CHANNEL ALLOCATION & POWER CONTROL IN IS-136

BY DAVID PANDIAN

A thesis submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Master of Science

Graduate Program in Electrical and Computer Engineering

Written under the direction of

Professor Roy D. Yates

and approved by

____

____

____

____

New Brunswick, New Jersey

May, 1999

ABSTRACT OF THE THESIS

Channel Allocation & Power Control in IS-136

by David Pandian

Thesis Director: Professor Roy D. Yates

In the North American TDMA digital cellular standards , the downlink is constrained in that the base station transmit power has to be kept constant over a frame. This implies that the base has to transmit over the entire duration of a frame even if there is just one active user in that frame. This introduces additional interference on the downlink in the unoccupied slots in a partially occupied frame and hence substantially reduces capacity. Ideally, we would like to vary the base station transmit power every slot and turn off during idle slots.

We evaluate the penalty in keeping the base transmit power constant over a frame as opposed to the ideal scheme of varying the powers every slot. We explore schemes like carrier packing to reduce the additional downlink interference and investigate various combinations of channel allocation, power control and handoff algorithms and their effectiveness in mitigating the above problem.

Acknowledgements

I am extremely grateful to Professor Roy Yates for his invaluable guidance and encouragement all the way and for giving me the opportunity to be a part of his research. I would like to thank Professor David Goodman for kindling my interest in wireless communications and getting me started on my research at WINLAB. A special note of thanks to Dr Mandayam, Dr Ogielski and everyone in the WiPPET group for all the help and valuable insights and to everyone at WINLAB for the wonderful atmosphere to work in

Dedication

To God for making all things possible, Mum & Dad for all the sacrifices and a very special friend for always being there for me

Table of Contents

List of Tables

List of Figures

Chapter 1

Radio Resource Management in North American TDMA (IS-136)

1.1 Introduction

The driving force behind cellular systems is the need for increased capacity. Operators of existing cellular systems have been able to reuse the available spectrum more efficiently by adopting new transmission technologies. One of the technologies proposed was Time Division Multiple Access in which each frequency channel is further multiplexed in time to accomodate more users. The first version of a North American digital cellular system employing TDMA was the Interim Standard 54 (IS-54) which used the AMPS analog control channels. IS-136, also called the North American TDMA system(NA-TDMA) is a further refinement of IS-54 which specifies a digital control channel and hence incorporates features like short message services and sleep mode to conserve battery power etc. In IS-136, each 30KHz channel in the older analog system AMPS is multiplexed in time into six half rate slots and shared by three users to increase capacity. These six slots comprise a frame and each user transmits only in the two slots assigned to him in each frame. These two half rate slots comprise one full rate digital traffic channel. A channel in IS-136 is thus completely specified by a combination of a frequency carrier and a time slot.

Capacity is directly tied to spectrum efficiency which measures how efficiently we can reuse the available bandwidth while keeping the interference at tolerable levels. This is the essence of radio resource management in all cellular radio systems. In IS-136, interference is a major issue, especially on the downlink because the base station transmit

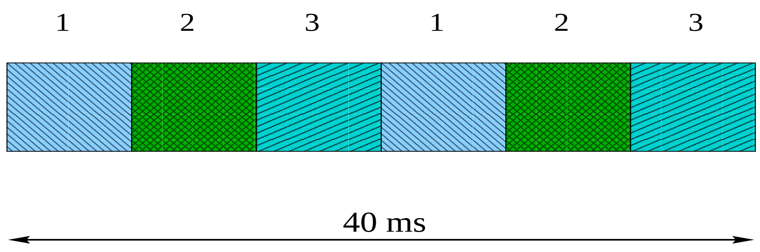

Figure 1.1: An IS-136 frame comprises six half rate slots which are shared by 3 users. The frame length is 40ms

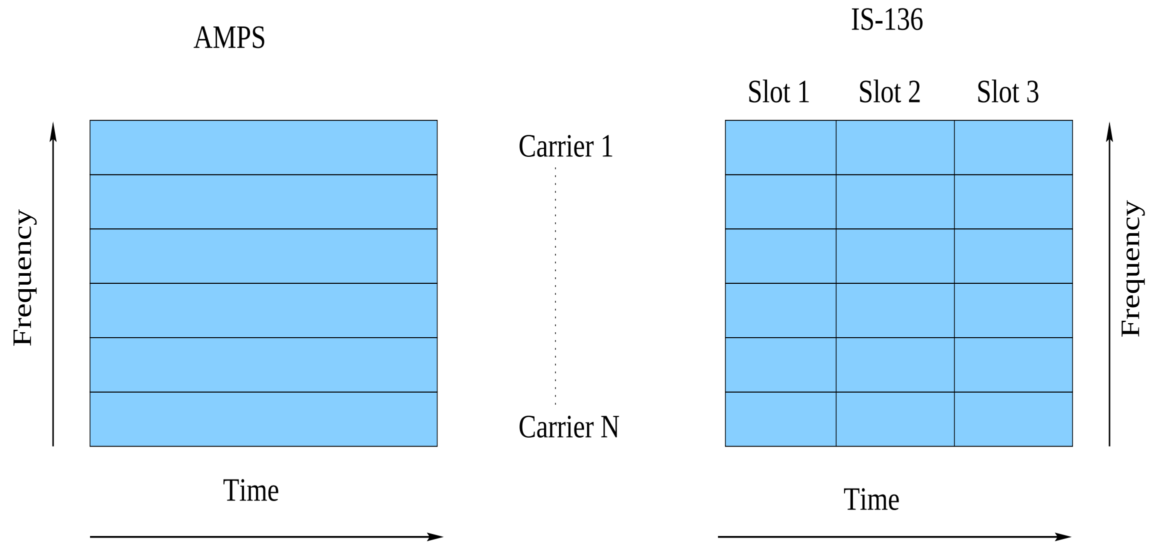

Figure 1.2: The concept of a channel is different in AMPS and IS-136. In AMPS a channel is a 30KHz frequency band. In IS-136, a channel is a time slot within that frequency band. Each IS-136 carrier can thus carry 3 users as opposed to just 1 in the case of AMPS. The six half-rate slots are shown as three full rate slots.

power has to be kept constant over a frame during digital mode operation [1]. A consequence is that the base has to transmit in the idle slots if there is atleast one occupied time slot in a carrier. This introduces additional interference on the downlink.

Earlier studies [2] have addressed this problem but within the context of fixed channel allocation and without considering the effect of mobility, handoffs and power control. An adaptive channel allocation algorithm similar to channel segregation [3] has been studied in [4] but the effects of various combinations of channel allocation, power control and handoff schemes as well as factors like carrier packing have not been considered.

Efficient channel allocation and power control schemes help reduce co-channel interference and this study explores various combinations of channel allocation and power control algorithms and their effectiveness in mitigating the effects of the downlink constraint in an IS-136 system. We also evaluate the penalty we have to pay in keeping the base station transmit power constant over a frame as opposed to varying it on a slot by slot basis.

1.2 Power Control

Power control reduces interference to other mobiles or bases transmitting on the same channel (co-channel interference) in other cells thereby ensuring that the signal quality is maintained above a certain threshold on all links. It has been shown that power control in conjunction with dynamic channel allocation improves the capacity of the system [5, 6].

Uplink Power Control: Power control has been an actively researched topic in cellular mobile systems with the initial focus being on maintaining a constant received power level at the receiver [7, 8]. This was found to be effective in providing better adjacent channel protection especially in spread spectrum systems but not too effective in reducing co-channel interference. SIR balancing schemes which strive to attain a common SIR over all radio links were found to maximize capacity [9, 10] but these centralized schemes required a central controller and was computationally expensive. It has been shown that distributed SIR balancing schemes which use local measurements can achieve similar performance [5, 6].

We employ one such synchronous distributed constrained SIR balancing power control scheme on the uplink as proposed in [11]. The scheme is based on SIR measurements made locally by each base station and does not require a central controller. All mobiles in the system assigned to the same time slot update their powers synchronously to attain a target SIR. The scheme is called constrained since we impose a peak power constraint on the mobile transmitter power. The power control iteration at the ith step is given by

$$P_i = \min\{P_{\max}, P_{i-1} * \frac{\gamma_t}{\gamma_{i-1}}\} \tag{1.1}$$

where γ_t is the target SIR and γ_{i-1} is the current uplink SIR measured by the base station.

Downlink Power Control: In IS-136, the downlink is constrained in that the base station transmit power has to be kept constant over a frame. This implies that the base has to transmit over the entire duration of a frame even if there is just one active user in that frame. The key issue is what that constant power should be since there could be up to three users in each carrier and each of them may require different powers to meet the target SIR on the downlink. Two approaches which are studied are the MAX and MIN schemes which are outlined below. The ideal scheme which is not currently specified in the standards is the PERSLOT scheme.

- **MAX Scheme**: For each occupied slot, calculate the base station transmit power required to meet the target SIR. This computation is carried out using the same synchronous power control iteration described above for the uplink. The base station transmit power over the next frame would be the maximum of the required powers thus computed in the previous frame.
- **MIN Scheme**: For each occupied slot, calculate the base station transmit power required to meet the target SIR. The base station transmit power over the next frame would be the minimum of the required powers thus computed in the previous frame.
- **PERSLOT Scheme**:For each occupied slot, calculate the base station transmit power required to meet the target SIR and update the base station transmit power accordingly. The transmitter at the base is turned off during idle slots

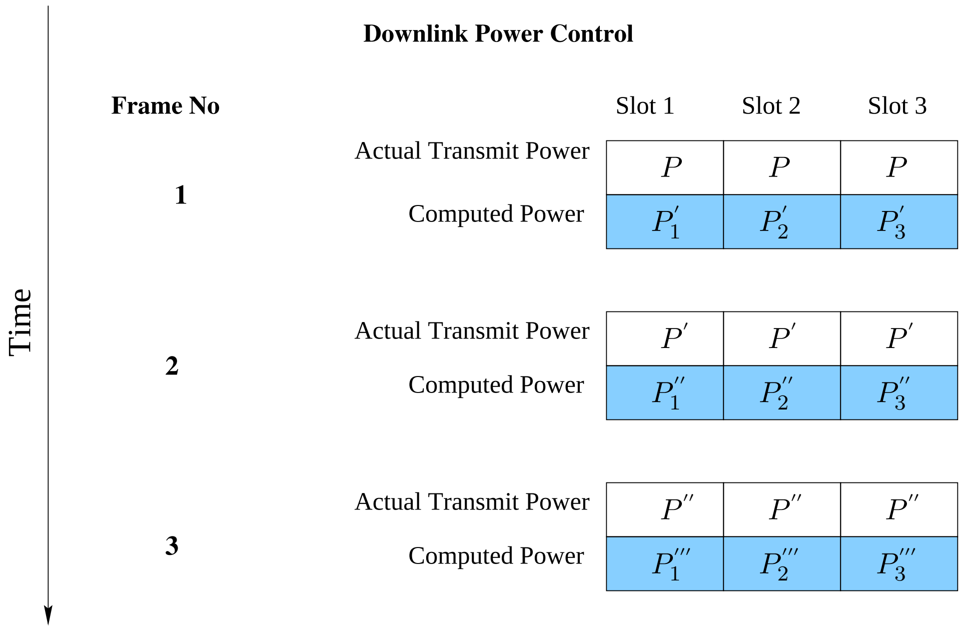

Figure 1.3: Downlink power control using the MAX,MIN or AVG schemes. Lets assume that three new mobiles are assigned to slots 1, 2 and 3 of a carrier with required base station transmit powers P_1, P_2 and P_3. With the MAX scheme, the base station transmit power over the next frame would be $P = \max(P_1, P_2, P_3)$. During this frame, for each occupied slot, calculate the base station transmit power required to meet the target SIR (γ_t). The required base station transmit power P_i' in the i_{th} slot is computed as $P_i' = P * \gamma_{target}/\gamma_{current}$. The base station transmit power over the next frame would be the maximum of the required powers thus computed in this frame i.e $P' = \max(P_1', P_2', P_3')$. This procedure is repeated every frame.

1.3 Channel Allocation

The partially occupied carriers introduce additional interference on the downlink as discussed earlier. One approach to reducing the number of partial slots, and consequently the additional interference on the downlink, is to pack the carriers, by assigning as many mobiles to a single carrier before picking the next carrier for channel assignment as shown in Fig 1.4.

Channel allocation is carried out both during call admission when a mobile is admitted into the system and also during handoffs when it breaks its connection to the old base station and communicates with a new base station. Before we get into the channel allocation algorithms, we describe two mobile assisted procedures which are salient features of IS-136.

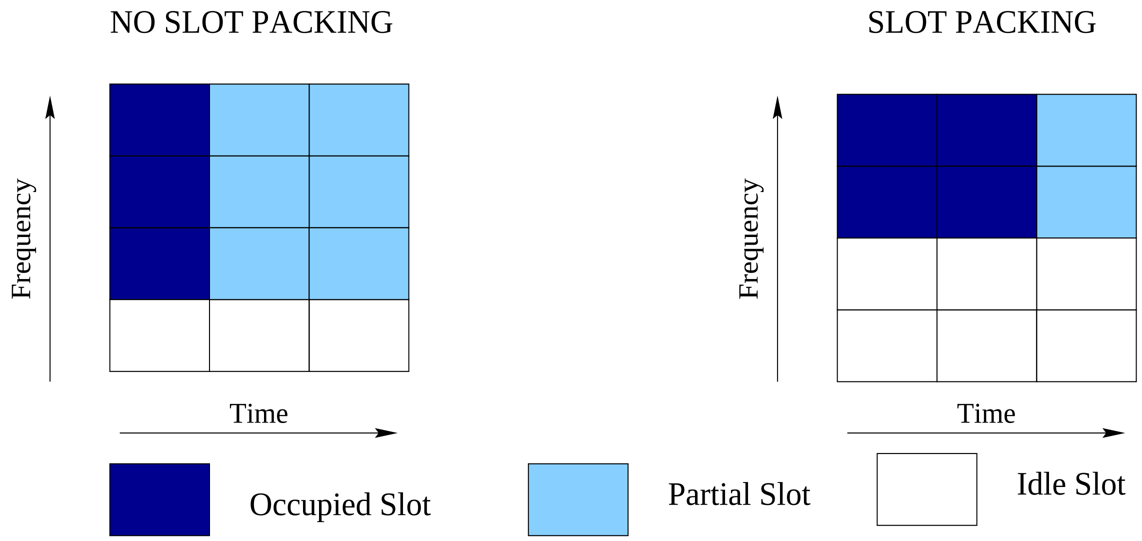

Figure 1.4: Comparison of slot assignment schemes with and without slot packing. A partial slot is a slot which is not occupied but is part of a carrier in which there are occupied slots. Partial slots contribute to additional downlink interference because the base has to transmit during the partial slots too even though they are effectively idle. We observe that the number of partial slots reduces from 6 without packing to 2 with packing thereby reducing the additional interference on the downlink

An advantage of using TDMA is that the mobile terminal does not have to receive or transmit all the time as in AMPS but only during its assigned time slot within a frame. This aspect is exploited in IS-136 and the mobile thus tunes to traffic channels during idle slots and reports signal strength measurements on these channels back to the base. The base station can thus use these measurements in its resource allocation algorithms and optimize the downlink too as opposed to only the uplink in an analog system like AMPS. There are two types of these mobile assisted procedures in IS-136. In both cases, the mobile reports signal strength measurements which the base uses for channel allocation for call admission and handoffs.

Mobile Assisted Channel Allocation: This procedure is carried out by a mobile without a call in progress, which is called the camping state in IS-136. In the camping state, the mobile tunes to and receives information on a digital control channel (DCCH). The base station sends a list of channels to the mobile which includes the current DCCH and optionally other channels. When it is not receiving information on the control channel, the mobile tunes to these channels and reports signal strength measurements back to the base. The base uses these measurements for channel allocation during call admission.

