite at a time then the number of guard bands required is significantly reduced. Assuming that this can be achieved then guard bands can be eliminated between all uplink carriers on a given satellite transponder by pre-shifting uplink frequencies according to the received frequency of a pilot signal from the satellite. This adds complexity to the mobile terminal but the improvement in spectral efficiency is very worthwhile for TDMA and absolutely essential for the viability of FDMA as multiple access schemes.

Downlink frequency guard bands are reduced by splitting the satellite coverage footprint into spot beams created by the satellite antenna system. The size and shape of the spot beams are determined by the satellite antenna patterns and can be tailored to meet the specific requirements of minimizing Doppler spread across the spot beam. In section 4.4 the worst LEO case, with 6 spot beams, was chosen as a baseline for this thesis. If this is the worst case, the maximum spread in relative velocity across a spot beam in FPLMTS will be 2.6km/s. This corresponds to a downlink Doppler frequency spread of 19kHz. Therefore downlink frequency guard bands of 19kHz should be allowed between FD/TDMA satellite downlink carriers used by different satellites and 9.5kHz between satellite and terrestrial downlink carriers.

For both uplinks and downlinks, a suitably designed CDMA scheme can eliminate the need for frequency guard bands between FPLMTS carriers altogether. Only one guard band is then required, the 9.5kHz guard band between the satellite FPLMTS downlinks and other non-FPLMTS services at 2.2GHz.

Receivers will need to be able to acquire carriers across the ±49kHz range and to track the carrier once in lock. The maximum acceleration relative to a terminal on Earth that can be expected is for the 769km LEO satellite passing directly overhead and is 64m/s/s. As this is an order of magnitude greater than the acceleration of a car, this figure will determine the tracking loop bandwidth. For the 2.2GHz downlinks the loop must be capable of tracking a 470Hz/s change in Doppler frequency shift. To eliminate guard bands between carriers used by the same satellite, the mobile terminal must also be able to measure the received carrier's Doppler frequency shift, calculate the offset required for its transmission and pre-shift its own carrier frequency by the correct amount, which will be in the range ±49kHz. The accuracy of this operation will determine how much residual guard band needs to be left between carriers within the same satellite carrier group.

7.5. Design Objectives

So far this thesis has concluded with the following design objectives that affect channel assignment mechanism design:

- Satellite terminal cost must be low
- Satellite service needs very high availability

- Rapid handover between different satellites in the sky is necessary to evade shadowing because it is unlikely that satellite links would be endowed with sufficient margin to ensure communications through weak reflected signals
- FESs may control multiple satellites and share satellites with other FESs
- To improve spectral efficiency FDMA and TDMA carriers used on the same satellite are best assigned in groups next to each other in the frequency spectrum.

Criteria for assessing the suitability of these handover mechanisms are the spectral efficiency of the channel assignments and the grade of service that can be offered using them, in terms of the probability of a call being dropped and the transparency of handovers.

7.6. Satellite System Model

To simplify discussion the satellite FPLMTS network is defined with functional channels rather than physical radio channels. In a practical implementation, some of these functional channels may be combined or further sub-divided and the physical radio channels would be separated from each other in one or more of the time, frequency and code domains. The four functional channels used in this description of traffic channel assignment are:

- Paging Channel, broadcast by the FES to all mobile terminals in idle mode within a satellite spot beam. It contains a variety of information including a list of mobiles which have incoming calls, a pointer to the Access Channel, Traffic Channel assignment information and the FES's location area identity

- Access Channel, a random access uplink channel monitored by the FES for mobile terminals to respond to Paging Channel pages, request outgoing calls and notify the network of its changing location;

- Downlink Traffic Channel, a data bearer carrying a customer's traffic from the FES to the mobile terminal;

- Uplink Traffic Channel, the corresponding data bearer to the Downlink Traffic Channel carrying the customer's traffic from the mobile terminal to the FES.

The Uplink and Downlink Traffic Channels are assigned on demand as a Traffic Channel Pair from a Traffic Channel Pool, which is radio channel resource for many Traffic Channels. Traffic Channel Pairs are held for the duration of the channel assignment (including during speech pauses) and thereafter are returned to the Traffic Channel Pool. Traffic Channels are fixed bit-rate for the duration of their assignment.

Paging Channels and Access Channels are permanently held by an FES.

7.7. Fixed Frequency Re-use Planning

There are a number of approaches to selecting traffic channels. Channel assignment within first and second generation terrestrial macro cellular networks is usually done according to a fixed frequency re-use plan where channels are assigned to cell sites according to a fixed channel assignment plan. This is planned when the system is installed, to ensure a sufficiently large minimum frequency re-use distance to guarantee that co-channel interference is always below an acceptable level.

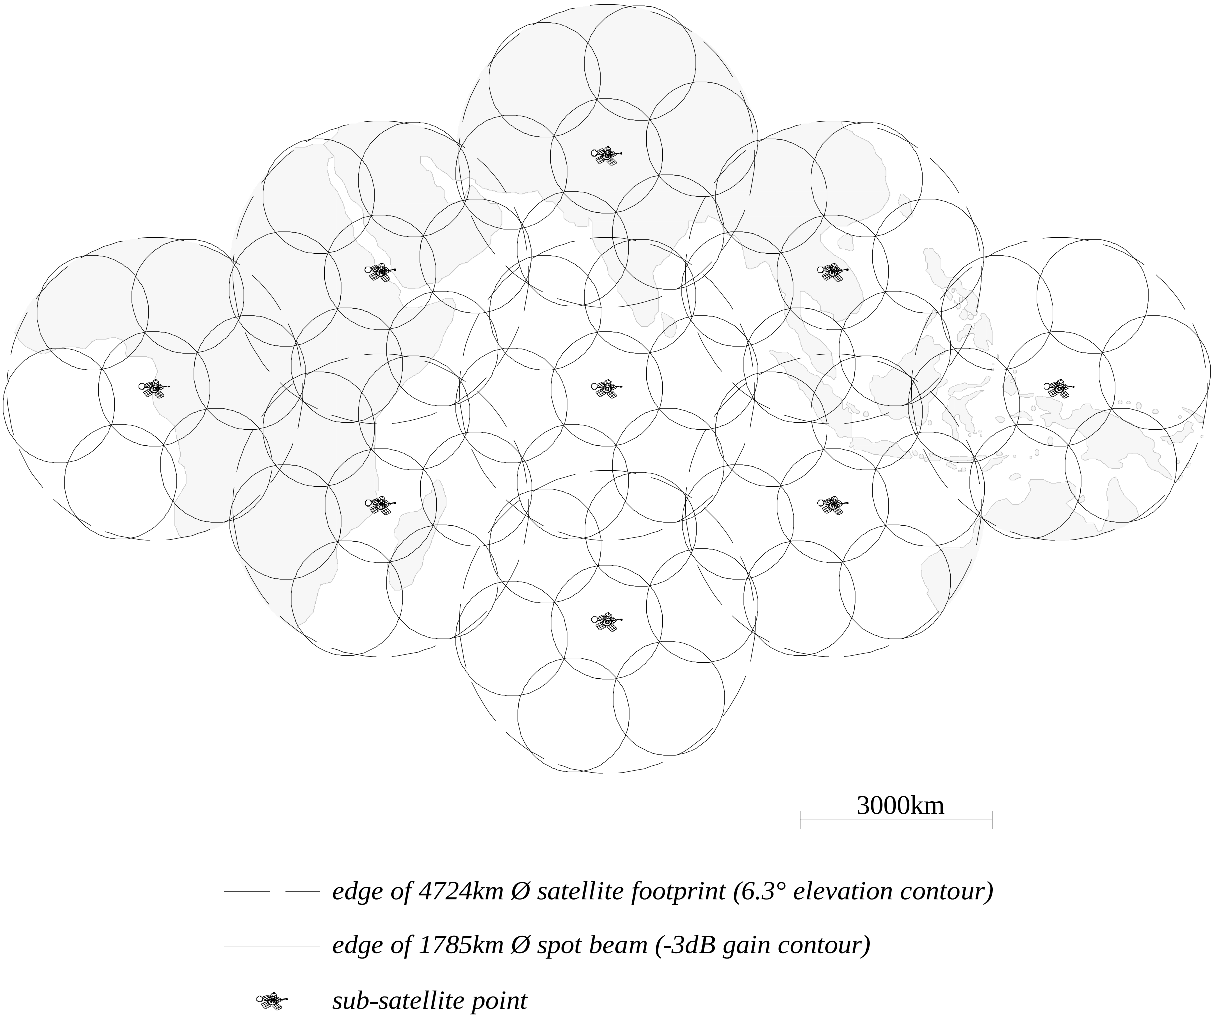

Figure 39 Tessellation of beam patterns over the Indian Ocean at equatorial latitudes

Because of the continual convergence and divergence of satellites in a constellation of non-GEO satellites, the coverage footprints of satellites will often overlap and the area of overlap will be changing all the time. Figures 39, 40 and 41 show an example of this happening for the 769km altitude LEO satellites of figure 15, each with the seven fixed spot beams of figure 18(a). Because it is difficult to use these moving satellite coverage areas as the cellular basis for a fixed frequency reuse plan, a geographically fixed cellular pattern is favoured. Satellites' coverage footprints move over the geographically fixed cellular pattern so satellite spot beams need to take on the characteristics of whichever cells they illuminate as they move into them. The spot beam pattern formed by the satellite's antennas can either be a fixed pattern, in which case a cell's calls will be handed from beam to beam as the satellite moves over the cell, or the spot beams could be steered to each track an individual cell in the coverage footprint. The latter solution ensures that the mobile terminals are located in the centre of the spot beam where the link budget is best and that interference to mobiles in other cells is reduced, although this reduction is so slight that it does not enable a shorter minimum frequency re-use distance to be specified. It also almost eliminates the need for handovers between beams on the same satellite.