Mobile Assisted Handoff: This procedure is carried out by a mobile with a call in progress. The base station sends to the mobile, a list of forward RF channels. The mobile tunes to these channels during its idle slots in a frame and reports signal strength measurements on these traffic channels as well as word error rate measurements on the current traffic channel. The base uses these measurements for channel allocation during handoff.

A digital traffic channel in IS-136 is a combination of a frequency and a time slot. The base runs through a channel allocation algorithm to pick the best channel from the available resources. There are two broad classes of channel allocation schemes.

Fixed Channel Allocation (FCA): The available spectrum, in terms of the total set of channels available in the system is divided into N smaller groups and reused in cells which are sufficiently separated geographically so as to keep the co-channel interference to desired levels. This ensures that adjacent base stations do not use the same set of channels. The reuse factor is said to be N. FCA however has not been observed to perform efficiently under varying traffic loads since the frequency reuse pattern is fixed by design.

Dynamic Channel Allocation (DCA): Under DCA, each base station has access to all channels in the system. The reuse factor is 1. There are various flavors of DCA. An interference adaptive DCA scheme assigns channels based on real time interference measurements. It has been shown that interference adaptive DCA schemes can substantially increase capacity [12]. In particular, it has been shown that DCA outperforms FCA at light traffic loads although it may inefficiently deploy channels at higher traffic loads [13, 14]. In our study, we investigate a few variations of the interference adaptive DCA scheme.

We have implemented a variant of the Mobile Assisted Channel Allocation (MACA) procedure in IS-136 in the channel allocation algorithms we have developed and which are discussed below. We make the assumption that all channels are digital and that a mobile without a call in progress can make signal strength measurements on forward traffic channels during its idle slots.

- **LI-DCA (Least Interference Based DCA) algorithm:**
  1. The base measures the uplink interference on all channels and compiles a list of the N best uplink channels.
  2. The base requests the mobile to measure the downlink interference on these N channels.
  3. The mobile sends back the downlink interference on these channels to the base.
  4. The base picks the first channel c in the list which meets the SIR threshold on both the uplink and the downlink. The base determines the required transmit powers on the uplink and downlink as follows

$$P_{req}^{(u)}(c) = 10^{\frac{\gamma_{new(up)}}{10}} \cdot \frac{gain^{(u)}}{P_I^{(u)}} \tag{1.2}$$

$$P_{req}^{(d)}(c) = 10^{\frac{\gamma_{new(down)}}{10}} \cdot \frac{gain^{(d)}}{P_I^{(d)}} \tag{1.3}$$

     where $P_I^{(u)}$ is the uplink interference value on channel c for the `MS` at the parent `BS` and $P_I^{(d)}$ is the downlink interference value on channel c at the `MS` and $gain^{(u)}$ and $gain^{(d)}$ are the uplink and downlink gains repectively. Note that the simulation parameter $\gamma_{new(up)}$ is the required uplink signal to interference ratio and $\gamma_{new(down)}$ for the downlink.

- **Power Differential Algorithm:**
  1. The base sorts the channels in increasing order of uplink interference. It then separates them into two sets, slots in partially occupied carriers and slots in idle carriers. It then compiles a list of N channels with as many partial slots as possible and then fills up the rest with slots in idle carriers.
  2. The base requests the mobile for the downlink interference on these N channels.

3. The mobile sends back the downlink interference on these channels to the base. The base picks those channels from this list which meet the SIR threshold on both the uplink and the downlink according to equations 1.2 and 1.3.

4. For each channel i, the base computes a factor called the power differential for each of these channels as

$$\Delta P(i) = P(i) - P'(i) \tag{1.4}$$

where

$P(i)$ = Transmit power on the carrier, the channel is located in

$P'(i)$ = Required base station transmit power on the channel

5. These channels are now sorted in decreasing order of the power differentials $\Delta P(i)$. This ensures that the positive power differentials, implying no additional interference, come out on top. The negative power differentials, implying either idle carriers or partial carriers in which additional interference is going to be introduced, come later. This ensures that if we cannot get a channel that introduces no additional downlink interference, then we pick the channel which introduces the least additional interference irrespective of whether its in a partially occupied carrier or an idle carrier.

1.4 Call Admission

When a mobile initiates a new call, the mobile first scans a block of digital control channels and locks on to the one with the strongest received signal strength (RSS). The mobile thus communicates with the base from which it receives the strongest pilot tone. During call setup, the base runs through one of the channel allocation algorithms mentioned above and assigns a particular time slot on a particular carrier frequency from a pool of channels available at that base. We use an admission SIR threshold policy to determine if a call is admitted or blocked. A call is admitted only if both the mobile and the base can achieve the call admission SIR threshold γ_{new} on the uplink and downlink respectively with transmitter powers less than their respective peak power constraint, P_{max}.

1.5 Handoffs

One of the features of IS-136 is mobile assisted handoffs (MAHO) as described under channel allocation. There are different types of handoffs depending on the criterion used to initiate a handoff. We investigate two kinds of handoffs in this study.

Inter-Cell Handoffs: If the RSSI of a neighboring base is higher than that of the current base by a hysteresis threshold T_{hy} dB, an inter-cell handoff is initiated to the new base. Since the RSSI is directly related to the distance of the mobile from the base station, an inter-cell handoff occurs at cell boundaries. A handoff is successful only if the new base can run through the channel allocation algorithm and find an available channel and the mobile and new base station can achieve the hand-in SIR threshold on the uplink and downlink respectively with transmitter powers less than their respective peak power constraint, P_{max}.

Intra-Cell Handoffs: As opposed to inter-cell handoffs in which RSSI is used to hand the mobile off to a new base station, intra-cell SIR based handoffs try to transfer the mobile to a channel with a higher signal quality in the same base station.There are two scenarios which trigger intra-cell handoffs in our simulations. In the first scenario, a handoff is initiated if the uplink or downlink SIR of the active channel falls below a threshold for a specified period of time called the outage time. The second scenario involves slot packing. In order to reduce the number of partial slots, we try to repack the carriers in each base on a slot by slot basis and thus mobiles are moved to new channels to fill up vacant slots in partially occupied carriers.

Chapter 2

WiPPET$_{\mathrm{IS-136}}$: A parallel simulator for IS-136

2.1 Introduction

The Wireless Propagation and Protocols Evaluation Testbed(WiPPET [15]) is a virtual testbed developed at WINLAB for parallel simulation of wireless networks. The simulator has been written in the Telecommunications Description Language (TeD) developed at Georgia Tech [16, 17]. WiPPET [15] models session or call level traffic, distance loss and shadow fading propagation, mobility and resource management algorithms-call admission, power control and handoff algorithms. It also provides system performance statistics in terms of call blocking, call dropping and average number of handoffs per call.

WiPPET$_{\mathrm{IS-136}}$ is a flavor of WiPPET [15] which has been modified to simulate IS-136. Significant changes which have been incorporated in WiPPET$_{\mathrm{IS-136}}$ include

1. A time slot structure has been introduced and mobiles are active only in their assigned time slots.

2. Downlink power control has been introduced, since suppressing downlink interference is our primary objective as described earlier. The MAX and MIN schemes for downlink power control have been implemented.

3. Call admission algorithms have been modified to include admission thresholds on both the uplink and downlink. Mobile Assisted Channel Allocation(MACA) which is a feature of IS-136 has been incorporated [18].

4. Handoff algorithms have been modified to include hand-in thresholds on both

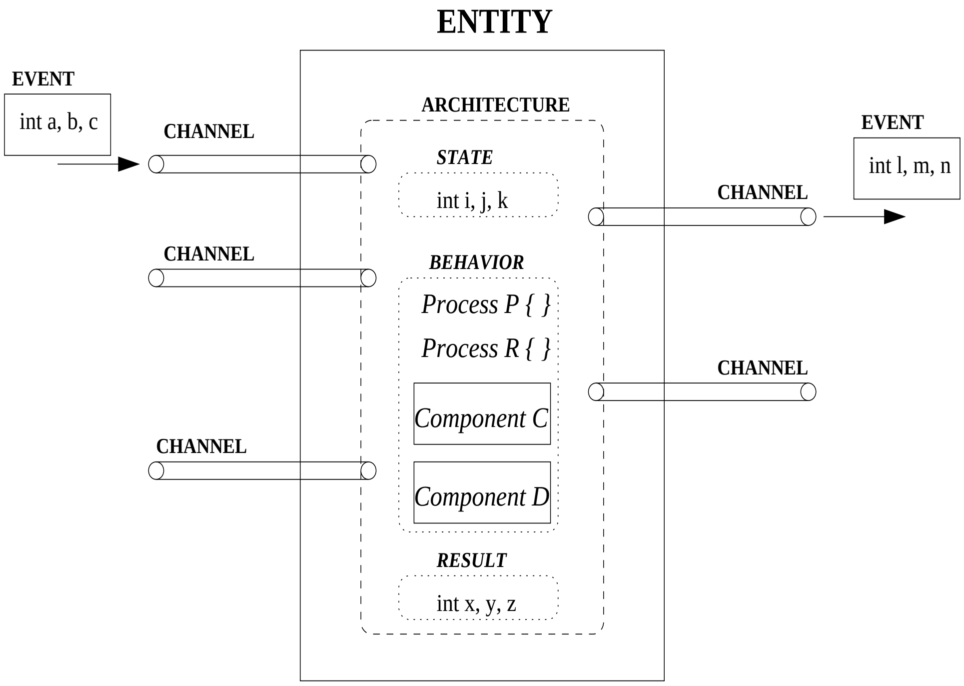

Figure 2.1: Basic framework of a TeD entity. An entity has an architecture which contains states, processes, IN (incoming) and OUT (outgoing) TeD channels, etc. A process of the architecture sends or receives events on these TeD channel.

the uplink and downlink. Mobile Assisted Handoff(MAHO) which is a feature of IS-136 has been incorporated [1].

5. Channel allocation algorithms such as the power differential algorithm which has been described in section 1.3 have also been incorporated.

6. Features pertinent to IS-136 like slot repacking have been implemented.

7. System performance statistics like time slot occupancy, SIR and transmit power statistics are computed.

2.2 Fundamentals of TeD

This section gives a brief overview of the Telecommunications Description Language(TeD) and explains some of the TeD terminology which we will be referring to in the rest of the chapter.

Entities: Since TED has been tailored to model telecommunications networks [16], it is very easy to model the different components of a cellular network as entities. In our case, we model mobiles, bases, switches etc as separate entities. Entities are highly modular. Their states and behavior are highly encapsulated and their only interaction with other entities is through TED *channels*.

Channels: TED channels are the only means of dynamic interaction between entities. An output TED channel of one entity may connect to the input TED channel of another entity. This is called an external TED channel. The TED channels of an entity can also be mapped to themselves, in which case they are called internal TED channels. Entities share data between themselves by sending *events* on TED channels.

Events: Events are information units which flow through TED channels. An event which is sent on an external TED channel is called an Inter-Entity event(Inter-EI event) and an event which is sent on an internal TED channel is called an Intra-Entity event(Intra-EI event).

Architecture: The core of each entity is the architecture which consists of a set of one or more processes. A process defines the actions which the entity should take in response to event arrivals on its input TED channels. The processes generate events on one or more output channels as part of their action. Processes within an entity exchange events through internal TED channels. Architectures can be inherited as we shall see later in this chapter. In WiPPET, each layer of the protocol stack is modeled as an architecture which inherits from the architecture(protocol) immediately above it.

2.3 Entities in WiPPET$_{\mathrm{IS-136}}$

Since TED has been tailored to model telecommunications networks [16], it is very easy to model the different components of a cellular network as entities. The entities in WIPPET$_{\mathrm{IS-136}}$ are the same as in WiPPET [15].

1. Mobile Station (`MS`): This entity participates in call admission, power control and handoff protocols. It also schedules power control updates on the downlink.

2. Base Station (`BS`): This entity participates in call admission, power control and handoff protocols. It runs through a specified channel allocation algorithm during call admission and handoffs and assigns a channel to the mobile. It also schedules power control updates on the uplink, as well as handoff checks.

3. Mobile Switching Center (`switch`): This entity takes part in call admission and handoff algorithms.

4. Call generator and statistics collector (`callgen`): This entity generates session or call traffic. It also collects system performance statistics like call blocking, call dropping, average number of handoffs per call etc.

5. Geography zone (`zone`): This entity represents the geography of the network as a collection of points along with their link gains to all base stations. One of the primary tasks of the zone is reporting interference measurements back to the base and mobile. It takes part in mobility, call admission, power control and handoff protocols.

A more comprehensive description of entities in WiPPET is presented in [15].

2.4 Architecture of Entities

Since WiPPET$_{\mathrm{IS-136}}$ is derived from WiPPET, it is structured very similarly. The general hierarchy of the WiPPET$_{\mathrm{IS-136}}$ entity architecture is shown in Fig 2.2. There are eight distinct layers in the hierarchy. Protocols are implemented by inheritance of architectures. Each of the entities described above has a protocol stack with all or a subset of the protocols depending on the type of entity. Each layer or protocol in the protocol stack has a distinct architecture which inherits the variables and functions from the layer above it in the protocol stack. The other design option is to model each protocol as a sub-entity within the larger entity. The reason for opting for inheritance

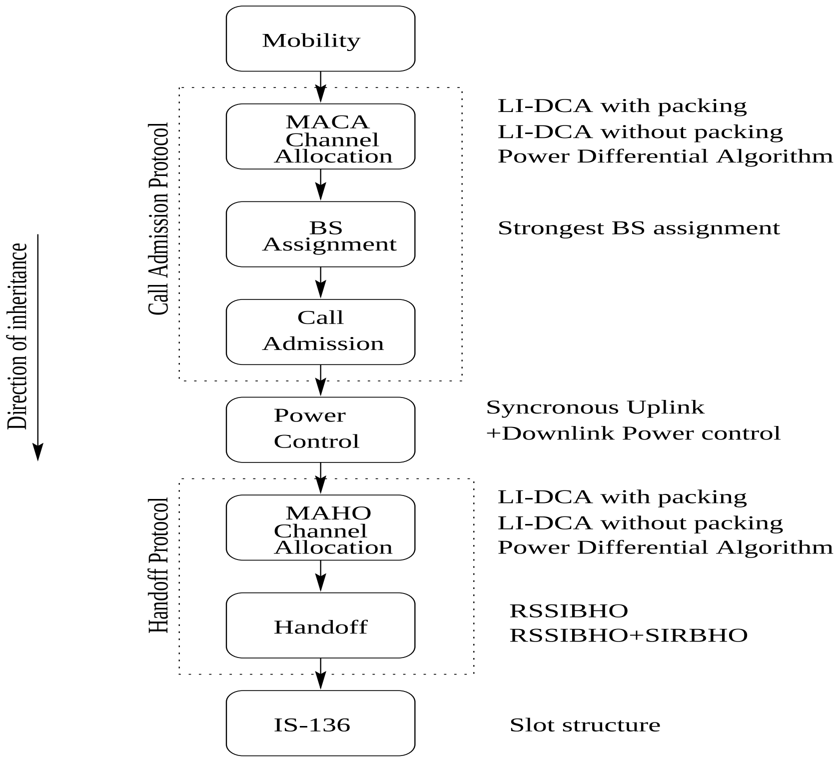

Figure 2.2: General hierarchy of inherited architectures in an entity. Down arrows show the direction of inheritance in hierarchy of architectures. Mobility is the primary architecture and IS-136 inherits all states and processes defined in architectures above it.

of architectures is that modelling each protocol as a sub-entity would involve exchange of events between the sub-entities to share data, and events are quite expensive in TeD. WiPPET$_{\mathrm{IS-136}}$ has an additional IS-136 protocol at the bottom of the stack as shown in Fig 2.2 to incorporate IS-136 functionality like synchronizing time slots, updating the carrier powers, collecting time slot statistics, slot repacking etc. Call admission, handoff and channel allocation algorithms for IS-136 were implemented by modifiying the respective architectures in each of the entities. Since the mobile assists in channel allocation during call admission and handoffs, the protocol stack for the mobile in WiPPET$_{\mathrm{IS-136}}$ includes modules for channel allocation during call admission and handoffs.

2.5 Protocols and Algorithms

WiPPET$_{\mathrm{IS-136}}$ models session or call level traffic with mobility, call admission, power control and handoff protocols within the context of a time slotted structure to simulate IS-136. Most of these protocols are discussed in sufficient detail in [15]. The algorithms and protocols which were developed or modified for IS-136 are discussed in this section.

2.5.1 IS-136

This is the lowermost layer in the inheritance hierarchy discussed in section 2.4. It makes higher layers in the protocol stack like power control and handoffs work within the context of a time slotted frame structure. It manages system wide time slot synchronization, collects time slot occupancy statistics, updates carrier transmit powers on the downlink and runs through a slot repacking algorithm to determine slots which can be moved to pack carriers and reduce partial slots. Each of these sub-tasks is discussed below

Time Slot Synchronization: For simplicity, the six half-rate slots in IS-136 have been reduced to three full rate slots in our simulations. Every **Slot_Length** secs(13 ms in our simulations), the IS-136 architecture in the switch sends out

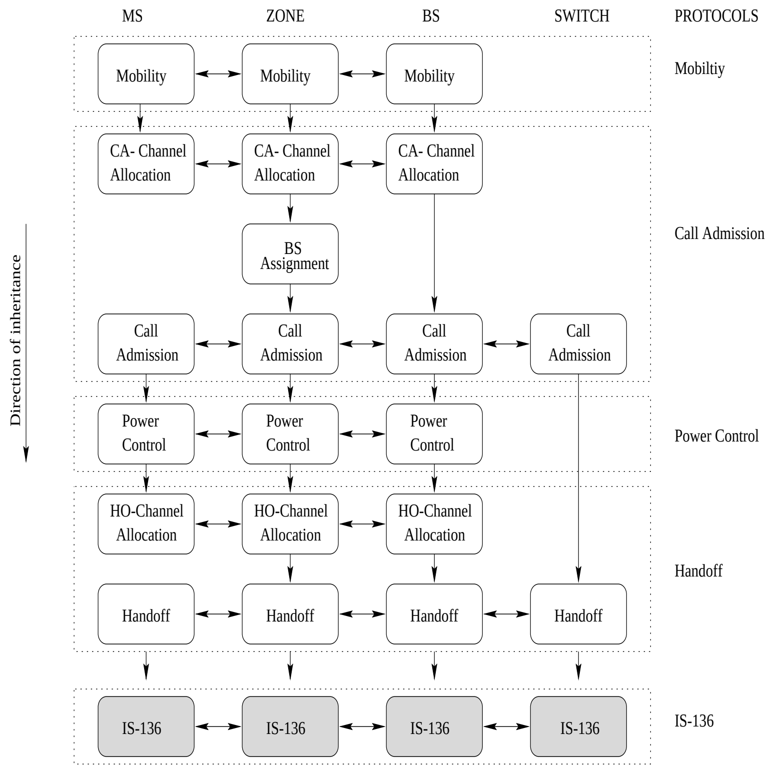

Figure 2.3: Hierarchy of inherited architectures in WiPPET$_{\text{IS}-136}$ entities. Vertical down arrows show the direction of inheritance. Each entity supports one or more protocols, i.e. protocol functionality distributed among entities. For example, `zone` has all eight layer of architecture inheritance, so, it support all protocols of WiPPET$_{\text{session}}$. On other hand, `switch` doesn't support power control and mobility protocols. Protocols are presented by one or more architectures in an entity. Horizontal arrows show the communication between one architecture to its peer architecture in other entities. Shaded architectures are the last architectures in the hierarchy and bound to their respective entities.

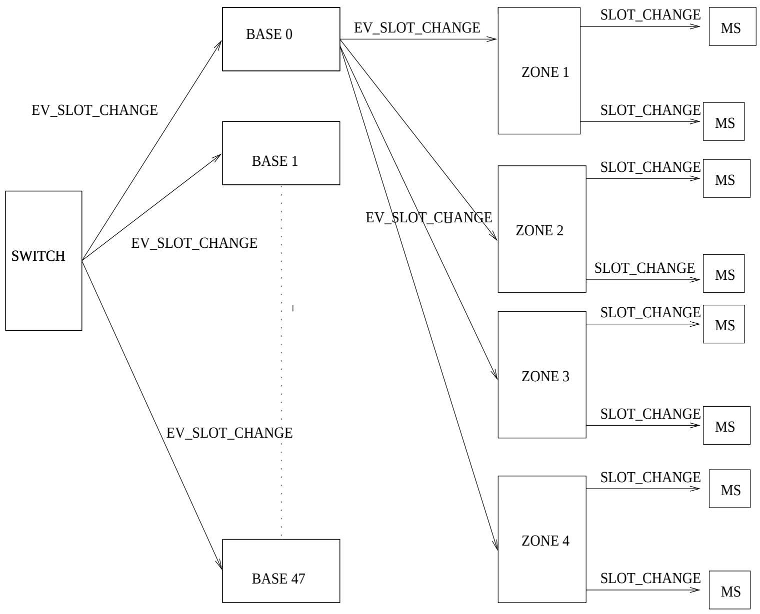

Figure 2.4: Time Slot Synchronization in WiPPET$_{\text{IS}-136}$

an EV_SLOT_CHANGE event to the IS-136 architecture in all base stations. The current slot number is sent along with the event. The slot counter is maintained by the switch and cycles from 1 to 3. Each base thus keeps track of the current slot number in the frame. One of the bases, BS 0 for convenience, sends the event to all zones who in turn update their slot number and send the event to all active MSs in their geography. Thus all entities in the system keep track of slot changes and maintain time slot synchronism as shown in Fig 2.4. System wide time slot synchronism is one of the key assumptions in our simulations.

Carrier Transmit Power Updates: When the event EV_SLOT_CHANGE is received, each BS updates the transit powers on all its carriers if the slot number is 1, indicating the beginning of the frame. The transmit powers are updated, depending upon the downlink power control algorithm specified. These algorithms were described earlier in section 1.2.

Scheduling Power Control and Handoff Checks: When the event EV_SLOT_CHANGE is received, each BS runs through the list of MSs it communicates with and compares the assigned slot number for each of them with the system slot number. If

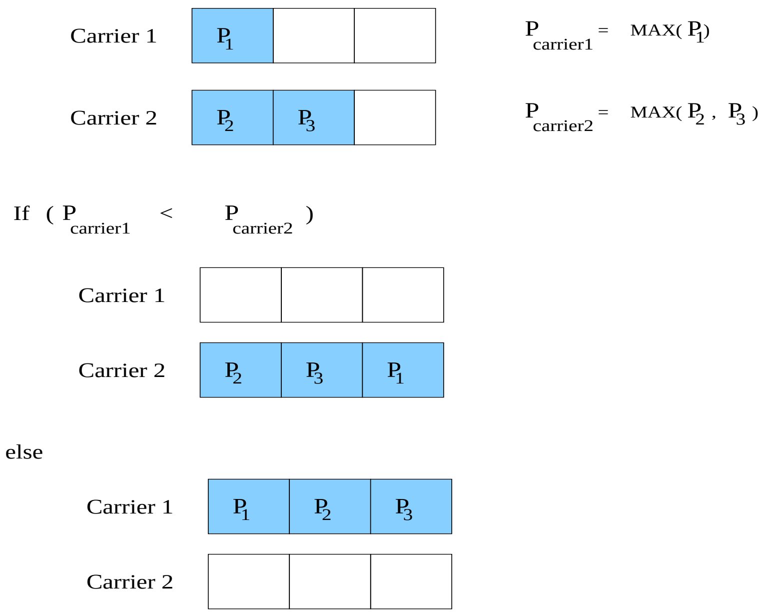

Figure 2.5: Each BS periodically scans the carriers and tries to combine a pair of carriers with 1 slot and 2 slots occupied respectively into a carrier with 3 slots occupied. The algorithm tries to rearrange slots such that the carrier transmit power remains the same

they match, the MS is active in that slot. The BS then sends an EV_UPPWR_CRTL event to the power control architecture to schedule a power control iteration on the uplink for that MS. A EV_HANDOFF_CHECK is also sent to the handoff architecture to trigger a handoff check for the MS. A handoff check determines if a handoff is necessary and if so, initates a handoff. An uplink power control iteration and handoff check are thus scheduled by the BS for all its MSs which are active in that slot. When the MS receives EV_SLOT_CHANGE, it similarly determines if it is supposed to be active in that slot and sends the EV_DWNPWR_CTRL event to its power control architecture to schedule a downlink power control iteration.

Slot Repacking: One of the ways to reduce the number of partial slots and hence the additional downlink interference is to pack the carriers. Although our channel allocation algorithms try to pack the carriers as best as possible, the packing is not optimum because handoffs and call dropping introduce vacant slots in carriers. We thus try to enhance the carrier packing by searching for partially occupied carriers periodically and trying to move mobiles to these vacant slots.

At the beginning of each slot, each BS runs through a slot repacking algorithm and compiles a list of source and target channel pairs. A source channel is a currently occupied slot which could be moved to a target channel which is a vacant slot in a partially occupied carrier. The algorithm is described below:

1. The BS scans all carriers and compiles two lists, carriers with 1 slot occupied and carriers with 2 slots occupied.

2. Each of the lists is then sorted in decreasing order of carrier transmit powers.

3. Pick the first element from both lists. We have a carrier i with 1 slot occupied and another carrier j with 2 slots occupied.

   If $P_c(i) < P_c(j)$

   $$source\_slot = OCCUPIED\_SLOT(Carrier_{1slot})$$

   $$target\_slot = VACANT\_SLOT(Carrier_{2slot})$$

   else

   $$source\_slot = MAX\_SLOT(Carrier_{2slot})$$

   $$target\_slot = VACANT\_SLOT(Carrier_{1slot})$$

   where $P_c(i)$ is the transmit power on carrier i.

   This ensures that moving a mobile to another carrier does not increase the carrier transmit power on that carrier and hence the downlink interference.

4. We thus compare corresponding elements from the lists and compile a list of source and target channel pairs. The two lists may be of unequal length. In such cases, we do the pairwise comparison between the two lists only until the point where they are of equal length. We then continue with the longer list , running through step 3 as before but with both the carriers from the same list. If the list containing carriers with 2 slots is longer, we try to convert 2 slot carrier pairs into a 3 slot carrier and 1 slot carrier. If the list containing carriers with 1 slot is longer, we try to convert 1 slot carrier pairs into a 2 slot carrier and an idle carrier.

Collecting Time Slot Statistics: Each BS determines the number of idle slots, occupied slots and partial slots at the beginning of each frame and sends them to

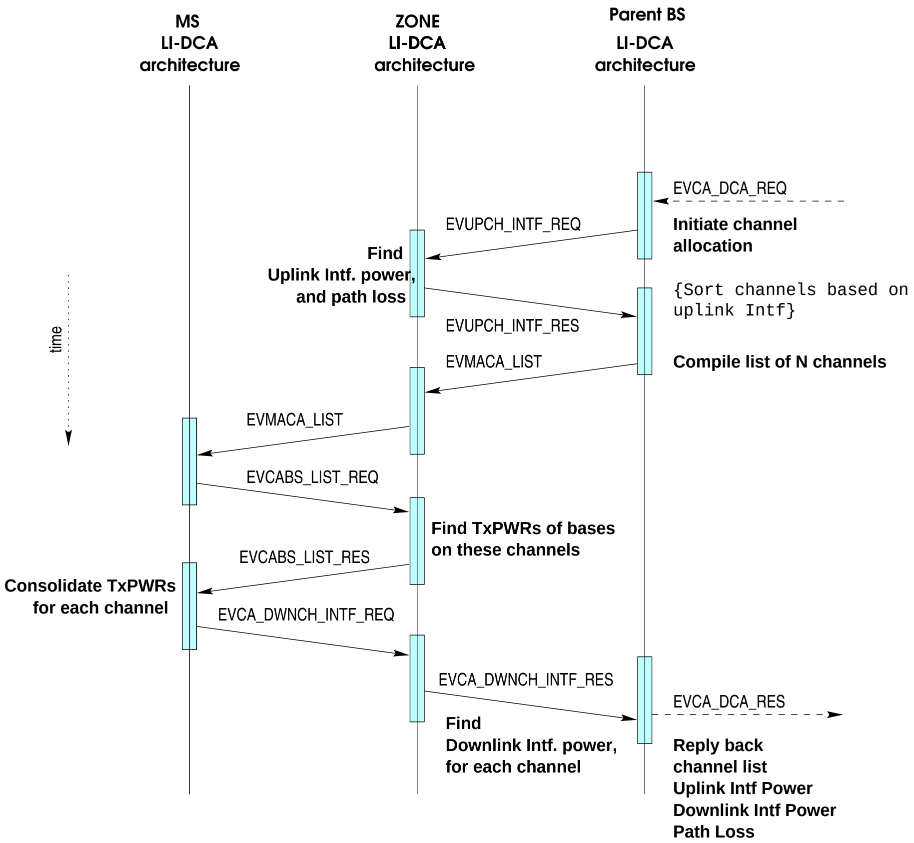

Figure 2.6: Event flow in least interference based dynamic channel allocation (LI-DCA) algorithm. Vertical shaded strips represent processes of channel allocation architectures in a `zone` and a `BS`, which receive and/or send Inter-EI and/or Intra-EI events. Solid-line arrows show the direction of flow of Inter-EI events from a channel allocation architecture to its peer in other entity. Broken-line arrows show the direction of flow of Intra-EI events from the channel allocation architecture in `BS` to an architecture within an entity.

the `switch` as part of a EV_SLOT_STATISTICS event. The `switch` maintains a running average of these statistics.

2.5.2 Call Admission

A `MS` entering the system has to be initially assigned a `BS`. When the `MS` initiates a call, it has to be then assigned a radio channel by the `BS`. In our simulations, base station assignment and channel allocation are carried out together as part of call admission. We do not simulate control channels and mobiles moving around without a call in progress. WiPPET$_{\mathrm{IS-136}}$ simulates two channel allocation algorithms: (1)Least-Interference Based Dynamic Channel Allocation and (2) The Power Differential Algorithm. These algorithms are discussed in more detail later in the section. Base station assignment is done using the strongest pilot algorithm. Although the channel allocation algorithms in WiPPET$_{\mathrm{IS-136}}$ are different from those in WiPPET, the overall event flow and interaction diagrams for call admission remain the same and are presented in detail in [15].

Mobile Assisted Channel Allocation (MACA): WiPPET$_{\mathrm{IS-136}}$ implements a variation of the MACA procedure specified in [18]. The BS sends an internal event EV_DCA_REQ to the channel allocation architecture. The channel allocation architecture then sends an EVUPCH_INTF_REQ event to all the zone entities. The zones send back uplink interference measurements on all channels to the BS. The BS sorts these channels in increasing order of uplink interference. The BS then compiles a list of N channels to be sent to the mobile. The way in which the list is compiled depends upon whether the channel allocation algorithm used employs slot packing or not. When LI-DCA without slot packing is specified, the BS picks the first N channels from the sorted list. If the power differential algorithm or LI-DCA with slot packing is used, the BS scans the sorted list and separates them into two sets: slots in idle carriers and slots in partially occupied carriers. It then compiles the list of N channels with as many partial slots as possible and the rest with idle slots. The BS then sends the list of channels as part of a EV_MACA_LIST event to the parent zone which in turn sends it to the MS.