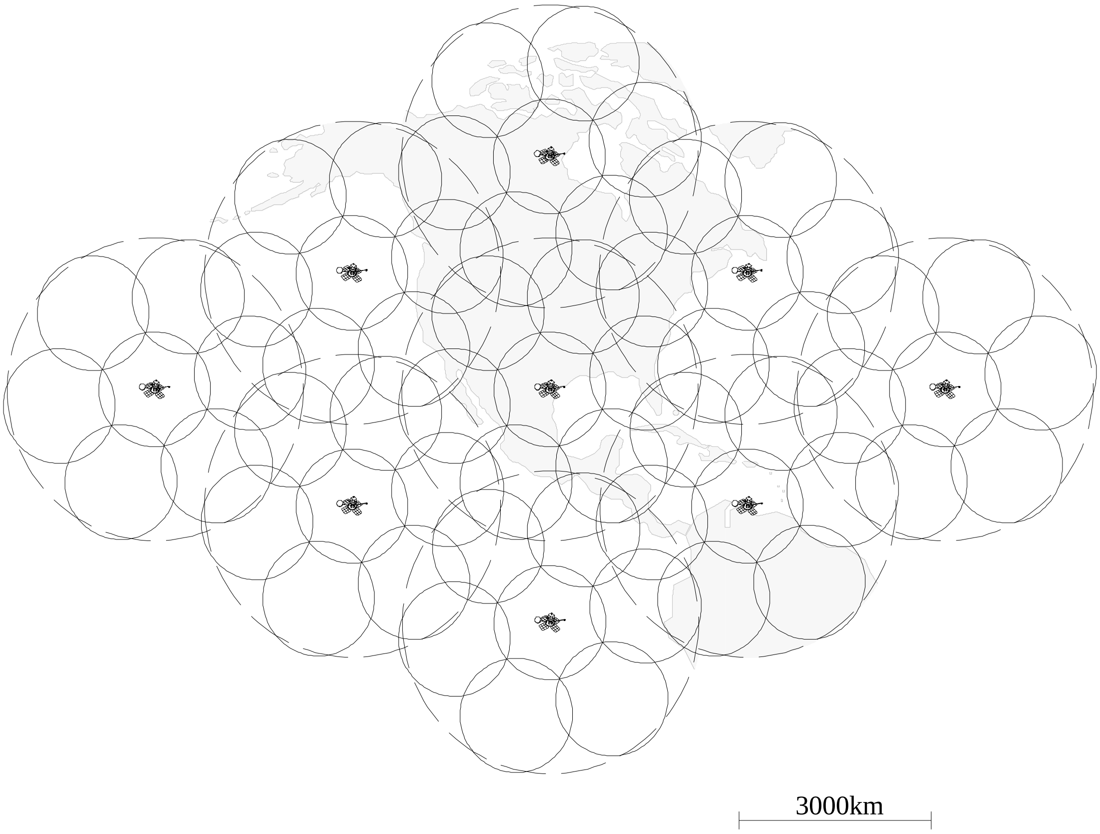

— — *edge of 4724km Ø satellite footprint (6.3° elevation contour)*

—— *edge of 1785km Ø spot beam (-3dB gain contour)*

sub-satellite point

Figure 40 Convergence of beam patterns at 30°North over the Americas

A frequency reuse plan is planned by permanently assigning channels to cells based on worst-case interference scenarios such that so long as the plan is followed, interference is guaranteed to be low enough not to affect communications. Some cells may be assigned more channels than others to accommodate traffic from cells where demand is expected to be greater than average. When channel assignment is required, the FES will select one of the free channels assigned to the geographic cell in which the mobile terminal is located and can begin to use it immediately.

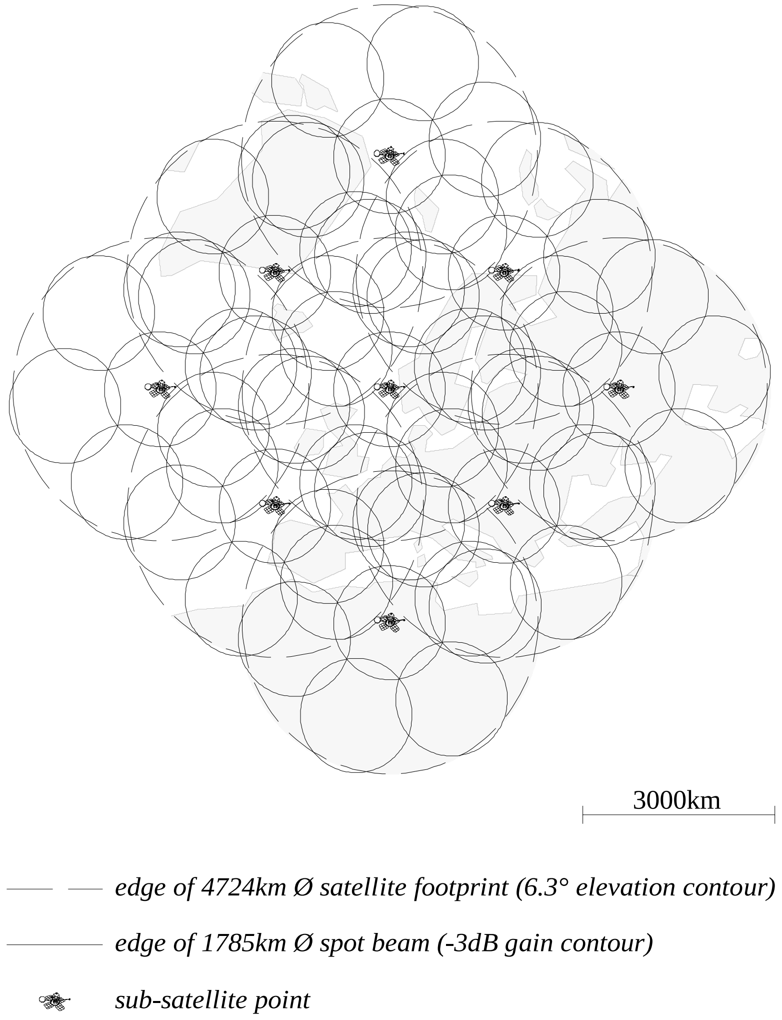

Figure 41 Convergence of beam patterns at 60°North over Europe

However, as with terrestrial cellular networks, spectral inefficiency results from

- an over-cautious minimum re-use distance being specified for the worst case condition (i.e. best-case propagation condition), which fixes an artificially low level of frequency re-use across the whole system, and

- the lack of capability to adapt to uneven traffic demand by allowing unused traffic channels assigned to one cell to be used in another cell where there is more offered traffic than there is Traffic Channel resource.

7.8. Dynamic Channel Assignment

DCA was proposed by Cox and Reudink in 1972 to improve efficiency [COX1, COX2, COX3, ENGEL]. DCA is currently used only in cordless pico-cell systems, such as CT-2 and DECT [ERIKSSON], where a vacant channel is selected at the start of a call and held for the duration of the call. It is not currently in use in terrestrial macro-cell systems but is being proposed for FPLMTS as a means of increasing spectral efficiency by circumventing the two constraints described above [BECK]. In principle, it can take advantage of locally poor propagation conditions (in built-up areas, for example) to allow frequencies to be reused more often by allowing the dynamic selection of traffic channels according to real C/I ratios measured at the time of channel assignment on the radio link that would be used for communication. In this way traffic channel assignments also follow customers' actual traffic demands rather than a static pre-determined plan and it can be shown that large spectral efficiency gains can be achieved [NETTLETON]. DCA is particularly effective where traffic demand is not uniformly distributed geographically, when DCA will concentrate resources on the most heavily used spot beams. Purely DCA can be conceived where all mobile terminals and FESs can select any traffic channels from a single pool of channels. Also hybrid schemes can be considered, where some channels are assigned to FESs in a fixed re-use plan and there is a pool of traffic channels for dynamic assignment to any FES where there is too much offered traffic for the fixed assignment of traffic channels to handle. Hybrid schemes can be useful to limit the large processing overhead of choosing a traffic channel from a very large pool of traffic channels.

7.8.1. Introduction to DCA

Without information from all other FESs, both an FES and its mobile terminal must monitor their receive channels to determine if a particular Traffic Channel is in use. In addition Uplink and Downlink Traffic Channels need to be paired such that, for any two-way link, if traffic can be detected on one channel then the return traffic is guaranteed to be on the paired channel.

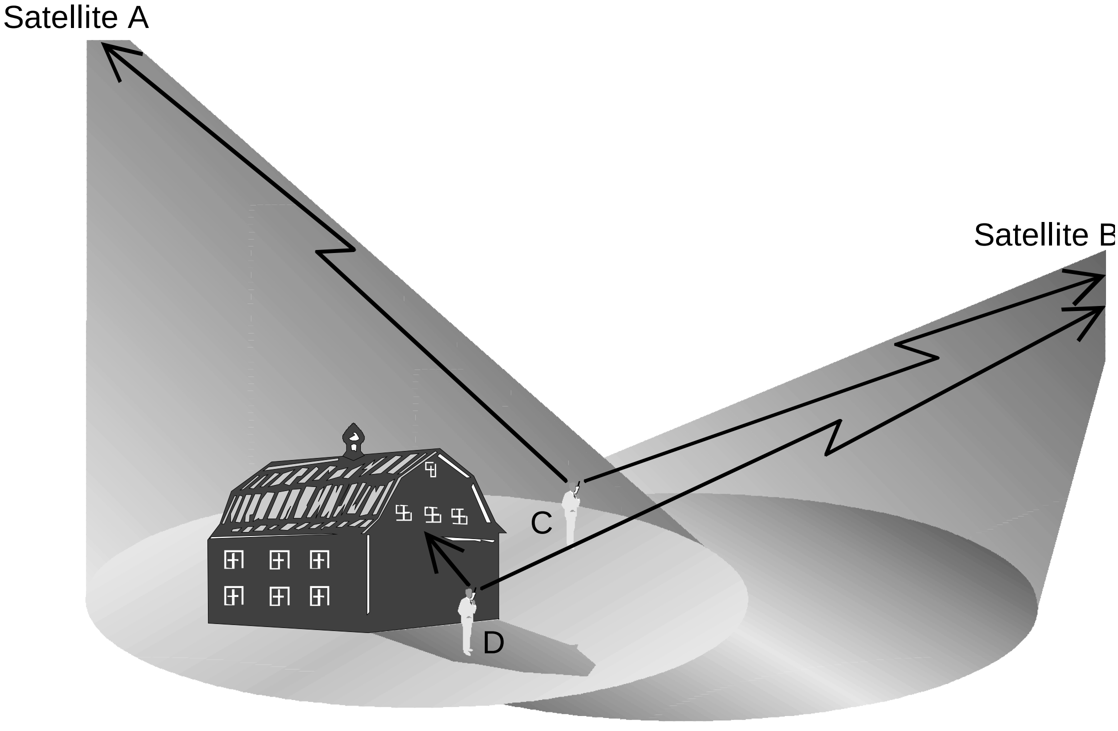

Figure 42 Uplink carriers from mobile terminals in overlapping satellite coverage areas, one in a shadow

Consider the uplinks example in figure 42. An FES is attempting to assign a Traffic Channel through satellite B for communications with mobile terminal D. It can deduce from its reception of C's uplink transmissions that the proposed uplink and downlink pair must already be in use between C and some other satellite (in this case A) and avoid co-channel interference by selecting another Traffic Channel pair.