To determine the downlink interference on a particular channel, we need the transmit powers of all BSs active on that channel. Since the BSs are distributed across all zones, we need to query all the zones to get the consolidated list of BSs transmitting on that channel. The MS sends the list of channels in a EVCABS_LIST_REQ event to all zones and receives in turn a list of transmit powers of BSs transmitting on each of those channels from each zone in a EVCABS_LIST_RES event. The MS consolidates the transmit powers into a single list and sends them along with the channel list to the parent zone as part of a EVCA_DWNCH_INTF_REQ event. The zone calculates the downlink interference on each of the channels and sends the interference measurements back to the BS as part of a EVCA_DWNCH_INTF_RES event. The BS uses these interference measurements to run through the channel allocation algorithms described earlier in section 1.3 and assign a channel to the MS.

2.5.3 Strongest Pilot Base Station Assignment

A MS entering the system has to be assigned a BS to communicate with. In our simulations, the BS with the strongest RSSI signal at the MS is chosen. The event flow and interaction diagrams are presented in [15].

2.5.4 Mobility

The MS moves around in a Manhattan style grid(2-D Macrocell). The mobility is modeled as a biased random walk. The speed of the MS is a simulation parameter. A more comprehensive description of the geography model and the mobility mechanism is given in [15].

2.5.5 Uplink Power Control

In the synchronous uplink power control protocol, all mobiles in the system which transmit during a particular slot update their tranmsitter powers simultaneously at the beginning of the slot.

When a BS receives an EV_SLOT_CHANGE event, it determines the MSs in its domain which are active in that slot. It then sends an EVUPPWR_CRTL event for each of the MSs to the power control architecture to schedule an uplink power control update for each of them. The power control architecture then sends EVUPPWR_CRTL_REQ to the zone. If there is more than one zone, then it sends it to all zones. Here, we assume a

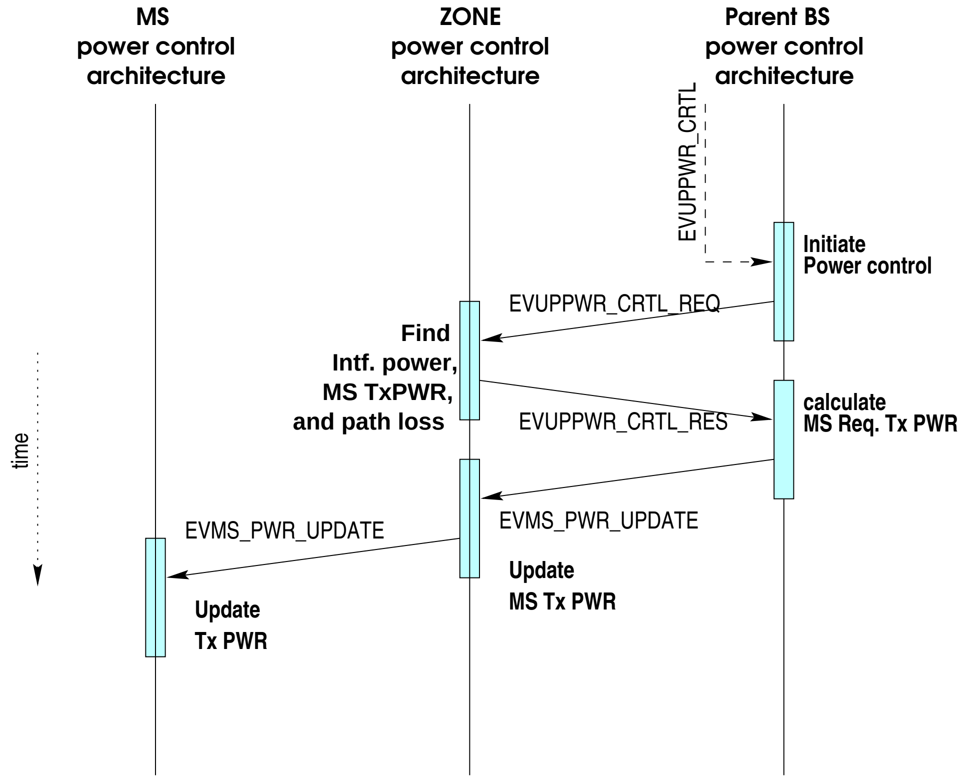

Figure 2.7: Event flow and entity interaction diagram of synchronous distributed uplink power control protocol. A vertical shaded strip represents a process of a power control architecture in an entity, which receives and/or sends Inter-EI events to the processes of its peer architectures or Intra-EI events to the processes of architecture within its entity. Solid-line arrows show the direction of flow of Inter-EI events from a power control architecture to its peer in other entity. Broken-line arrows show the direction of flow of Intra-EI events from a power control architecture to an architecture within an entity

one `zone` simulation. The `zone` sends back EVUPPWR_CRTL_RES including the uplink path loss value $gain^{(d)}$ from the `MS` to its parent `BS` and the interference value $P_I^{(u)}$ on the `MS` uplink channel. The `BS` receives EVUPPWR_CRTL_RES and calculates the new required `MS` transmitter power, $P_{req}^{(u)}(c)$, to meet the uplink target SIR. The required `MS` transmitter power is calculated as follows:

$$P_{req}^{(u)}(c) = 10^{\frac{\gamma_{t(up)}}{10}} \cdot \frac{P_I^{(u)}}{gain^{(u)}}$$

The new required `MS` transmitter power is computed as:

$$P_{new}^{(u)}(c) = \max\left(\mathtt{MIN\_MS\_PWR}, \min\left(\mathtt{MAX\_MS\_PWR}, P_{req}^{(u)}(c)\right)\right)$$

Note that `MIN_MS_PWR`, the minimum `MS` transmitter power, is 0 Watt, `MAX_MS_PWR`, the maximum MS transmitter power, is 1 Watt, and $P_{new}^{(u)}$ is the required MS transmitter power on the uplink channel.

In a multiple zone simulation, there may be one or more MSs present in every zone using the same uplink channel. Here we consider the fact that a MS transmitting to its parent BS on a radio channel in one zone introduces uplink interference on the same radio channel at BSs in other zones. Hence, to take into account the interference due to the interfering MSs, the parent BS requests all zones for interference values on a given MS uplink channel. The parent BS sums up the interference values from all zones and finds the total interference value on a MS uplink channel, $P_I^{(u)}$. Then, the parent BS computes the new required MS transmitter power to meet the uplink target SIR. The BS sends EVMS_PWR_UPDATE with the new required MS transmitter power to the parent zone. The zone updates the MS state and forwards the updated power to the MS using EVMS_PWR_UPDATE event.

2.5.6 Downlink Power Control

Downlink power control iterations are scheduled by the MS. When a MS receives the EV_SLOT_CHANGE event, it sends an EVDWN_PWR_CRTL event to the power control architecture. On receiving this event, the power control architecture in the MS sends an EV_BS_LIST_REQ event to all zones requesting a list of other BSs transmitting on the same channel as the parent BS,the base which the MS communicates with. For a multiple zone simulation, each zone sends back a list of BSsand their transmit powers in its domain which transmit on the same channel, as part of an EV_BS_LIST_RES event. The MS receives this event from each zone and combines them into one consolidated list of BSs transmitting on the channel. The MS sends this consolidated list of co-channel BS transmit powers to the parent zone,the zone in which the mobile is located, in an EVD-WNPWR_CRTL_REQ event. The parent zone sums up the interference contributions from the co-channel BSs and sends the total downlink interference power $P_I^{(d)}$ on the MS downlink channel c back to the BS. The BS receives EVDWNPWR_CRTL_RES and calculates the required BS transmitter power, $P_{new}^{(d)}(c)$, to meet the downlink target

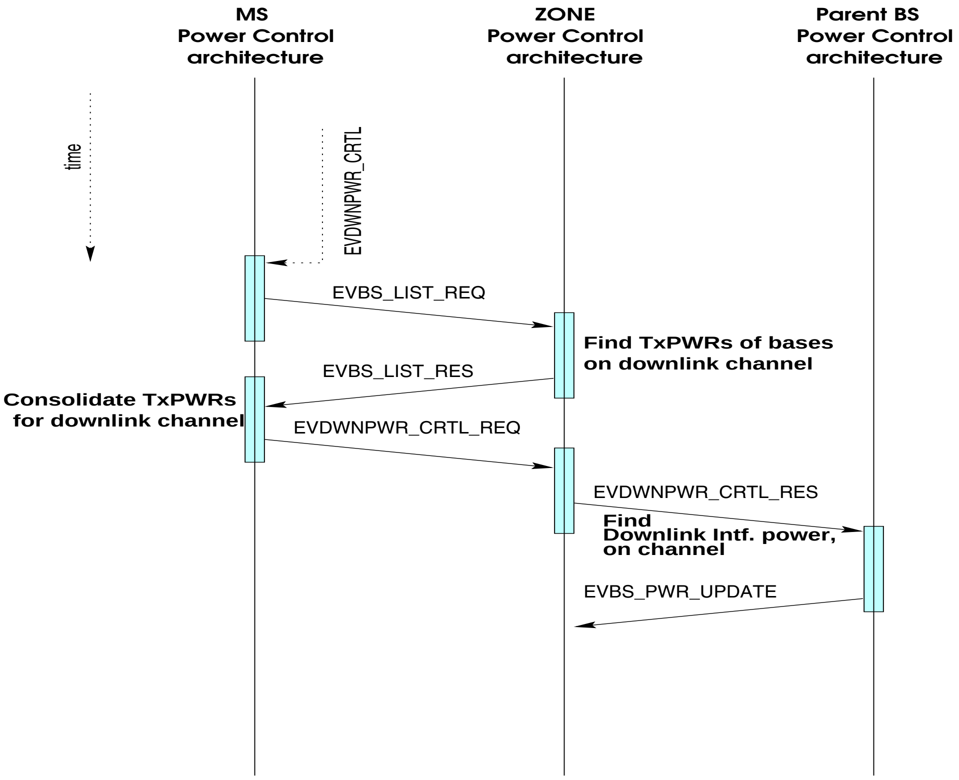

Figure 2.8: Event flow and entity interaction diagram of synchronous distributed uplink power control protocol. A vertical shaded strip represents a process of a power control architecture in an entity, which receives and/or sends Inter-EI events to the processes of its peer architectures or Intra-EI events to the processes of architecture within its entity. Solid-line arrows show the direction of flow of Inter-EI events from a power control architecture to its peer in other entity. Broken-line arrows show the direction of flow of Intra-EI events from a power control architecture to an architecture within an entity

SIR. The required BS transmitter power is calculated as:

$$P_{req}^{(d)}(c) = 10^{\frac{\gamma_{t(down)}}{10}} \cdot \frac{P_I^{(d)}}{gain^{(d)}}$$

The new required BS transmitter power is computed as:

$$P_{new}^{(d)}(c) = \max\left(\texttt{MIN\_BS\_PWR}, \min\left(\texttt{MAX\_BS\_PWR}, P_{req}^{(d)}(c)\right)\right)$$

Note that MIN_BS_PWR, the minimum BS transmitter power, is 0 Watt, MAX_BS_PWR, the maximum BS transmitter powe, is 1 Watt, and $P_{new}^{(d)}$ is the new required BS transmitter power on the downlink channel.

The new required base station transmit power is then sent to the parent zone in an EVBS_PWR_UPDATE event.

An important point to be noted is that on the uplink , the new mobile transmitter power calculated is sent back to the `MS` which updates its power immediately. The situation is different on the downlink since in IS-136, downlink power control can only be done at the carrier level and not at the slot level. Thus the required base transmitter powers computed every slot are just stored and used by the specified downlink power control algorithm at the beginning of the next frame to determine the carrier transmit power.

The sequence of events described above is required because unlike uplink interference calculation, computation of downlink interference cannot be partitioned among `zone`s. This is because for computing the downlink interference for a `MS` located at a particular point in the geography, we need the transmit powers of all co-channel `BS`s as well as the link gains from that point to each of them. The parent `zone` has the link gains but it knows only the transmit powers of co-channel bases in its domain. Hence, we have to consolidate the list of co-channel `BS`s from all `zone`s and send them back to the parent `zone` for computing the downlink interference.

2.5.7 Handoffs

In our simulations, handoffs can be triggered by three factors

1. When the RSSI of a neighbouring base becomes stronger than that of the current base, an inter-cell handoff is triggered.
2. When the quality of the active link drops below a threshold for a specifies amount of time, an intra-cell handoff is triggered.
3. When mobiles needs to be moved around to fill vacant slots in to a slot in carriers for slot repacking, an intra-cell handoff is triggered.

We can configure the simulation to include any or all of the above three types of handoffs. We assume all three types of handoffs can occur in the description of the handoff mechanism in WiPPET$_{\mathrm{IS-136}}$ which follows.

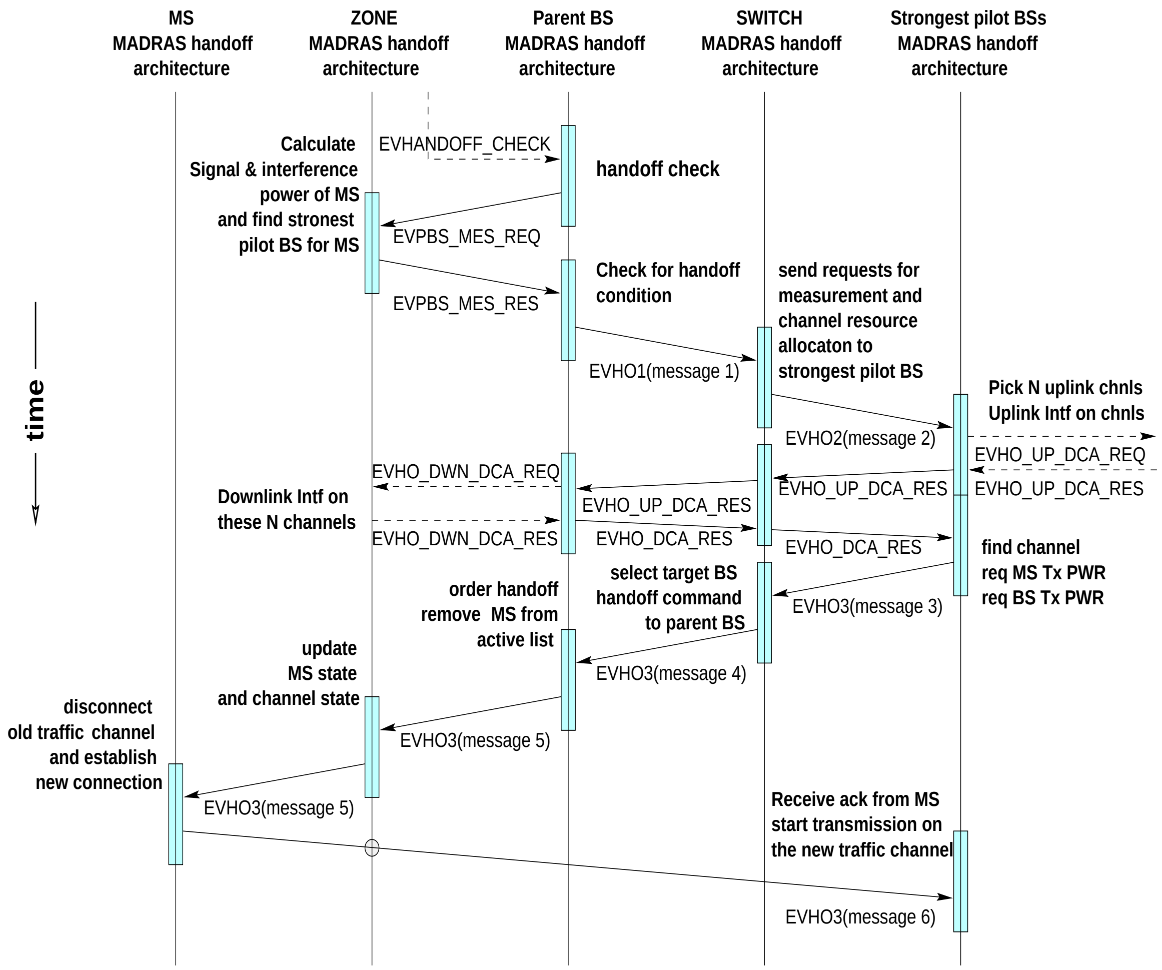

Figure 2.9: Event flow and entity interaction diagram of MADRAS handoff protocol. A vertical shaded strip represents a process of a handoff architecture in an entity, which receives and/or sends Inter-EI events to the processes of its peer architectures or Intra-EI events to the processes of architecture within its entity. Solid-line arrows show the direction of flow of Inter-EI events from a handoff architecture to its peer in other entity. Broken-line arrows show the direction of flow of Intra-EI events from a handoff architecture to an architecture within an entity. A shaded circle represents a process which acts as an event repeater in which the process receives the event and sends it to its destination without doing any event processing. Here, `zone` has a event repeater process, because, there are no direct TeD channels from any `MS`s to any `BS`s.

Handoff check: The first event in the IS-136 handoff is EVHANDOFF_CHECK. This is an Intra-EI event sent on an internal TED channel within the parent BS. This event is sent by every BS for every MS which is active in that particular slot. In the handoff check procedure, the parent BS checks the possibility of the given MS handing off in response to any of the three factors discussed earlier. The handoff check is carried out synchronously for all MSs in the system which are active in the current slot.

The parent BS initiates the handoff check procedure by sending EVPBS_MES_REQ to the handoff architecture of the parent zone requesting the uplink interference on the MS channel. In case of a multiple zone simulation, the parent BS sends the request to all zones. For simplicity, here, we consider a single zone simulation. The parent zone receives the request and calculates the uplink interference. It also selects a candidate BS the strongest pilot BS for the MS corresponding to the MS location. An involved discussion of the interference calculation in WiPPET is presented in [15].

If there is no stronger BS for the given MS other than parent BS, the parent zone assigns NOBS , implying no candidate base station, as the candidate BS. Now let us assume that the parent zone found a candidate BS. In this case, the parent zone compares the path loss $gain_{new}$ from the strongest BS to the MS and the path loss $gain_{old}$ from the parent BS to the MS. If

$$\texttt{RSSI\_HY(dB)} \leq 10 \log_{10} \left(\frac{gain_{new}}{gain_{old}}\right)$$

the parent zone assigns the strongest BS as the candidate BS for the MS. Otherwise, the parent zone assigns NOBS as the candidate BS. RSSI_HY is the hysteresis threshold and is a simulation parameter.

The parent zone sends back EVPBS_MES_RES to the handoff architecture of the parent BS. This event includes the uplink interference value , a list of candidate base stations , the current MS transmitter power and the path loss from the MS to its parent BS. In our handoff algorithms, there is only one candidate BS and that is the strongest pilot BS.The parent BS receives EVPBS_MES_RES. In a multiple zone simulation, the parent BS collects the uplink interference power values from all the zones and adds them up to get the total uplink interference.

The parent BS checks for a candidate BS by examining the ListbsID parameter sent back by the zone. A candidate BS implies that an inter-cell handoff can be initiated. The parent BS then sends EVHO1 with payload message type **1** to the handoff architecture of the switch requesting a handoff for the MS. Let us first consider the case when a candidate base is available.

Inter-Cell Handoff

Handoff request: The switch receives the handoff request EVHO1 from the parent BS. EVHO1 has payload message type **1** including the gain, MS transmitter power , total uplink interference on the MS channel, and the candidate BS ID (ListbsID). Then, the switch sends a request to the handoff architecture of the candidate BS as an event EVHO2 with payload message type **2** requesting a traffic channel for the MS.

The candidate BS receives EVHO2 and initiates the channel allocation algorithm by sending EVHO_UP_DCA_REQ to its channel allocation architecture. The channel allocation architecture sends back EVHO_UP_DCA_RES with the N best uplink channels $c[N]$, the uplink interference power $P_I^{(u)}[N]$ on each of these channels at the candidate BS and the uplink path loss value $\text{gain}_{new\_base}^{(u)}$ from the MS to the candidate BS. The channel allocation architecture of the candidate BS sends the event to the handoff architecture as an Intra-EI event. The handoff architecture of the candidate BS then passes this event back to the switch which passes it back to the parent BS. The parent BS then gets the downlink interference on each of these N channels. The sequence of events is exactly the same as described in section 2.5.2 with just the CAs in the event names replaced by HOs. Table 2.1 summarizes these corresponding events

The parent BS then sends EVHO_DCA_RES to the switch which passes it on to the candidate BS. The payload for this message contains the N best uplink channels $c[N]$, the downlink interference power $P_I^{(d)}[N]$ on each of these channels at the MS location and the downlink path loss value $\text{gain}_{new\_base}^{(d)}$ from the candidate BS to the MS. The uplink and downlink gains are assumed to be the same. The candidate BS now knows the uplink and downlink interference powers on these N channels as well as the new gain.

Call Admission	Handoffs
EVCA_UPCHINTF_REQ	EVHO_UPCHINTF_REQ
EVCA_UPCHINTF_RES	EVHO_UPCHINTF_RES
EV_MACA_LIST	EV_MAHO_LIST
EVCA_BS_LIST_REQ	EVHO_BS_LIST_REQ
EVCA_DWNCHINTF_REQ	EVHO_DWNCHINTF_REQ
EVCA_DWNCHINTF_RES	EVHO_DWNCHINTF_RES

Table 2.1: Corresponding events in call admission and handoffs

It then runs through the specified channel allocation algorithm described in section 2.5.2 and picks the best channel c_{new}.

The `switch` then sends EVHO3 to the handoff architecture of the parent `BS` with payload message type **4** including the candidate `BS` ID, the new channel c_{new}, the required `MS` transmitter power on the channel, $P^{(u)}_{new}$, and the required base station transmit power on the channel $P^{(d)}_{new}$.

Handoff order: The parent `BS` receives EVHO3 with payload message type **4** and checks the possibility of a handoff by comparing the new traffic channel ID (c_{new}) with `NO_CHANNEL`. If there is no traffic channel available at the candidate `BS`, the handoff algorithm fails.

At this point, the parent `BS` checks if the call should be dropped or not. If the handoff fails, the parent `BS` calculates the current uplink SIR, $SIR^{(u)}_{current}$, and the current downlink SIR, $SIR^{(d)}_{current}$, as follows

$$SIR^{(u)}_{current} = 10\log_{10}\left(\frac{gain^{(u)}.P_{ms}}{P^{(u)}_I}\right) \tag{2.1}$$

$$SIR^{(d)}_{current} = 10\log_{10}\left(\frac{gain^{(d)}.P_{carrier}}{P^{(d)}_I}\right) \tag{2.2}$$

where $gain^{(u)}$ and $gain^{(d)}$ are the uplink and downlink gains respectively and P_{ms} and $P_{carrier}$ are the current `MS` transmit power and carrier transmit powers respectively.

If $SIR^{(u)}_{current}$ or $SIR^{(d)}_{current} < \gamma_{drop}$ (drop threshold), the `MS` *outage time* is incremented. In case there is a new traffic channel available at the candidate `BS`, the handoff architecture at parent `BS` calculates the new uplink SIR and downlink SIR for the `MS` at the new BS as

$$SIR_{new}^{(u)} = 10\log_{10}(\frac{gain^{(u)}.P_{ms}}{P_I^{(u)}}) \tag{2.3}$$

$$SIR_{new}^{(d)} = 10\log_{10}(\frac{gain^{(d)}.P_{carrier}}{P_I^{(d)}}) \tag{2.4}$$

where the gains and transmit powers pertain to the channel c_{new} allocated by the new BS.

If $SIR_{new}^{(u)}$ or $SIR_{new}^{(d)} < \gamma_{drop}$, the outage time is updated. The parent BS has an outage timer for each active MS. The parent BS uses this timer to keep track of the MS outage time. Whenever the uplink or downlink SIR goes above the drop threshold (γ_{drop}), the parent BS resets the outage timer to zero. If the outage time $\geq$ MAX_OUTAGE_TIME, the maximum outage time allowable, the MS is considered to be dropped and the handoff architecture at the parent BS sends EVCALL_DROP to the call admission architecture. The drop threshold (γ_{drop}) and maximum outage time (MAX_OUTAGE_TIME) are simulation parameters. EVCALL_DROP is an Intra-EI event sent by the handoff architecture of the parent BS to its call admission architecture. By receiving this event, the call admission architecture removes the MS from the active list. The call admission architecture of the parent BS in turn sends EVCALL_DROP to the call admission architecture of the parent zone and the switch.

If the MS is not dropped, the parent BS orders the MS to disconnect the present connection and establish a new connection with the candidate BS on the new traffic channel by sending EVHO3 with payload message type **5**. Now, the candidate BS becomes the target BS. Since there is no direct TED channel from a BS to any MS, EVHO3 is routed through the parent zone.

Handoff over: The handoff architecture of the parent zone receives EVHO3 with message **5** and updates the MS states: the new channel and the new parent BS ID and the new uplink and downlink transmitter powers , $P_{new}^{(u)}$ and $P_{new}^{(d)}$. Then, the zone passes on the same event to the MS with same message ID. The handoff architecture of the MS receives EVHO3 with payload message type **5** and updates its states: the new traffic channel and the new parent BS ID and the new uplink transmitter power,

$P^{(u)}_{new}$. Then, the MS breaks the present connection to its parent BS and makes a new connection to the target BS. It informs the target BS that it is communicating on the new traffic channel using EVHO3 with message **6**. Since there is no direct TeD channel from the MS to the new BS, EVHO3 is routed through the parent zone to the target BS. The target BS receives EVHO3 with message **6** from the MS. It updates its channel state and puts the MS in its active list. The MS will start transmitting in its next active slot. The IS-136 inter-cell handoff algorithm ends here.

Intra-Cell Handoff (SIR based)

If there is no candidate BS an inter-cell handoff cannot be initiated. We then compute the current uplink and downlink SIR using equations 2.1 and 2.2. If the current uplink or downlink SIR , $SIR^{(u)}_{current}$ or $SIR^{(d)}_{current} < \gamma_{drop}$, the drop threshold SIR, then the outage timer is incremented. If the outage time $\geq$ MAX_INTRAHO_OUTAGE_TIME , the maximum outage time allowable before an intra-cell handoff is initiated, then the candidate BS becomes the parent BS and the sequence of events is the same as that for inter-cell handoffs. The only difference here is that we look for a channel within the same same BS.

Intra-Cell Handoff (Slot Repacking)

If the link quality is sufficient enough not to warrant an intra-cell handoff, we see if the MS could be moved to another carrier for more efficient slot packing. As discussed earlier under section 2.5.1, each BS maintains a list of source and target channels which is updated every slot. The MS scans the list to see if its assigned channel is a source channel and if so, picks up the corresponding target channel. This implies that the target channel is a vacant channel in a partially occupied carrier and the MS could be moved there for more efficient carrier packing. If the MS does not find its channel in the source channel list, slot repacking is not possible and hence an intra-cell handoff cannot be initiated for slot repacking. We shall consider the case when slot repacking is possible.

The BS sends EVSLOT_REARRANGE to the handoff architecture with the mobile ID, msID and the target channel c_{target}. The handoff architecture receives this event and sends EVSRUPCH_INTF_REQ with c_{target} as the payload to all zones to get the uplink interference power on c_{target}. The zone sends EVSRUPCH_INTF_RES including the path loss value $gain^{(u)}$ from the MS to its parent BS and the interference value $P_I^{(u)}$ on the uplink channel c_{target}. The BS then sends c_{target} in an EV_MASR_LIST event to the parent zone which in turn sends it to the MS. To determine the downlink interference on a particular channel, we need the transmit powers of all BSs on that channel. Since the base stations are distributed across all zones, we need to query all the zones to get the consolidated list of BSs transmitting on that channel. The MS sends c_{target} in an EVSRBS_LIST_REQ event to all zones and receives in turn a list of transmit powers of BSs transmitting on c_{target} from each zone in an EVSRBS_LIST_RES event. The MS consolidates the transmit powers into a single list and sends them to the parent zone as part of an EVSR_DWNCH_INTF_REQ event. The parent zone calculates the downlink interference on c_{target} and sends the interference measurements back to the BS as part of an EVSR_DWNCH_INTF_RES event. The BS thus knows the uplink and downlink interference on c_{target} as well as the link gains. The BS determines the required mobile and base station transmit powers, $P_{req}^{(u)}$ and $P_{req}^{(d)}$, to achieve γ_t(target SIR) using equations 1.2 and 1.3 respectively. If $P_{req}^{(u)}$ and $P_{req}^{(d)}$ are less than their respective peak power constraints, MAX_MS_PWR and MAX_BS_PWR, the channel is available for an intra-cell handoff. The BS then sends these parameters to the switch in an EVHO3 message.

Then, the switch sends EVHO3 to the handoff architecture of the parent BS with payload message type **4** including the candidate BS ID, $P_{req}^{(u)}$, $P_{req}^{(d)}$, c_{new}, and $gain_{new}$. The event sequence from this point is exactly the same as for inter-cell handoffs from the **Handoff Order** phase.

Chapter 3

Simulation Results

3.1 Assumptions

In this section we outline some of the key assumptions of our simulation model.

1. Time slots are synchronized in the system on both the uplink and downlink. On the uplink this implies that all mobiles assigned to a particular time slot transmit simultaneously and on the downlink all base stations are synchronized too.
2. A mobile without a call in progress can make signal strength measurements on forward RF channels. This is a slight variation of the procedure specified in the standards whereby signal strength measurements by the mobile before call setup are restricted to forward control channels.
3. We obtain fast measurements within the duration of a full rate time slot.

3.2 Simulation Parameters

In this section we describe the simulation experiments which were carried out using our simulator WiPPET$_{\mathrm{IS-136}}$ and the results obtained. The simulation parameters used in these experiments are listed in Table 3.2.

parameter [units]	**2D Macrocell**
Number of BSs	48
Number of Carriers	13
Number of Slots/Carrier	3
Frame Length [ms]	39
Slot Length [ms]	13
Geography Style	Manhattan
Geographical Area	$12km \times 12km$
Cell Radius [m]	1000
Antenna Height [m]	50
Distance Resolution [m]	5
Shadow Fading Standard Deviation [dB]	6
Shadow Fading Correlation Distance [m]	50
Noise Power [dBW]	-150
Maximum Mobile Transmit Power [W]	1
Maximum Base transmit Power [W]	1
Mobile Speed [m/sec]	25
Probability of Going Straight	0.6
Time Between Handoff Checks [msec]	39
Time Between Power Updates [msec]	39
BS Assignment Algorithm	Strongest Pilot BS
Call Admission Channel-Allocation Algorithms	LI-DCA,Power Differential
Uplink Power Control Algorithm	Synchronous,distributed, SIR balancing
Downlink Power Control Algorithms	MAX,MIN,AVG
Handoff Algorithms	Inter-Cell, Intra-Cell
Handoff Channel-Allocation Algorithms	LI-DCA,Power Differential
RSSI hysteresis (RSSI_HY) [dB]	9
SIR hysteresis (SIR_HY) [dB]	3
Admission SIR (Γ_{new}) [dB]	20
Target SIR (Γ_t) [dB]	20
Handoff SIR (Γ_{ho}) [dB]	17
Drop SIR (Γ_{drop}) [dB]	14
Max Outage Time for Dropping [sec]	5
Outage Time for Intra-Cell Handoffs [sec]	1

Table 3.1: WiPPET$_{\text{IS}-136}$ simulation parameters

3.3 Performance Measures

3.3.1 Call Blocking Probability

A call is said to be blocked if it cannot be admitted into the system either because a channel is unavailable or if it cannot meet call admission thresholds on an available channel. With LI-DCA, a call is usually blocked for the latter reason. The call blocking probability is the fraction of calls blocked. It is defined as

$$P_b = \frac{\text{\# of new calls blocked}}{\text{\# of calls attempted}} \tag{3.1}$$

3.3.2 Call Dropping Probability

A call is dropped if the SIR on its assigned channel in the forward or reverse direction drops below a threshold called the outage margin for a specified amount of time called the outage time. The call dropping probability is the fraction of calls dropped. It is defined as

$$P_d = \frac{\text{\# of active calls dropped}}{\text{\# of calls admitted}} \tag{3.2}$$

3.3.3 Average Number of Handoffs Per Call

The average number of handoffs per call is computed as

$$N_{ho} = \frac{\text{Total \# handoffs}}{\text{Total \# of calls admitted}} \tag{3.3}$$

This is computed separately for inter-cell and intra-cell handoffs.