If the Traffic Channel Pool contains a large number of channels, then the FES may have to check many channels before it finds a vacant one. This search would therefore be implemented as a background task that FESs perform continuously using a spare receiver. (If there is no spare receiver, then the FES is fully occupied and any call request would have to be blocked anyway.) Obviously the FES knows what Traffic Channels it is actively using itself, so it listens to the other Uplink Traffic Channels and compiles a shortlist of several candidate unused Traffic Channels ready for the next call request.

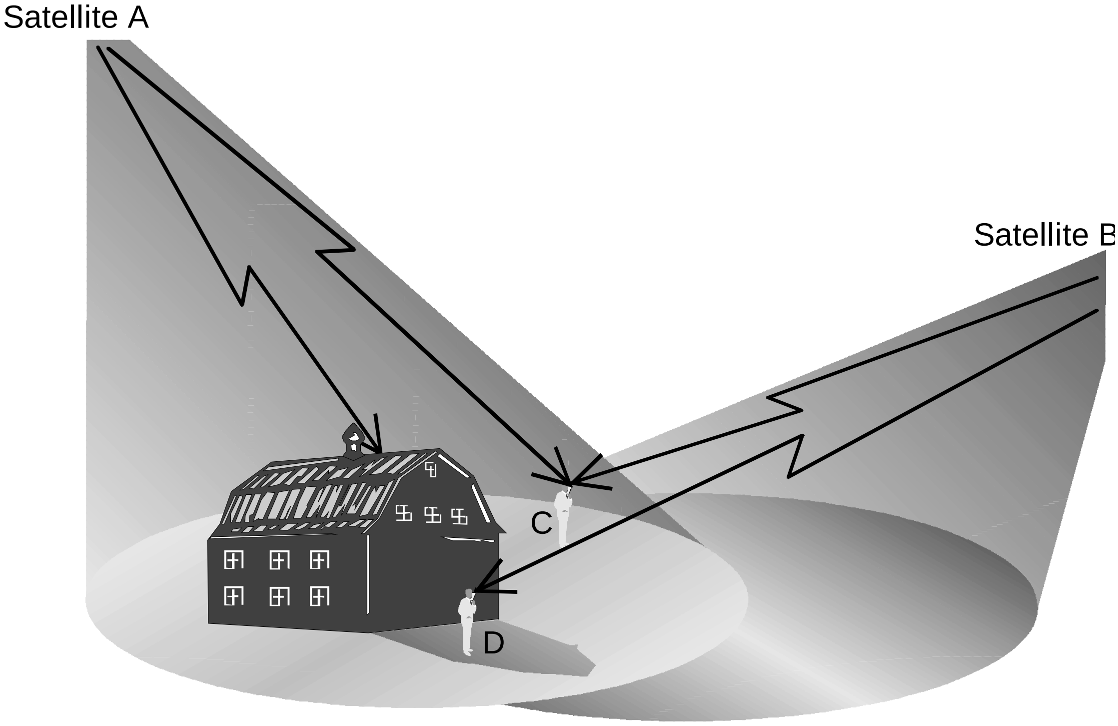

Figure 43 Downlink carriers to mobile terminals in overlapping satellite coverage areas, one in a shadow

Figure 43 illustrates the corresponding downlinks to figure 42. Consider what happens if satellite A receives a call request requiring communications with mobile terminal C. An FES using satellite A could not detect mobile terminal D's uplink because of an obstruction in the radio path (refer back to figure 42) and therefore does not know that the Traffic Channel is in use. Fortunately mobile terminal C is in a position to detect satellite B's downlink and can therefore determine for certain whether or not the Uplink and Downlink Traffic Channel pair are in use. From figure 42 note that mobile terminal D could not have detected that A was already using the proposed downlink because it is in a radio shadow - the FES must check that the Uplink Traffic Channel is unused and the mobile terminal must also check the Downlink Traffic Channel before any transmissions begin.

Hence on receipt of an *Origination Message* (or an incoming call *Page Response Message*) on an FES's Access Channel, the FES cannot immediately assign a Traffic Channel pair without first getting the mobile terminal to check that the proposed downlink is not in use. Therefore it sends a *Channel Assignment Message* to the mobile terminal on the Paging Channel and then listens to the Uplink Traffic Channel and waits for the *Traffic Channel Preamble*. When the mobile terminal receives the *Channel Assignment Message* on the Paging Channel it switches to listening to the assigned Downlink Traffic Channel and listens for a set period to ensure there is no traffic already using the channel. If the channel is vacant, it begins transmitting the *Traffic Channel Preamble* on the Uplink Traffic Channel, otherwise it returns to listening to the Paging Channel. If the FES has not received the *Traffic Channel Preamble* before a set time-out period it transmits another *Channel Assignment Message* on the Paging Channel to try a different candidate traffic channel pair. This process repeats as often as necessary until the mobile terminal is able to respond with the *Traffic Channel Preamble* on a vacant traffic channel. This sequence of events is shown in figure 44.

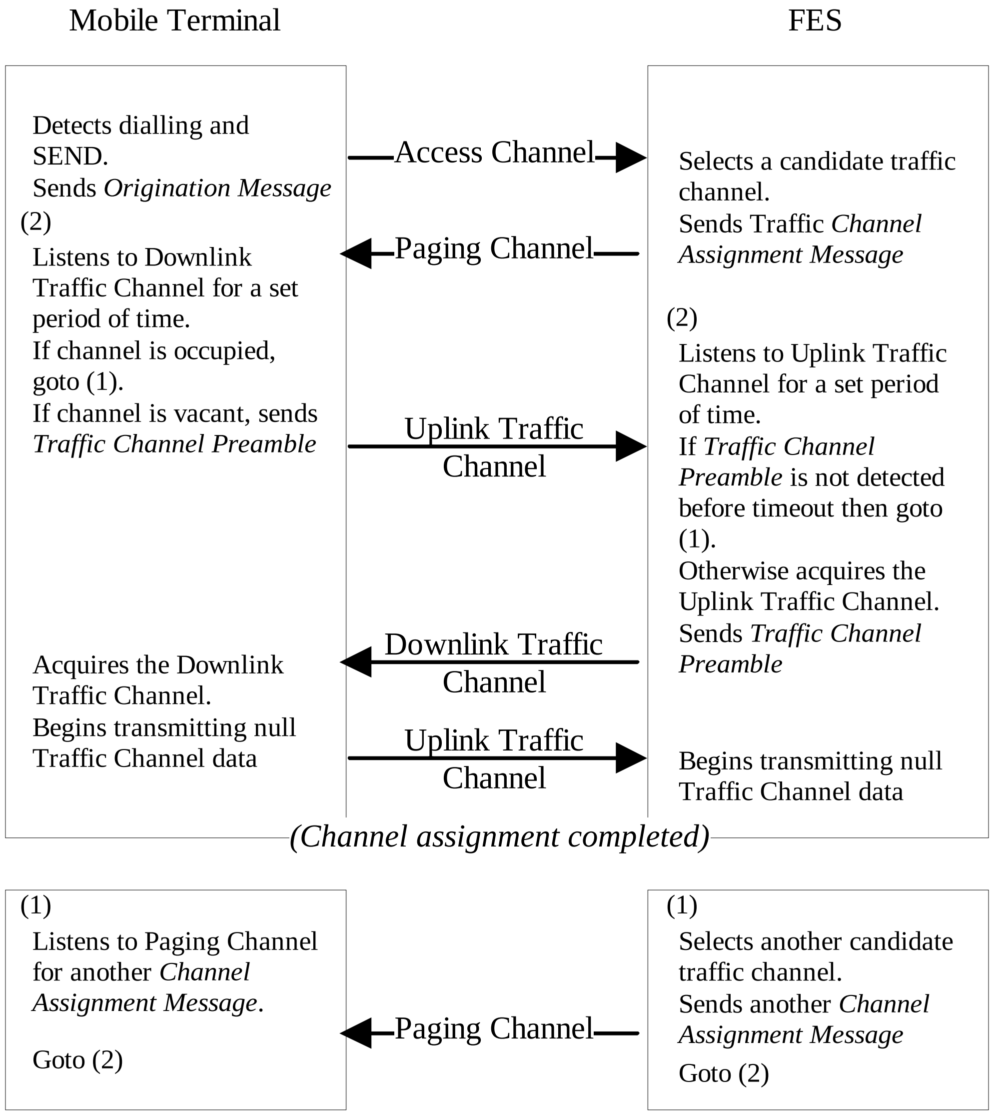

Figure 44   Call flow showing DCA for a mobile terminal originated call

For completeness, consider figure 45 depicting both mobile terminals in radio shadows. In this case it is impossible to determine if a proposed Traffic Channel pair is in use but it does not matter. The radio shadow ensures that no co-channel interference occurs and this actually increases frequency re-use. However, if either mobile terminal moves out of its shadow interference will result and a rapid handover will be required (as described in section 6.5.4). A slight increase in the number of handovers occurring may be one of the effects of using DCA.