3.3.4 Partial Slot Occupancy %

Two kinds of slots contribute to interference on the downlink in IS-136, the actually occupied slots and the partial slots which are idle but in which the base has to transmit anyway because of the downlink constraint. The partial slot occupancy percentage gives us the fraction of interference contributing slots which are actually contributing unnecessary interference as compared to the PERSLOT case. It is defined as

$$\%_{partial\_slots} = \frac{\sum_{allframes} N_{partial\_slots}}{\sum_{allframes} N_{partial\_slots} + \sum_{allframes} N_{occupied\_slots}} \tag{3.4}$$

where

$N_{partial\_slots}$ = Total # of partial slots in the system at a particular sampling instant

$N_{occupied\_slots}$ = Total # of occupied slots in the system at a particular sampling instant

3.4 Simulation Results

3.4.1 Handoffs

An important result which is evident from the graphs is that a system with inter-cell as well as intra-cell handoffs performs better than a system with only inter-cell handoffs as seen in Fig 3.1 This is because we could have a scenario in which a mobile in the middle of the cell could experience a long outage because a co-channel interferer in an adjacent cell is transmitting at high power. With inter-cell handoffs, the outage could last long enough for the mobile to be dropped before it reaches a cell boundary. With intra-cell handoffs however, the mobile could move to a better channel in the same cell.

To substantiate this conclusion, the SIR profile of a particular mobile in the system is plotted in Fig 3.2. Call dropping was disabled during this experiment and only inter-cell handoffs were allowed. The asterisks(*) denote the points at which inter-cell handoffs take place and the circles(o) denote the points where the call should have been dropped because of poor link quality. We observe from the profile that there are two points between the second and third inter-cell handoffs when the SIRs on the uplink and downlink combine to constitute an outage which would cause the call to be dropped because the mobile is still some way from the cell boundary. With intra-cell handoffs, we could avert this scenario by initiating a handoff after part of the outage time has elapsed.

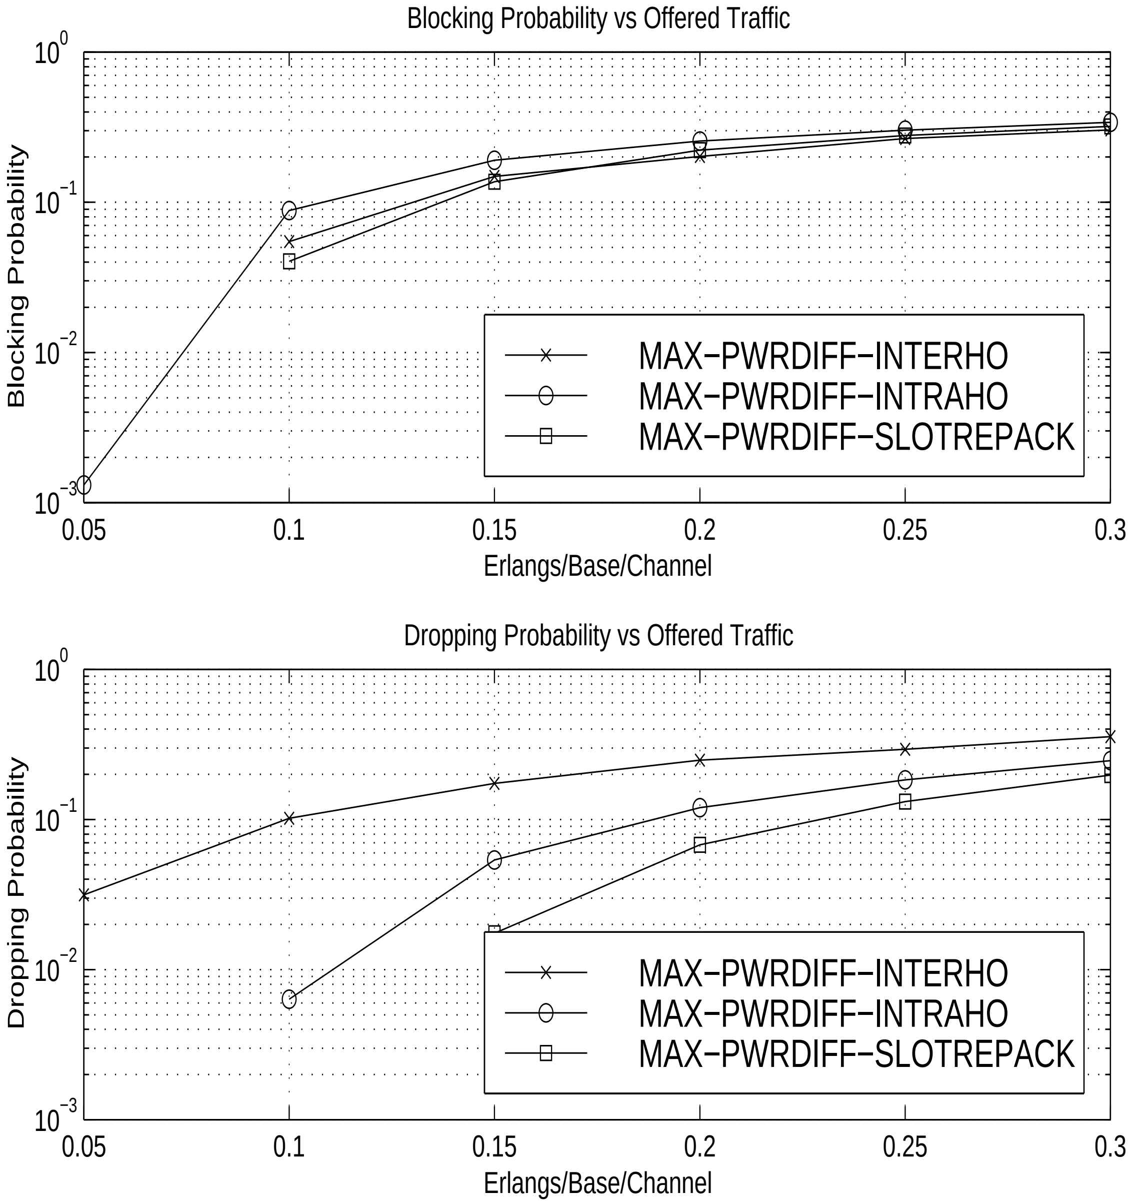

Figure 3.1: Comparison of blocking and dropping probabilities vs offered traffic between inter-cell handoffs and intra-cell handoffs using the MAX scheme for downlink power control and the power differential scheme for channel allocation

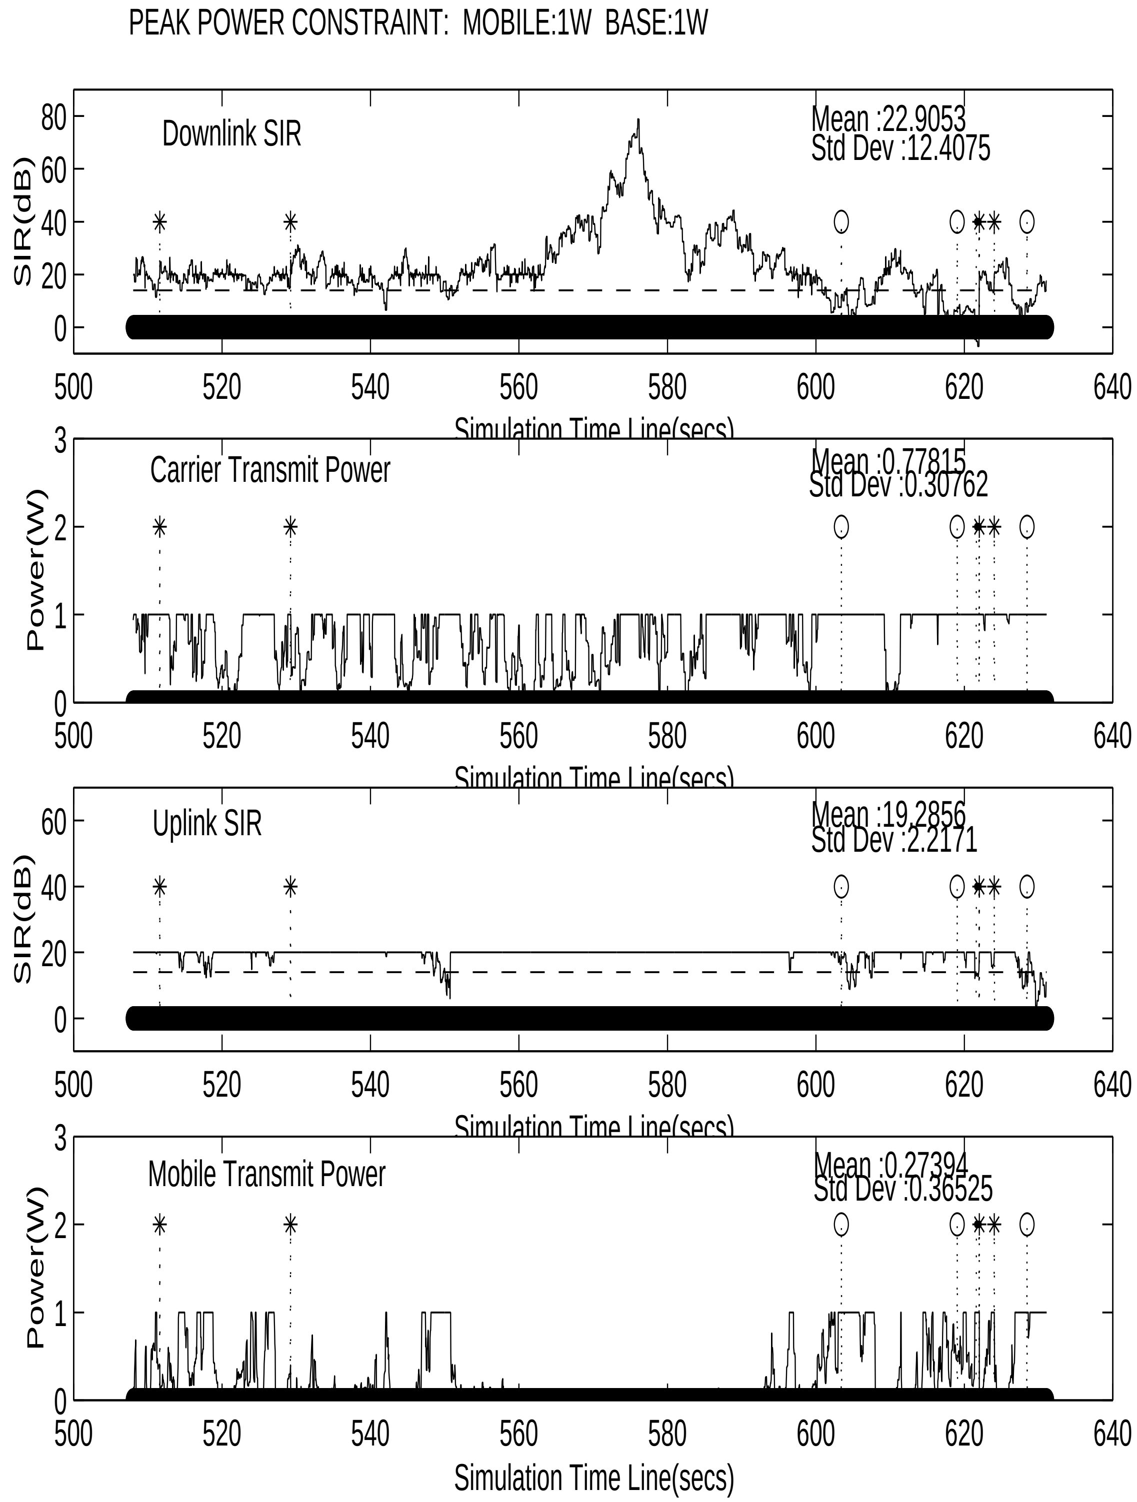

Figure 3.2: SIR trace during the call duration of a particular mobile in a simulation with just inter-cell handoffs. The (*) indicate the inter-cell handoffs and the (o) indicate points at which the call should have been dropped.

From Fig 3.1 we also observe a further improvement with slot repacking. With slot repacking, mobiles also make intra-cell handoffs whenever possible to occupy vacant slots in carriers and this reduces the number of partial slots and hence the additional downlink interference as seen in Fig 3.8. We observe that slot repacking results in the least partial slot occupancy percentage.

In terms of the average number of handoffs per call, we observe from Fig 3.3 that in general the MIN schemes have more number of intra-cell handoffs per call than the MAX schemes. In the MIN scheme, when three users are assigned to the same carrier, only one of them meets the target SIR while the other two are are below the target SIR and try to make intra-cell handoffs. The MIN scheme with slot repacking results in an unusually large number of intra-cell handoffs. This is because MIN and slot repacking both contribute to intra-cell handoffs.

3.4.2 Channel Allocation

Channel allocation algorithms come into play both during call admission and handoffs. The two algorithms investigated are Least Interference DCA(LI-DCA) and the Power Differential algorithm. We observe from Fig 3.7 and Fig 3.6 that LI-DCA yields the worst performance when we have partial slots. This is because LI-DCA is not designed to pack the carriers. It just assigns the best uplink channel which meets admission thresholds on the uplink and the downlink. Since uplink interference is the criterion, a slot in an idle carrier could be assigned when a partial slot is available if the uplink interference on the partial slot is worse. It thus optimizes the uplink but not the downlink which is the constraining factor in IS-136. The ideal case(PERSLOT) with LI-DCA however yields the lowest blocking and dropping rates as expected because of the absence of partial slots. We observe that performance improves with the power differential algorithm because it gives preference to partial slots. It tries to pack the carriers as far as possible in a way which introduces the least additional downlink interference. Fig 3.8 is an index of how efficient these algorithms are in reducing the number of partial slots. We observe that LI-DCA has the highest partial slot occupancy percentage. The

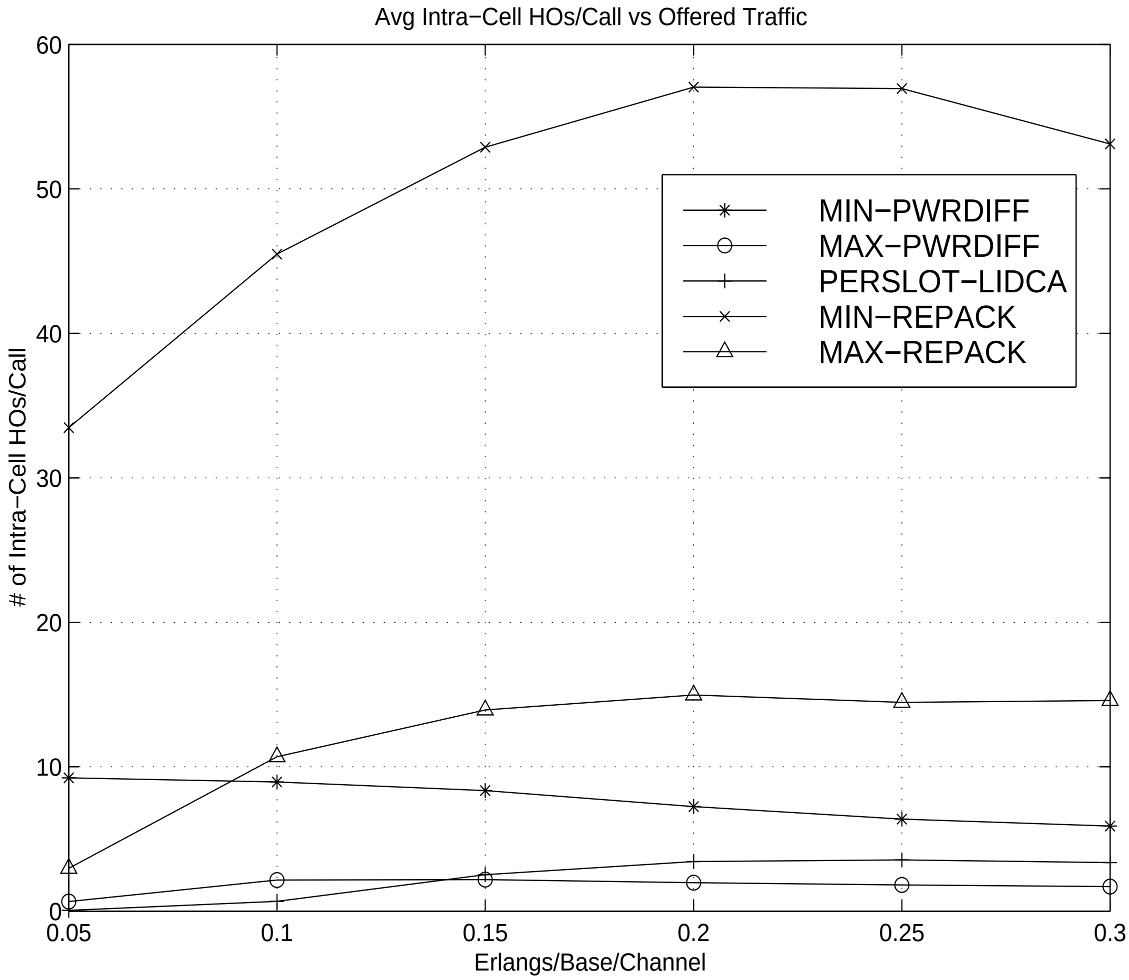

Figure 3.3: Average number of intra-cell handoffs/call as a function of offered traffic for the MAX and MIN schemes

power differential algorithm reduces the partial slot occupancy percentage considerably because it tries to pack the carriers.

3.4.3 Power Control

We observe from the blocking and dropping curves, that the closest we can get to the PERSLOT ideal case is with the MIN scheme for downlink power control coupled with slot repacking. The average carrier transmit powers as well as the average downlink SIRs are plotted in Fig 3.4 for the MAX and MIN schemes. We observe from the graphs that the advantage of using MIN is that the carrier transmit powers are relatively less than with the MAX scheme and consequently the co-channel interference is less. The

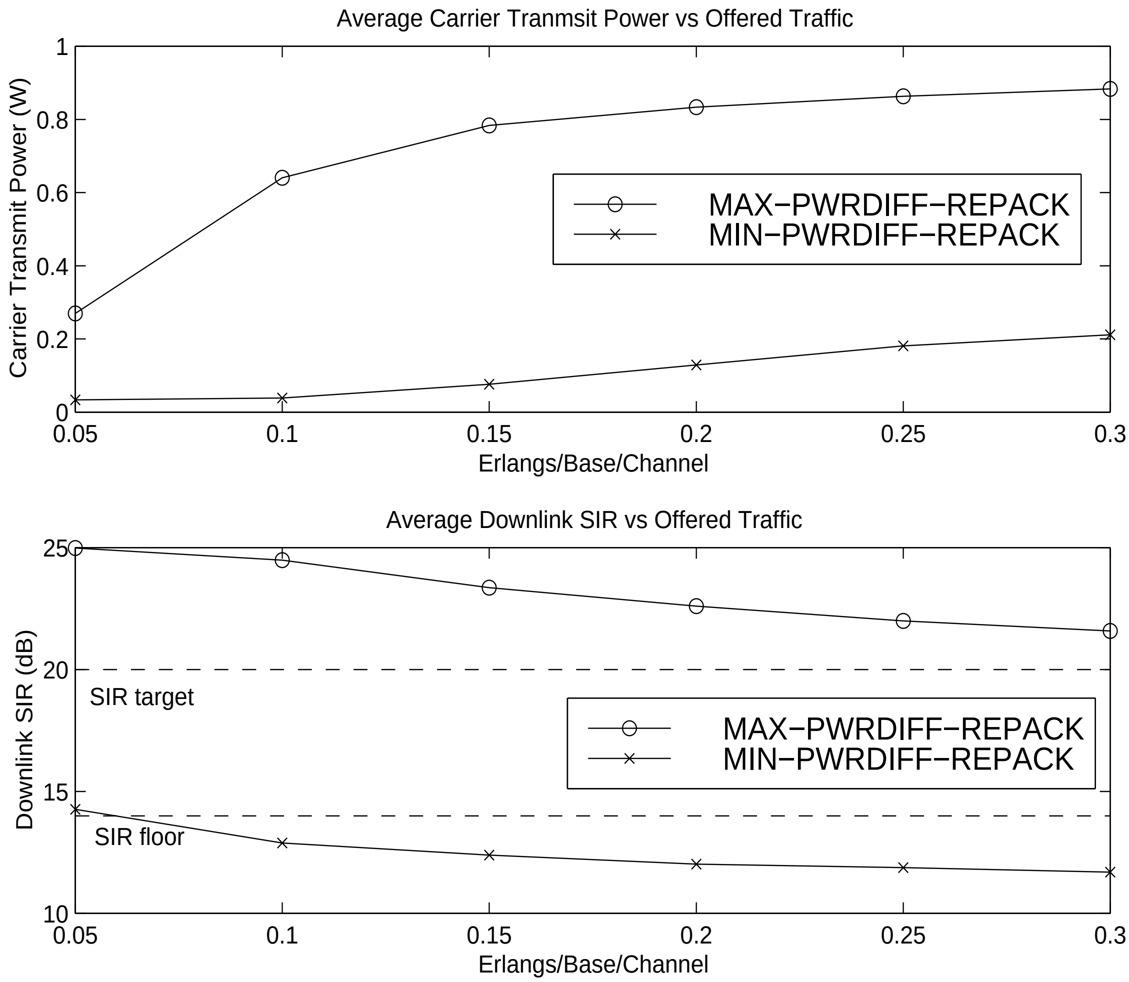

Figure 3.4: Average Carrier Transmit Power and Average downlink SIR vs offered traffic for the MAX and MIN schemes

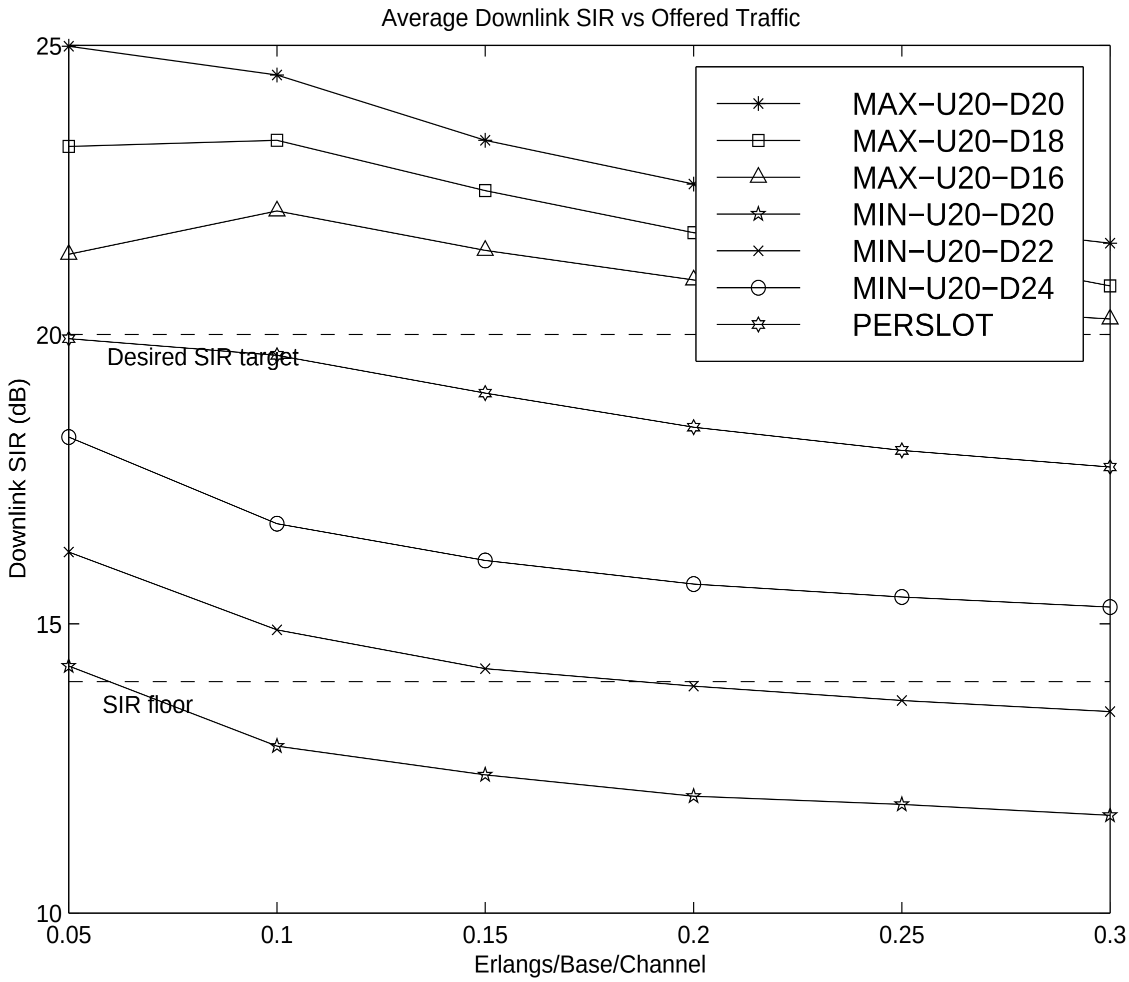

Figure 3.5: Average downlink SIR vs offered traffic for the MAX and MIN schemes for varying downlink target SIRs. It is observed that the target SIR should be set below the actual desired target for MAX and above the actual desired target for MIN

other aspect which was discussed earlier is that only one mobile in a carrier satisfies the downlink target SIR while the other two are below the target SIR looking to handoff. This is reflected in Fig 3.4 where we observe that the average downlink SIR is much less than the target SIR. With this scenario, we would expect more calls to be dropped since the dropping criterion is based on the link quality. Interestingly enough, the performance enhancement with MIN and slot repacking is achieved at the expense of an unusually larger number of intra-cell handoffs per call as seen in Fig 3.3 because mobiles keep hopping from carrier to carrier to prevent getting dropped. The performance enhancement seems to be strictly in terms of capacity if we can tolerate a poor link quality and an unusually large number of handoffs per call.

Although the average downlink SIR with MAX is more than the target SIR, we observe that the carrier transmit powers are higher and consequently the co-channel interference is larger. We observe from the above results that in both the schemes, the downlink SIR and co-channel interference are tied together. Ideally, we would like to maximize the SIR and minimize the interference, but with carrier level power control, we could opt for just one of the above and pay a penalty in terms of the other parameter.

A point to be noted with the MAX scheme is that even with perfect packing, we still pay a penalty for keeping the base station transmit power constant over a frame. This is because if we have more than one mobile on the same carrier, one of them would require maximum downlink transmit power which is used as the carrier transmit power. The transmit power in the other occupied slots will be higher than what is actually required by them and this translates into additional interference contributed by these slots when compared to PERSLOT.

The results discussed above were obtained with a target SIR of 20dB on the uplink and downlink. From Fig 3.4, we observe that for the MAX scheme, the average system SIR on the downlink is considerably higher than the target SIR. This is because only one mobile in a carrier operates at the target SIR while the others are above the target. Keeping the uplink target SIR fixed at 20dB, we observe the effect of decreasing the downlink target SIR in Fig 3.5. For the MAX scheme, the average system SIR

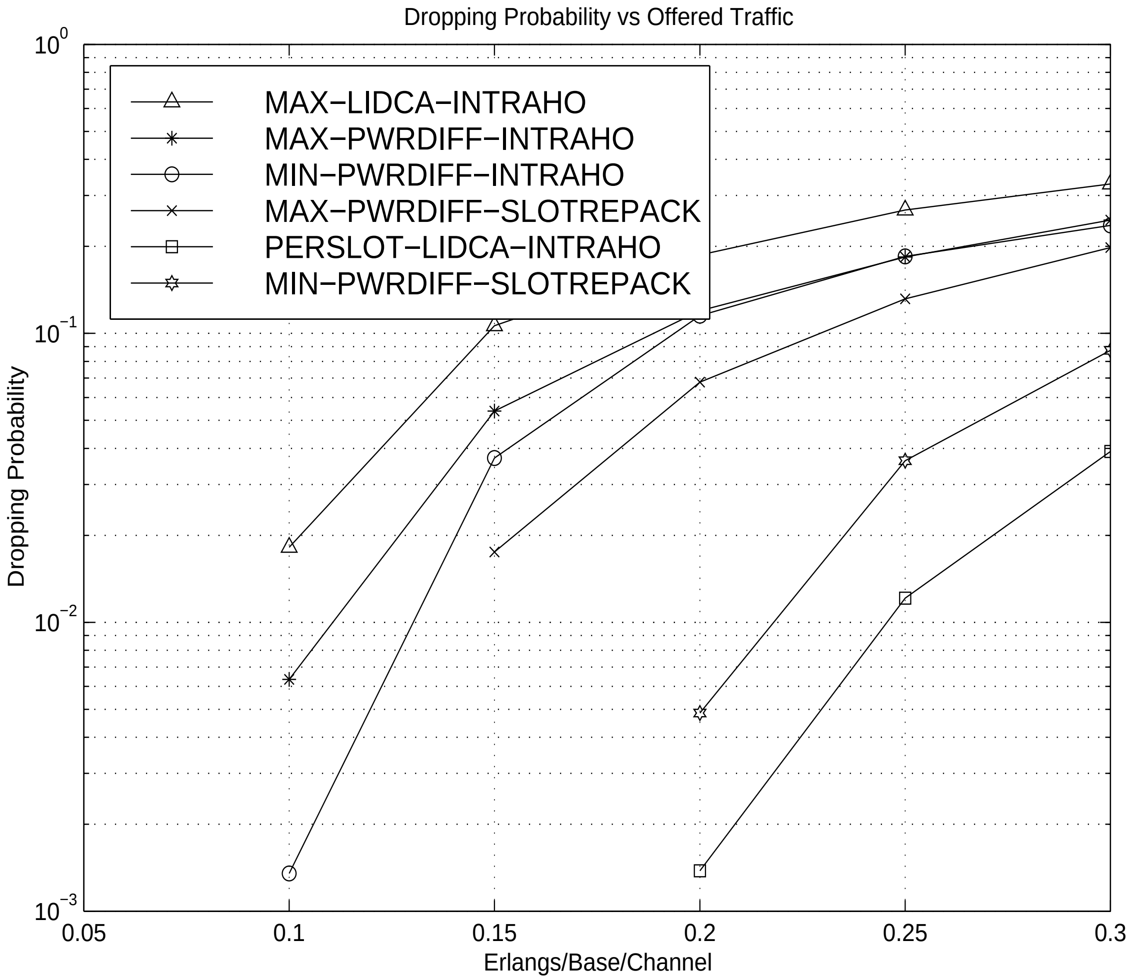

Figure 3.6: Dropping probability vs Offered Traffic for various combinations of power control, channel allocation and handoff algorithms investigated

on the downlink is always a couple of dB greater than the target SIR, consequently increasing the carrier transmit powers and hence the co-channel interference. Blocking and dropping rates also improve as seen in Fig 3.10 and Fig 3.9 . We conclude that we can do better with the MAX scheme by setting the downlink target SIR a couple of dB lower than what we actually want to achieve.