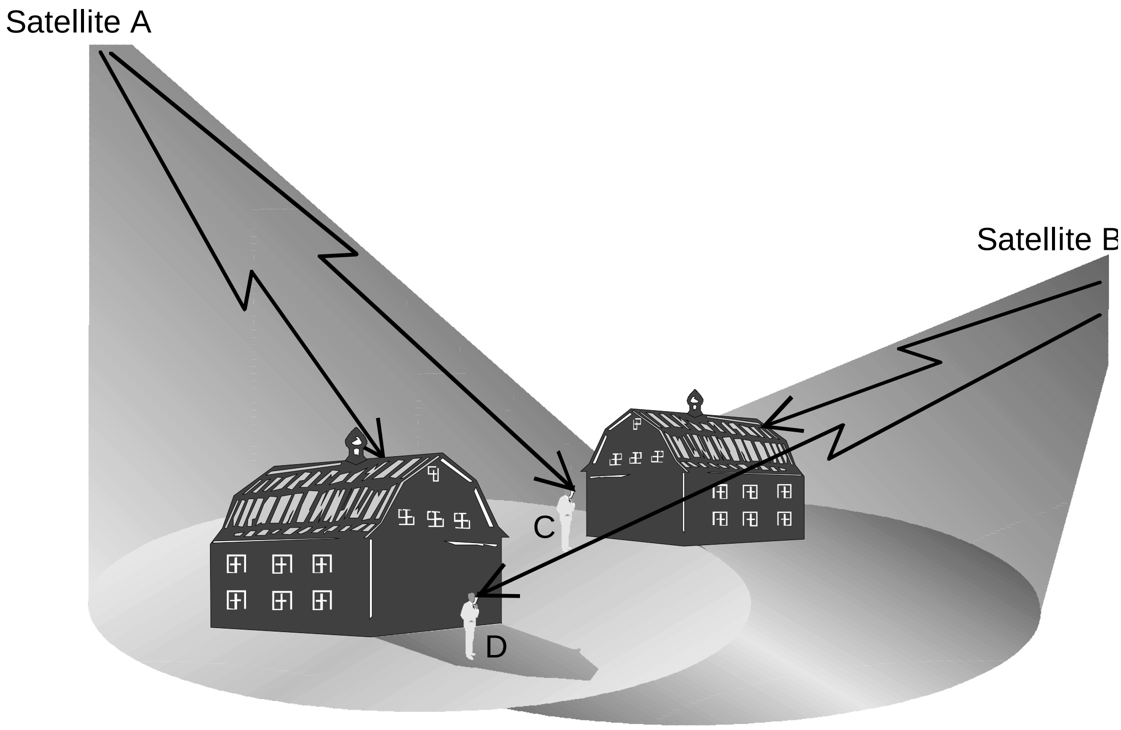

Figure 45 Downlink carriers to mobile terminals in overlapping satellite coverage areas, both in shadows

7.8.2. Simulation Model

To quantify the capacity gains attainable using DCA in place of fixed frequency re-use plans, a computer model of a 100 Traffic Channel DCA system operating on the 769km baseline LEO satellite network was constructed. Details of the model used are presented in appendix A. Over thirty million call requests were simulated on a model of the LEO system and DCA algorithms, written in K&R C. Thirty scenarios were simulated to cover a range of different offered traffic intensities at the centre of each of the spot beam patterns shown in figures 39 (Equator), 40 (30°N) and 41 (60°N), showing the effects of varying amounts of satellite footprint and spot beam overlap. For comparison, the performance of a geographically fixed cell frequency re-use plan was studied analytically for the same satellite system.

The full results of the DCA simulations are presented in appendices B and C and are discussed further later in this chapter. The results show the capacity of the simulated system to be approximately 4.3 channels per 1,000,000km^2 for a 100 channel pool, producing satisfactory blocking and dropping probabilities.

The capacity of this model was compared with the capacity of a fixed frequency re-use plan that was constructed analytically for the same network model (see appendix A). To ensure the same channel characteristics as the DCA algorithm, a 27 cell frequency re-use pattern is required for the model, offering 1.79 channels per 1,000,000km^2 for the 100 channel pool.

These results demonstrate the opportunity to potentially increase the satellite system capacity by a factor of 2.4 over the fixed re-use plan capacity by employing the simulated DCA algorithm. This increase is a similar order of benefit to those quoted for the terrestrial cellular DCA simulation work, for example the increase by a factor of 3 that has been reported in [NETTLETON].

In summary, DCA in the satellite environment is possible and produces efficient use of spectrum. A means of measuring received signal power and downlink C/I ratio is required at the mobile terminal.

7.8.3. Potential Complications

For the mobile terminal, measuring the received signal power of Paging Channels and the downlink C/I ratio in potential Traffic Channels are similar operations. Each involves the rapid re-tuning of the receiver frequency synthesizer to the frequency spectrum window in which the relevant channel should be and the integration of received energy in the window's bandwidth over a short period of time. For Paging Channels the end-value of this integration represents the power of the Paging Channel carrier, which would be recorded for comparison with other Paging Channels. For proposed Traffic Channels, the end-value of this integration represents interference power, which must be below a threshold defined by C/I_{block} and C (the expected received carrier power) for the mobile terminal to be able to start transmitting. It should be noted that frequency synthesizers capable of rapid re-tuning are expensive but their use is essential for DCA. However, a cellular system's location registration mechanism may require their use for Paging Channel scanning even if DCA is not employed.

Section 7.4 shows that for FDMA and TDMA multiple access schemes, wherever there is a possibility of different satellite footprints overlapping then very large frequency guard bands are required (up to 88kHz at 2GHz) between Uplink Traffic Channels used by different satellites. No guard bands need be left between channels assigned for communications with the same satellite. It is therefore preferable to always assign Traffic Channels using the same satellite adjacent to each other in the frequency spectrum. In a DCA scheme FESs can give priority to unused Traffic Channels that are adjacent to Traffic Channels active on that satellite by proposing these to the mobile terminal first. Since all FESs would be doing this, channel assignments will tend to group themselves into blocks of channels, one for each satellite, within which there is no need for frequency guard bands to be left. The simulations used a very rudimentary algorithm to group assignments and because the scheme is highly dynamic, these blocks fragmented and shifted around the frequency band. A better algorithm could be conceived to sort the candidate unused Traffic Channels into an order that fills in small gaps between existing assignments first and the resultant spectral efficiency should be quite acceptable.

A more significant consideration will be the relative Doppler shifting of Downlink Traffic Channel frequencies with respect to those used by different satellites as satellites move about the sky. This will either

- complicate the mobile terminal's monitoring of proposed downlink channels at the edges of a satellite's block of carriers so that the terminal widens its measurement bandwidth to include the guard band or

- cause some premature handovers to allow mobile terminals to escape from interference if another carrier drifts in from the guard band.

7.8.4. Limiting the Number of Traffic Channel Proposals

One major problem with DCA is that when the system is being heavily used, there is a high probability that a FES's candidate Traffic Channel Pair will be found to be in use when the mobile terminal tests the proposed Downlink Traffic Channel. This implies that there may be many recursions of the "propose a new Traffic Channel and test the proposed Downlink Traffic Channel" loop before a vacant channel is found, taking a significant increasing handover and call set-up time. How much time is consumed depends on the set period for which the mobile terminal must monitor the proposed Downlink Traffic Channel (which depends on how voice activation is being used and on the average shadow duration[1]), the traffic intensity being offered to the satellite system and the amount of radio path obstruction that is occurring.

The average number of channel assignment retries can be quantified to see what the mean speed of channel assignment will be. Channel assignment is the longest delay in the handover process unless the mobile terminal is equipped with two receivers such that conversation and the channel search can continue simultaneously. As such it directly affects how intrusive handovers are to the conversations of customers using low-cost single-channel versions of satellite compatible terminals. Without making bold assumptions about the frame structure of the Downlink Traffic Channel it is not possible to predict exactly what the set period for mobile terminal monitoring of the Downlink Traffic Channel should be but a period of around 250ms is ample for RF re-tuning and monitoring of the channel (as long as a voice de-activated carrier is still detectable) and can be used as a rough estimate.

The probability of any given channel being in use increases as traffic intensity increases. The FES improves the probability that its proposed Traffic Channels will be unused by previously checking the Uplink Traffic Channels. Therefore the probability that a proposed Traffic Channel is in use will also depend on the radio propagation conditions in the beams covering the mobile terminal - if there is much radio shadowing then the probability will be higher than if there is none. The only thing certain is that "Pr(in use)", the probability that a proposed Traffic Channel is already being used, increases as traffic intensity increases and increases proportionally to some measure of average radio path obstruction. The probability is the same for all proposed Traffic Channels so the probability of having to consider *n* different proposed Traffic Channels is

---

[1] The average duration of radio shadows is longer for higher altitude constellations and shorter for LEO satellites which move faster. When shadow durations are expected to be long (for example MEOs) then carrier monitoring can be for a short period and the channel assignment should be made to take advantage of the shadow for as long as it lasts (see section 7.8.1). Problems only occur in a very rapidly changing shadow where shadow durations are about the same as the time taken to complete the channel assignment. In this case, handover may be forced very soon after the channel assignment when interference begins at the edge of the shadow. This is more likely for LEO but not anticipated to be a major waste of resources.

$$\Pr(n) = \Pr(n\text{-1 unsuitable considerations}) \times \Pr(\text{1 suitable consideration})$$

$$= (\Pr(\text{in use}))^{n-1} \times (1 - \Pr(\text{in use})).$$

The mean number of Traffic Channel proposals before a vacant Traffic Channel is found is

$$\text{Mean number of proposals} = \sum_{n=1}^{N} n.\Pr(n) = (1 - \Pr(\text{in use})) \times \sum_{n=1}^{N} n.(\Pr(\text{in use}))^{n-1}$$

which, for N = a sufficiently large number of Traffic Channels in the Traffic Channel Pool, is an exponential rise in the mean time taken as traffic intensity increases. Evidently there will be a problem as Pr(in use) approaches 1, which happens when the system is approaching full occupancy. The number of retries on different Traffic Channels becomes exponentially larger as the number of vacant channels becomes smaller, resulting in pauses of many seconds before a Traffic Channel is found. In such circumstances the maximum number of proposals that can be tested on a satellite must be limited to that which can be done in 5 seconds, for example. Two options are considered: a lowest possible cost mobile terminal using a single receiver switchable between the communications modem and the interference power meter and a more expensive mobile terminal using a dual-channel receiver, one dedicated to communications and the other spare for interference power measurements[2]. If the channel assignment is for a handover then for the former, this represents a 5s break in communications. For the latter, there is no break in communications but the handover is delayed by 5s. Even if conversation can continue during the handover process, the Traffic Channel quality will be rapidly deteriorating as the satellite beam moves away from the mobile terminal or as the radio path obstruction changes and so long delays cannot be tolerated. Therefore a limit of 20 retries (5s delay) is set and

$$\text{Mean number of proposals} = \left((1 - \Pr(\text{in use})) \times \sum_{n=1}^{20} n.(\Pr(\text{in use}))^{n-1} \right)$$

$$+ 20 \times \left(1 - \left((1 - \Pr(\text{in use})) \times \sum_{n=1}^{20} \Pr(\text{in use})^{n-1} \right) \right)$$

which is plotted in figure 46, along with the probability of this limit being reached without having found a vacant Traffic Channel.

Figure 46 clearly shows the breakdown of independent channel assignment as the system becomes congested and how the last few percent of channels are difficult to use. Without experimental results for the obstruction of non-GEO satellite radio paths it is not possible to infer what traffic intensity would result in Pr(in use) = 0.8, which is perhaps the maximum acceptable delay (average of 5 proposals, or 1.25s delay) and failed assignment probability (1.2% of calls either dropped or moved to a new satellite on handover).

---

[2] A dual-channel receiver requires duplication of the baseband multiplier and frequency tuning but a second demodulator is not required - DCA requires measurement of total power in a given bandwidth, not of a specific carrier. Both channels would share the same RF front-end.

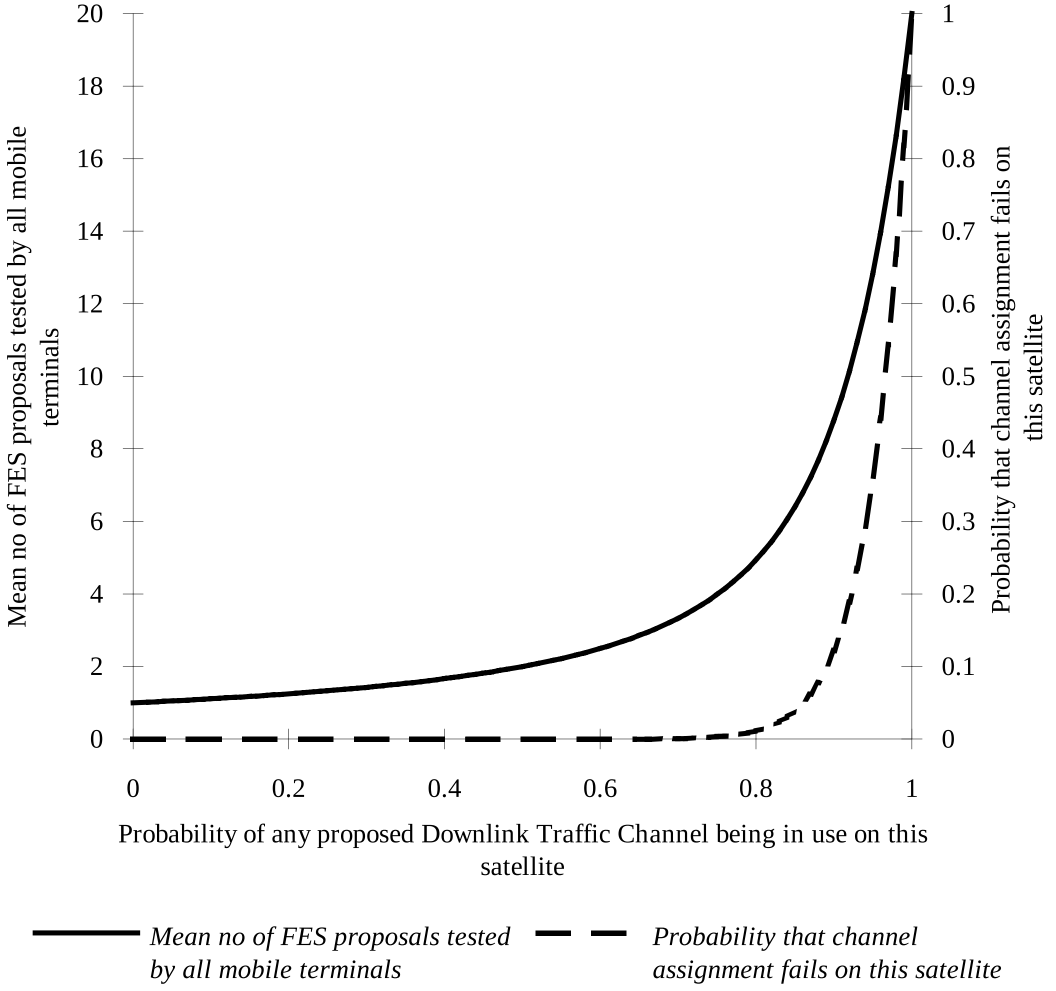

Figure 46 Mean number of proposals for mobile terminal to test and probability of Traffic Channel assignment failing vs. Pr(in use), maximum 20 retries

Recognizing that not many channel assignments go beyond the 5th channel proposal and that the delay if they did would inconvenience customers, figure 47 shows what happens to figure 46 if the retry limit is moved down from 20 channels to 5 channels, representing an absolute maximum delay of 1.25s before the call is either blocked, dropped or an inter-satellite handover is attempted. It appears that the probability of channel assignment failure is significantly raised in congested systems.

As indicated before, the relationship between offered traffic intensity and Pr(in use) is not clear. To clarify the relationship the computer model mentioned in section 7.8.2 was constructed with both 5 retry and 20 retry schemes and simulations were performed for a range of offered traffic intensities using both schemes. The complete results for a ceiling of 5 retries are presented and discussed in appendix B. The results for the 20 retry ceiling are presented and compared with the 5 retry results in appendix C. The results show that there is very little increase in call blocking and dropping probabilities when the number of channels that can be tested by the mobile terminal is reduced from 20 to

5, as can be seen in figure 51. In fact, the reduction is beneficial by slightly increasing the capacity of the system, by as much as 15% in the simulations of equatorial coverage (see figure C4 of appendix C). To understand this, it is instructive to re-draw the right-hand side of figure 46 with the mean number of proposals made to terminals that were successful in finding a spare Traffic Channel on the satellite, excluding proposals made to terminals that did not find a channel on the satellite. This plot is shown in figure 48 and is comparable with the simulations' records of numbers of proposals made leading up to successful handovers to new Traffic Channels on the same satellite, shown in figure 49. These graphs have similar shapes when the offered traffic intensity is plotted on a logarithmic scale.

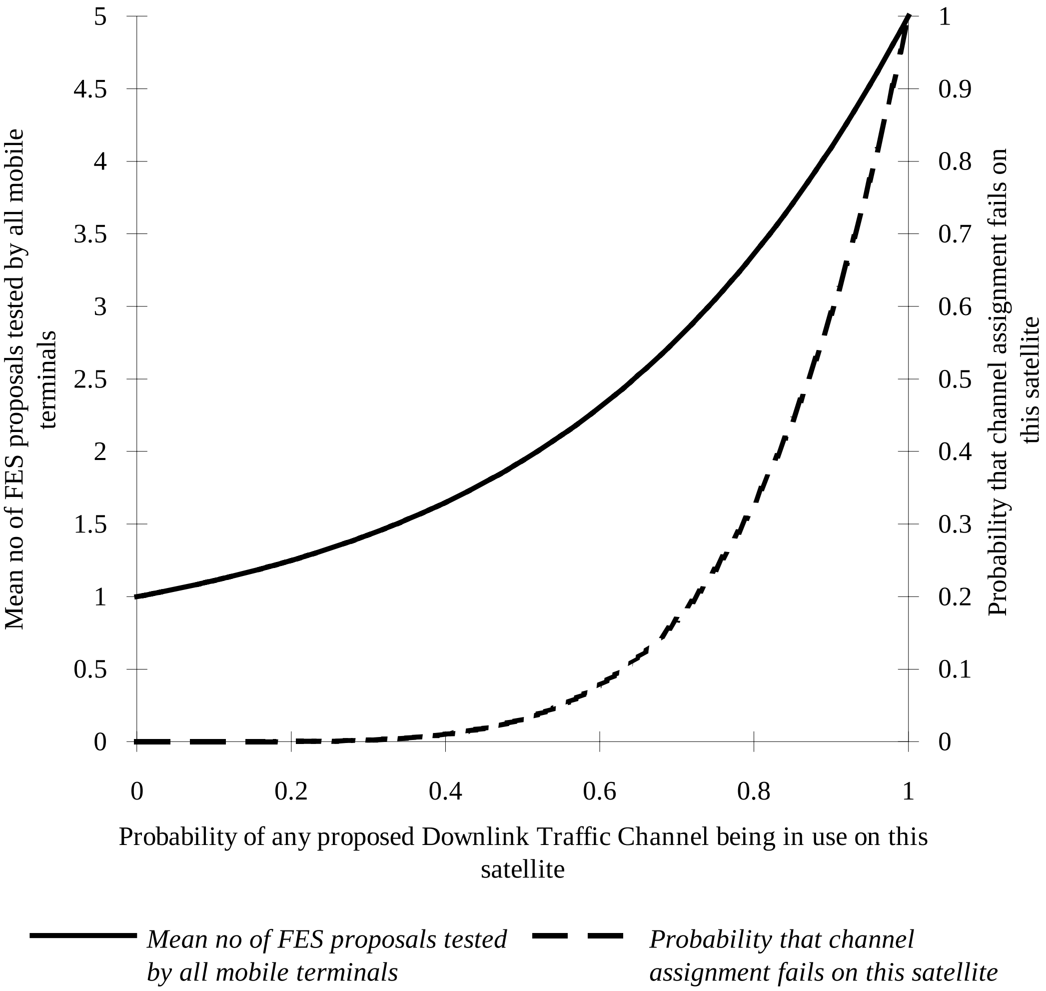

Figure 47 Mean number of proposals for mobile terminal to test and probability of Traffic Channel assignment failing vs. Pr(in use), maximum 5 retries

Figure 49 shows that the peak of the mean number of proposals varies depending on the interference from other satellites - the simulations at 60°North latitude have more satellite beam overlap than the simulations at 30°North and 0°North and all the curves have a different peak to the theoretical curve in figure 48. The curves show that on average, if a new channel has not been found within 5 proposals, it will not be found, which is why reducing the limit to 5 retries did not significantly raise the call dropping and blocking probabilities (see figure 51).

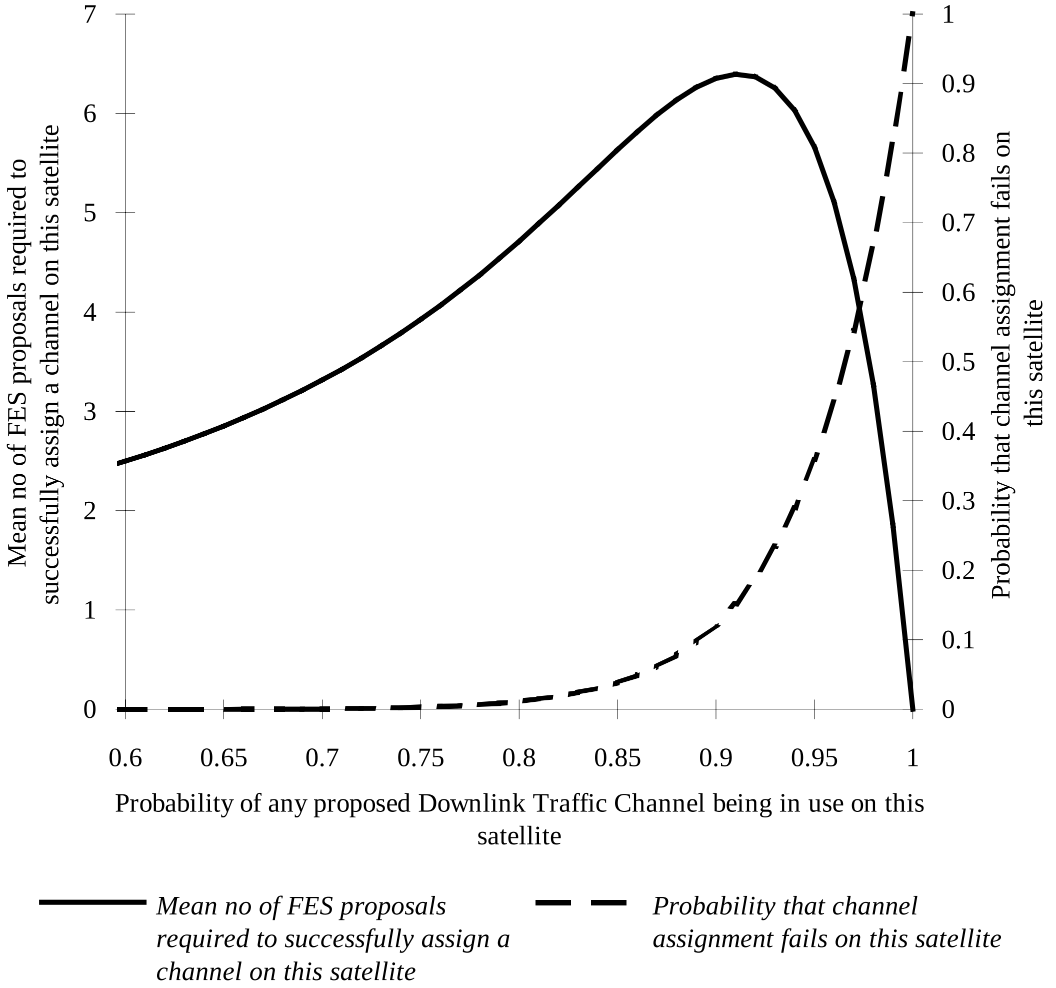

Figure 48 Mean number of proposals for mobile terminal to test, excluding failed assignment statistics, and probability of Traffic Channel assignment failing vs. Pr(in use), maximum 20 retries

As a result of these simulations, it is concluded that it is best to limit the number of Traffic Channel proposals to a small number and that increasing the number of retries allowed beyond this value is counterproductive. In the simulated system, the ideal value for this maximum appeared to be close to 5 retries. Initial channel assignment is less time-critical than handover assignments and more retries could be tolerated here but if the probability of a handover channel assignment failing is higher than that for initial assignment then the probability of calls being dropped at their first handover would rise. It is felt that a high blocking probability is better than a high call dropping probability and so exactly the same process should be used in both cases.

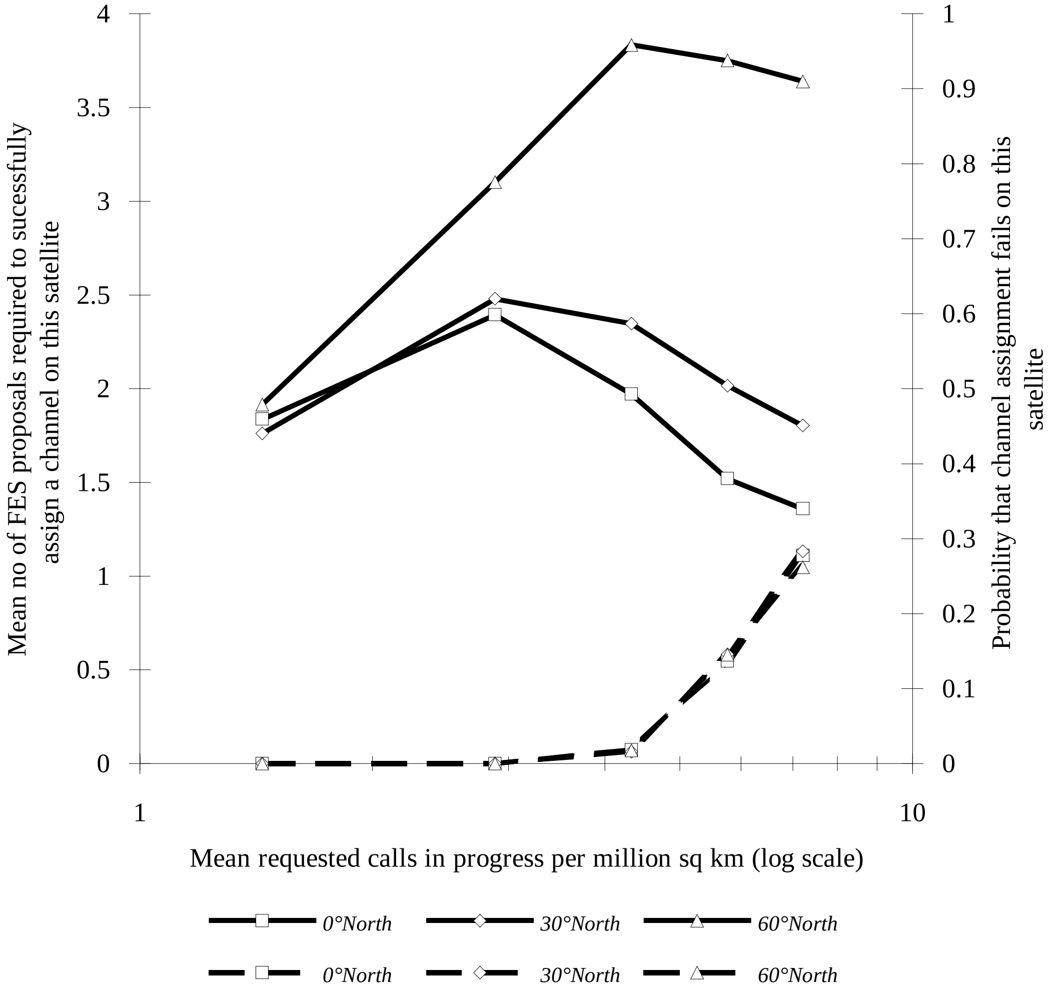

Figure 49 Simulation statistics for mean number of proposals for mobile terminal to test, excluding failed assignment statistics, (solid lines) and probability of assignment blocking (dotted lines) vs. traffic intensity on a logarithmic scale, maximum 20 retries

For dual channel mobile terminals the limit of 5 retries is still recommended because a call is less likely to be dropped during a rapid handover and the capacity of the network would be reduced if the limit were raised. The reason for this reduction in capacity when the retry limit is raised is thought to be that FESs' proposals beyond the optimum number of proposals are likely to be marginal cases where interference is only just acceptable. The addition of another mobile terminal to the interference environment causes more severe interference to other terminals than terminals using better Traffic Channels would. This interference then causes more rejections of other terminal's channel assignment requests, resulting in less traffic being carried by the network. In effect, poor quality calls are being allowed to continue, causing potentially lower-interference calls to be blocked. The conclusion drawn from the simulation is therefore that bad interferers are better dropped than rescued by more thorough channel searches. This is consistent with accepted wisdom, which shows an actual drop in throughput in severely congested systems. The classic graph is shown in figure 50.

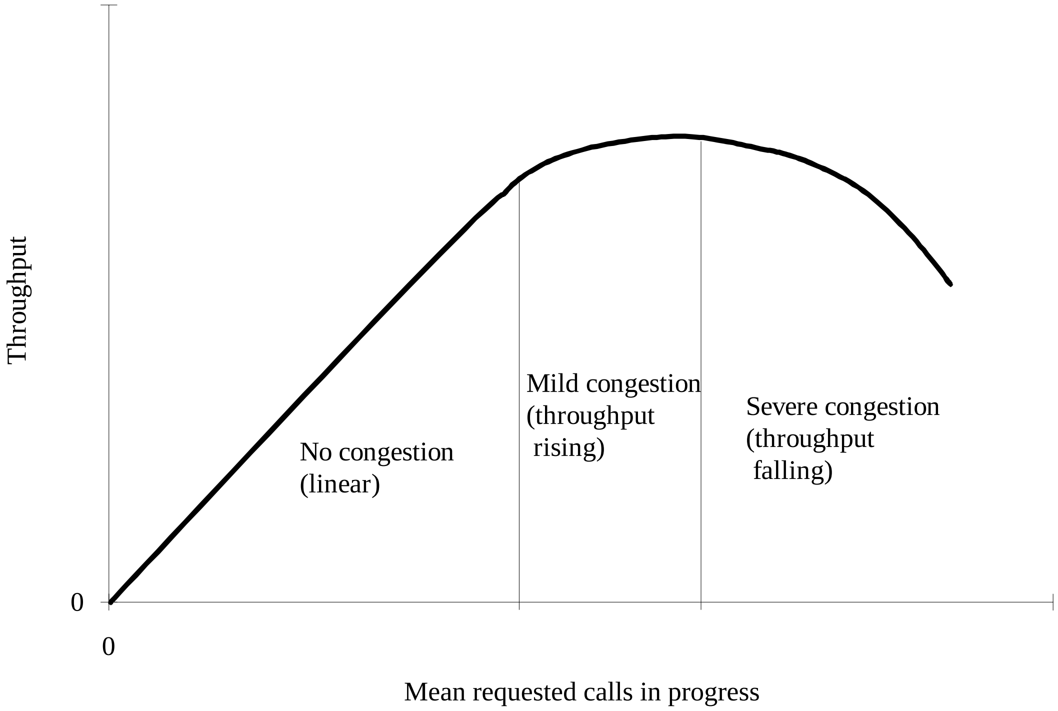

Figure 50   Expected fall in throughput in an overloaded system

In summary, FESs make Traffic Channel proposals for mobile terminals to measure before communications can commence and the delay introduced by this iterative process needs to be minimized. Therefore a maximum limit is placed on the number of proposals that an FES will make before it refuses the channel assignment on a satellite. There appears to be an optimum value for this maximum, below which too many channel assignments would be refused and above which the total capacity of the network begins to fall. For the simulated system, the optimum was close to 5, representing a 1.25s delay before handover. For other systems the optimum is likely to be different.

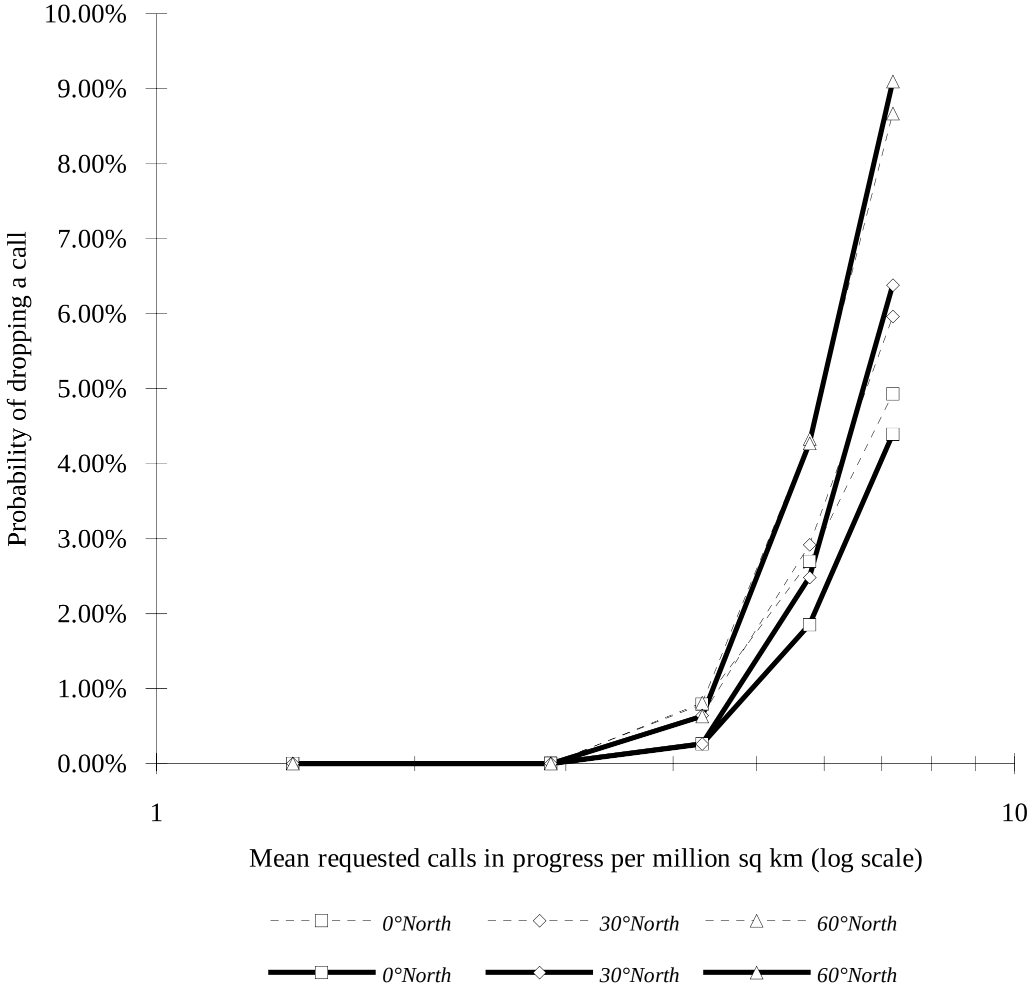

Figure 51 Probability of dropping a call vs. traffic intensity on a logarithmic scale for maximum 20 retry simulation (thick lines) and maximum 5 retry simulation (feint lines)

7.9. DCA Aided by Bulletin Boards

To speed the Traffic Channel selection process up a "bulletin board" system can be considered where each FES transmits the identities of all the channels that it is using on a control channel, such that a mobile terminal is able to check the usage of a Traffic Channel without having to tune into the Downlink Traffic Channel and listen to it directly. Note that the mobile terminal still cannot select the Traffic Channel by itself - referring back to the satellite downlinks, figure 43, mobile terminal D cannot even receive satellite A's Bulletin Board Channel, which may include traffic from a different FES to satellite B. If it was to try to use the same Traffic Channel as mobile terminal C is using to communicate with satellite A then satellite B's downlink to D would interfere with C's reception of A's downlink. However, if the FES was to suggest candidate channels for the mobile terminal to check then a FES using satellite B would never suggest the channel which C is using because it would have already checked the Uplink Traffic Channel and found that it was in use. The bulletin board scheme is therefore very similar to that previously described by figure 44 except that after sending the *Channel Assignment Message* the FES listens for a message on the Access Channel from the mobile terminal indicating whether the proposed Traffic Channel should be used or not and the mobile terminal uses the bulletin board to verify the proposed channel's vacancy instead of listening directly to its downlink. Once a suitable channel is found and the mobile terminal has sent a message on the aloha Access Channel confirming that it will use the proposed Traffic Channel, it begins to transmit the *Traffic Channel Preamble* on the Uplink Traffic Channel and the Uplink and Downlink Traffic Channels are acquired as before.

It is questionable which process would be more rapid: direct monitoring of the proposed Downlink Traffic Channel by the mobile terminal or checking that the proposed Traffic Channel does not appear on any of the Bulletin Board Channels that the mobile terminal can receive (there may be several - one for each FES in each spot beam of each satellite whose coverage area the mobile terminal is within). The answer to this question would depend most on how much additional complexity can be added to the mobile terminal equipment. If the mobile terminal can compile and continuously update a complete list of all the Downlink Traffic Channels that it cannot use, ready for checking against when a FES proposes a Traffic Channel, then the bulletin board system would be the most rapid. This would require sufficient memory to hold the list and the ability to keep the list current even whilst communications are taking place on a Traffic Channel (so that the list is ready for handovers) - i.e. a dual-channel receiver. If a single channel receiver is being used then it can only compile the list after the *Channel Assignment Message* is received then re-tuning and listening to each complete bulletin board for each satellite beam would usually take longer than directly monitoring several different proposed Downlink Traffic Channels[3].

Apart from adding to mobile terminal complexity the bulletin board system suffers from not being a direct measure of interference on the proposed radio link. The Bulletin Board Channel would be sufficiently close in frequency to ensure that the radio propagation is similar but there may be non-FPLMTS interference in the Traffic Channel bands that would be detected and avoided if interference was measured directly. Another drawback is that extra spectrum is required for the Bulletin Board Channel. The bulletin board should not be carried on the Paging Channel because it is best kept concise such that it can be read very quickly. Also, different Bulletin Board Channels would need to be transmitted in each satellite beam so that the information is local to the coverage of the mobile terminal. The Bulletin Board Channel would also require a higher link margin to ensure that its contents could be decoded and read in areas outside normal Traffic Channel coverage but where frequency re-use of channels would still cause unacceptable co-channel interference. Just because a Traffic Channel

---

[3] The best case is the same for both systems. If there is no overlap of satellite coverage areas then there is only one bulletin board to read but with no other satellites to interfere, the proposed Traffic Channel is certain to be unused. Hence both systems can confirm that the proposed Traffic Channel is unused very quickly. The worst case is also the same for both systems; it is where many satellite beams overlap. In this case a significant delay would occur during the sequential reading of multiple Bulletin Board Channels before the mobile terminal could determine if the proposed Traffic Channel is in use. However, with many overlapping beams the probability of the proposed Traffic Channel already being in use is high, so a number of proposals required from the FES before a free Traffic Channel is found may be high, as previously shown.

signal is not strong enough to be able to receive and decode it does not mean that it is too weak to cause unacceptable interference with the mobile terminal's proposed communications on the same Traffic Channel. DCA's direct measurement of interference power would detect this interference because it does not need to decode the Traffic Channel.

So, for a bulletin board to be of value in speeding up Traffic Channel selection it must consume significant satellite power and spectral resources and extra complexity is required in the mobile terminal. As the disadvantages of bulletin boards appear to outweigh the advantage, their use is not recommended.

7.10. Hybrid Schemes

As mentioned in section 5.5, once FESs' GCAs extend beyond the coverage of one satellite, FESs can begin to plan frequency reuse within their coverage areas. If information is shared between FESs then planning can take account of potential interference between FESs where coverage areas overlap. The objective of this extra planning is to speed up handovers by increasing the probability that a mobile terminal is able to accept an FES's proposed Traffic Channel. For example, if the two Traffic Channels of the two satellites in figure 42 are operated by the same FES then the FES controller software could calculate that B and D's Traffic Channel could be interfered with if A and C were to use it and not propose that Traffic Channel to C. If the Traffic Channels were operated by different FESs then information transfer between the FESs would be required to predict this interference scenario. Note, however, that whenever interference judgements are based on modelling rather than measurements, the full capacity of a pure DCA system cannot be used. For example, the shadowing in figure 45 would not be predicted and exploited, whereas DCA would improve network capacity by using it.

The possibilities for such hybrid schemes are numerous, including:

- Using a fixed frequency reuse plan for most of the time to guarantee that the first Traffic Channel proposal is acceptable, only resorting to DCA techniques when the system is congested.

- Using dynamically planned frequency reuse, where FESs share sufficient information to practically guarantee that the first Traffic Channel proposal is acceptable by calculating theoretical interference conditions for the proposed channels in real time. This does not necessarily imply that all FESs have complete knowledge of the channel usage and beam coverages of all other FESs. Interference will only occur from those FESs that are near to the given FES, so only knowledge of data for those FESs is required. Exact beam coverages are not essential because what is calculated is effectively a fixed channel re-use pattern that changes with time. Like terrestrial fixed re-use patterns, this is based on constant re-use distances corresponding to worst-case conditions, making precise simulation of the radio environment unnecessary. What is required are the channel assignments in each beam within interference range and the approximate locations and coverages of each of these beams. From this information, look-up tables of beam overlaps and beams closer together than the minimum channel re-use distance can be calculated. Again, DCA would need to be used when the system is congested to realize the full spectral efficiency gains enabled by locally measuring interference and exploit the increased frequency re-use derived from poor propagation conditions.

- Planning capacities in a geographically fixed re-use pattern to match the non-uniform traffic demands across the FES's GCA.

- Using the satellite's predictable motion to plan handovers in advance. The satellite beam pattern in the 769km altitude simulated LEO is moving at 27,000km/h. It is safe to assume that no mobile terminal will be moving faster than this and therefore that handovers will always be to one of the beams following the current beam in the pattern. In fact, the motion of the mobile terminal is always insignificant compared with the motion of the satellite for any of the non-GEO satellite orbits that are being considered as candidates for FPLMTS. Handovers no longer occur depending on the statistically random motion of mobile terminals but depend on the motion of the satellite patterns, which is known. Because the motion is regular, the underlying rise and fall in C/I (ignoring fading and shadowing) can be predicted and handovers could be planned before their initiation using this information as described in chapters 5 and 6. Although satellite motion is complex at the seams in a constellation where satellites orbit the Earth in opposite directions, the motion is still deterministic and useful for planning handover.

In all hybrid cases, DCA can be used, and its capacity gains can be realized, as long as mobile terminals have the ability to measure the interference on a Traffic Channel before using it. Because channel assignment is under the control of FESs, advances in channel assignment techniques could be implemented by network operators without changes to the installed base of customer's mobile terminals, allowing rapid deployment of new techniques in FES software. Such complicated hybrid schemes were not simulated. Specifically, it was not verified whether DCA could still realize the same capacity gains in a Traffic Channel Pool if all assignments were not made using DCA but some were made using planned assignments.

In conclusion, the FES can take control of nearly all handovers and execute them much more quickly than pure DCA can. The optimum channel assignment algorithm would be a hybrid scheme, combining DCA and FES planning within its GCA, with some of the features listed below:

- Initial channel assignment using DCA and subsequent handovers to be planned by an FES to maintain the same frequency reuse as far as possible
- Changes in coverage patterns are sufficiently deterministic that in many cases channel assignments can be determined before their being requested by mobile terminals, making handover very quick.
- Under unbalanced load, an FES can dynamically re-assign channels between satellites to match demand and re-configure carrier block assignments to maintain good spectral efficiency.
- Under heavy traffic conditions, DCA could take over from the faster handover algorithms to use the increase in capacity that it promises, which the planned channel re-use algorithms cannot achieve because they assume worst-case modelled propagation.

- FES control of handover implies simple upgrading of channel assignment algorithms as improved methods are found and even allows manual intervention to accommodate changing priorities in emergencies. DCA techniques themselves have a lot more scope for development with, for example, techniques using neural networks [SHIMADA] now being tried out. [JORDAN] presents an attempt to place an upper theoretical bound on the capacity gains that can be achieved.

7.11. DCA as a Means of Enabling Band Sharing Between Networks

Because DCA allows some reduced pre-assignment of frequency assignments, it enables FESs to use the same frequency bands without necessarily sharing information on frequency occupancy. This would enable multiple satellite network operators to operate and compete in the same frequency band whilst remaining completely independent of each other, providing that the expected received power levels and the handover thresholds were comparable. This last proviso stems from the requirement that the mobile terminal must be able to judge if its transmissions will interfere with another's, so assumptions about the receive power and the receive bandwidth of the other terminal are implicit. For example, two similar TDMA networks might be able to coexist in the same frequency band but equitable sharing between CDMA and TDMA systems would be difficult, unless equitable meant that CDMA use was low enough not to interfere with or be interfered with by the TDMA system. The sharing of frequencies by different network operators would require detailed agreement between the operators to prevent unfair channel assignments and would be the subject of regulatory rule-making.

7.11.1. Satellite DCA's Compatibility with Terrestrial DCA

DCA as proposed for terrestrial systems is very similar to the described satellite scheme, from the point of view of mobile terminal involvement. The satellite scheme successfully isolates the intelligence of the scheme at the FESs where additional complexity is less costly than at the mobile terminals. All that is required of the mobile is the ability to rapidly determine the interference level on a candidate Traffic Channel, a task that can readily be performed by measurement of the received power in the bandwidth of the proposed Traffic Channel over a short period of time. This method of sensing Traffic Channel occupancy is the same as that proposed in many terrestrial dynamic channel assignment schemes (e.g. CT-2 and DECT), so its implementation in mobile terminals could be common for terrestrial and satellite use.

The possibility of one channel assignment algorithm for all FPLMTS beckons consideration that the satellite frequency bands of FPLMTS might share the same pool of Traffic Channels with terrestrial FPLMTS networks. Figure 52 illustrates the application of DCA to terrestrial and satellite frequency band sharing.

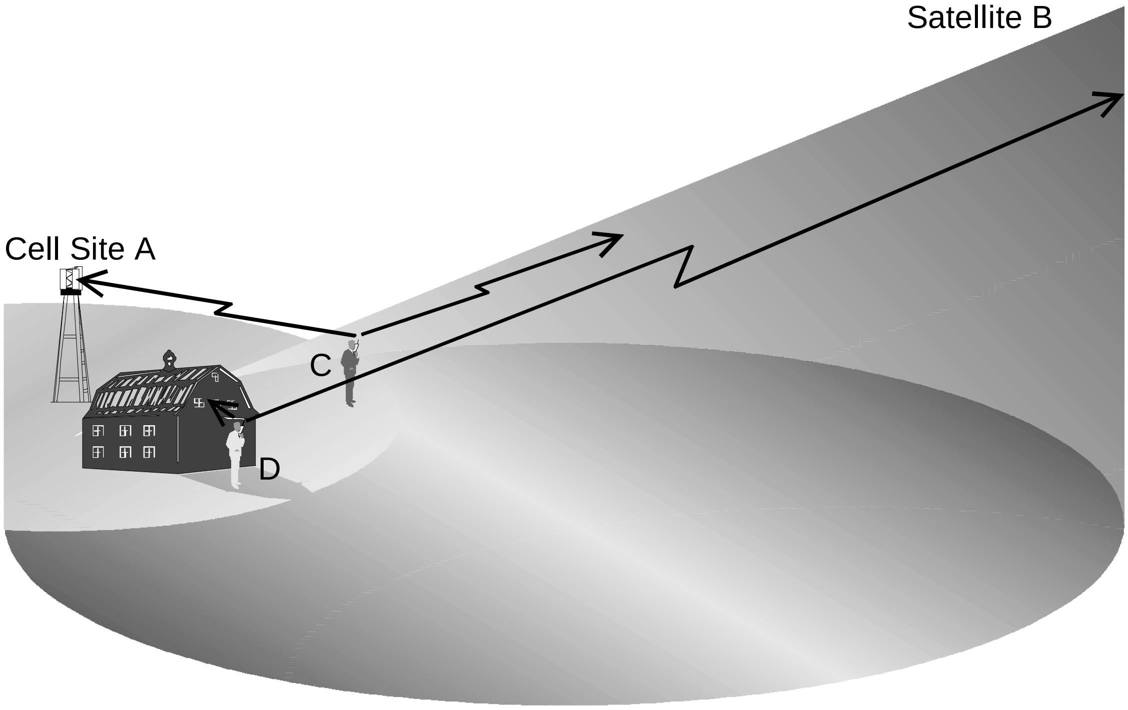

Figure 52 Uplink carriers from mobile terminals, one in a shadow, the other using a terrestrial base station

Figure 52 is the equivalent of figure 42 where cell site A is now a terrestrial cell site, not a satellite. If power control is to be used on Uplink Traffic Channels (as it almost certainly would be) then the transmission from C will be just sufficient for communications to cell site A. It is highly unlikely that satellite B would be able to receive the signal and so there is no means of either B or D determining that this Traffic Channel is in use between A and C. If satellite B was to start using the Downlink Traffic Channel to communicate with D then that would cause unacceptable interference at C, so this channel assignment is unsatisfactory. Hence terrestrial and satellite carriers must be segregated because the received power levels of satellite and terrestrial carriers are very different. If not, a high power transmitter could mistake a carrier occupied by a low power transmission as vacant and cause unacceptable interference to the low power transmission. Even if equal receive power carriers were possible, out of band interference from the numerous terrestrial base stations in a satellite beam might collectively be sufficiently powerful to cause interference when it spills over into a satellite channel.

This scenario could also be a problem in purely terrestrial systems where two or more different types of Traffic Channel carriers are being used which operate assuming different received power levels or bandwidths at the receivers. In this case, it becomes very complicated to determine whether a Traffic Channel can safely be used or not. In such cases it would be much simpler to assign a different pool of Traffic Channels to each type of carrier. It is likely, therefore, that pico cell systems will need to be segregated from macro and micro cellular systems as well. Co-ordination between different systems might be possible if the uplink transmit powers and downlink received powers are made the same for any system. However, this is too fundamental a design restriction for such a diverse set of systems as FPLMTS. For example with satellites it implies restricting the spot beam sizes and antenna sizes, preventing the adoption of improved satellite technologies as they become available.