With the MIN scheme, the exact opposite happens. The average system SIR on the downlink is much below the target SIR as seen in Fig 3.4. For a target SIR of 20dB, it even falls below the SIR drop threshold. In this case, we increase the target SIR to pump up the average downlink system SIR as seen in Fig 3.5

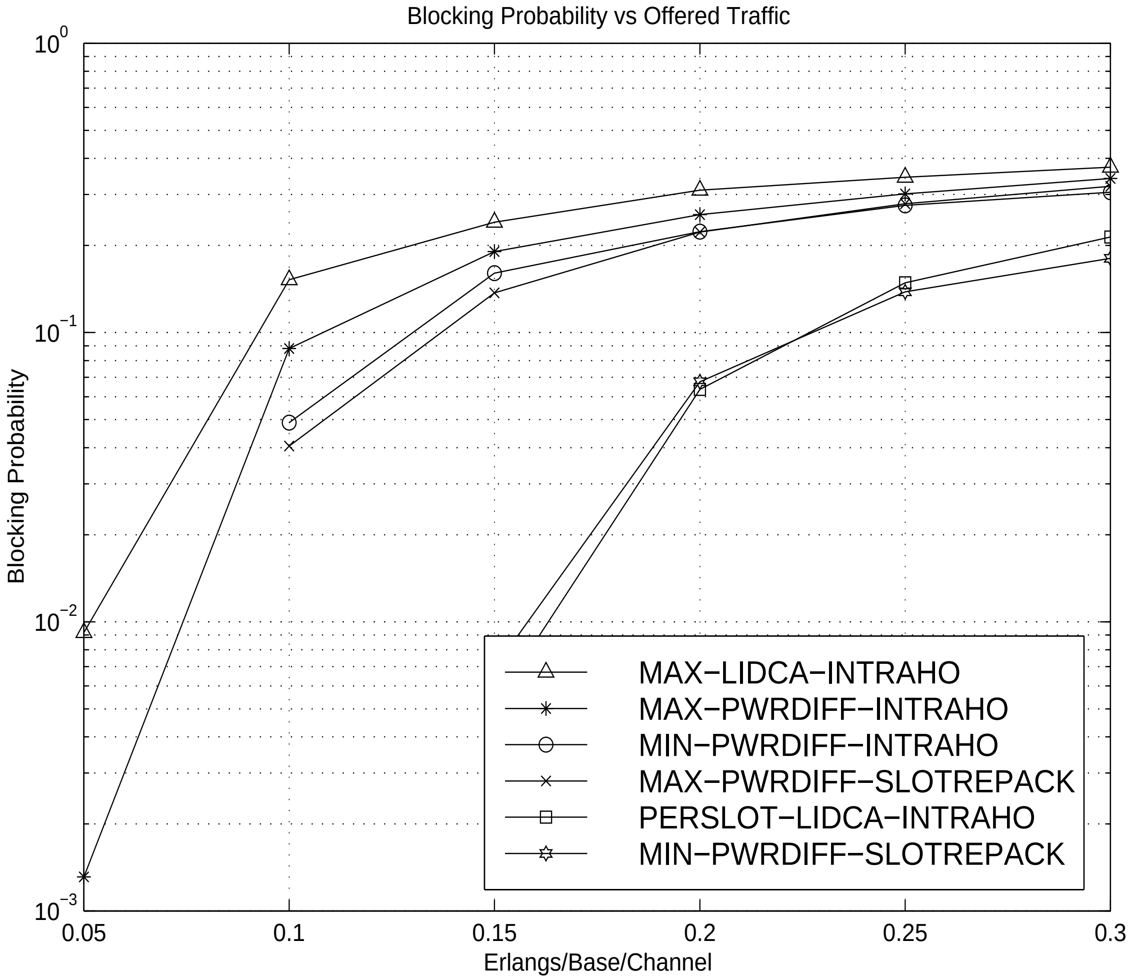

Figure 3.7: Blocking probability vs Offered Traffic for various combinations of power control, channel allocation and handoff algorithms investigated

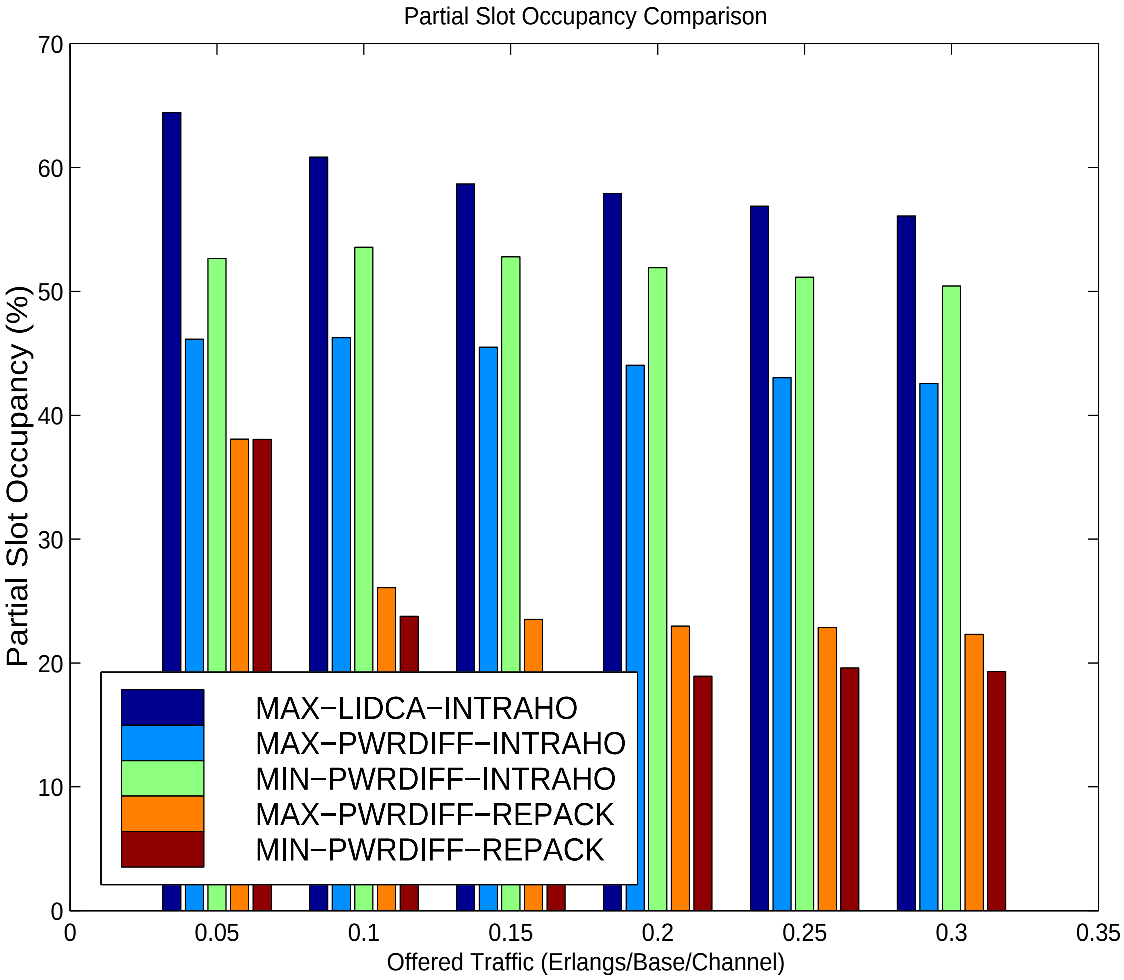

Figure 3.8: Average partial slot occupancy statistics versus offered traffic for various combinations of power control, channel allocation and handoff algorithms investigated

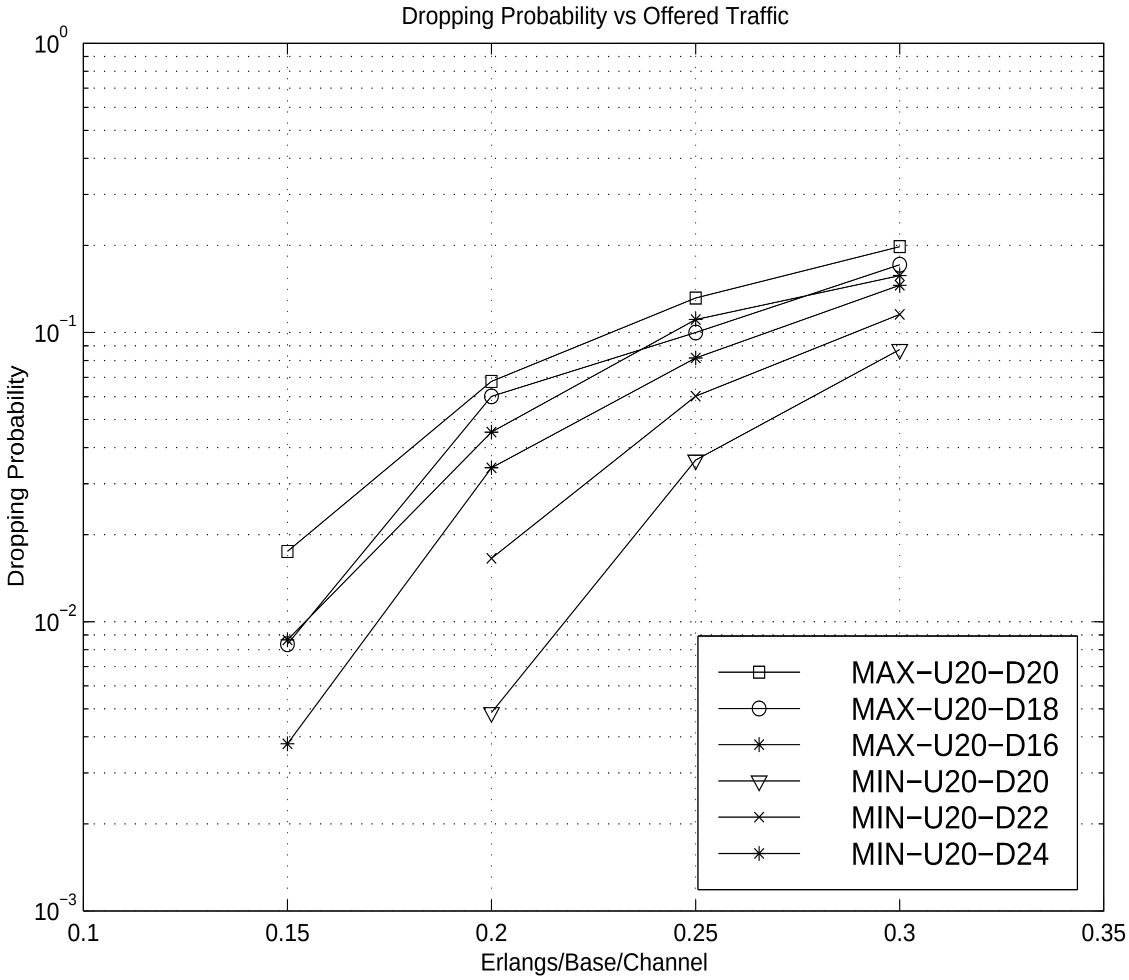

Figure 3.9: Dropping probability vs offered traffic for the MAX and MIN schemes for varying downlink target SIRs.

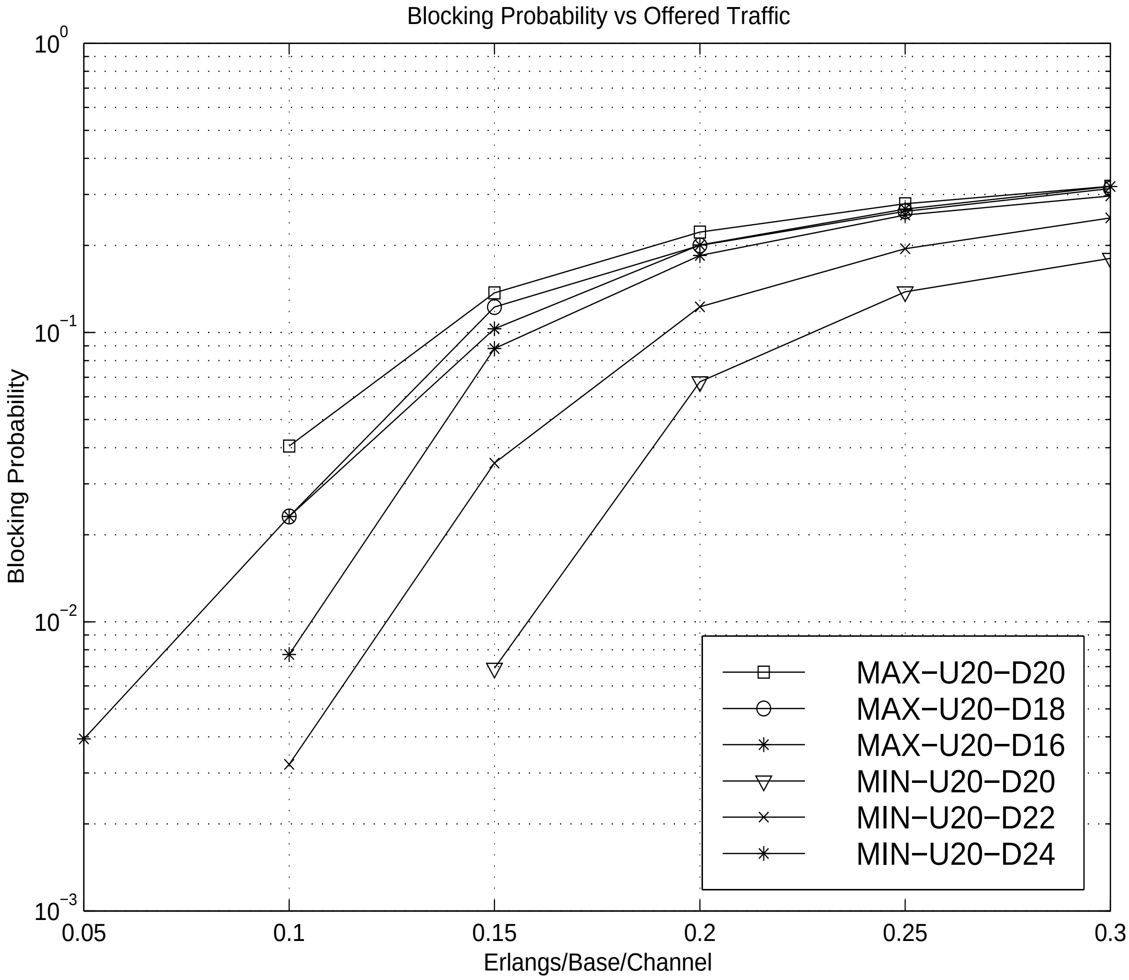

Figure 3.10: Blocking probability vs offered traffic for the MAX and MIN schemes for varying downlink target SIRs.

3.5 SIR Estimation

So far, one of the assumptions in our simulations is that we can obtain fast and accurate measurements of the SIR. We use these measurements in our power control, handoff and channel allocation algorithms. In practice, however, we cannot make perfect measurements of the SIR and hence have to estimate the SIR as accurately as possible. Algorithms for SIR estimation in TDMA systems have been proposed in [19, 20]. These algorithms vary in their computational complexity and the mean square estimation error, an index of how close the estimate is to the true SIR. We focus on the signal projection method proposed in [19] because it includes a model for an IS-54/IS-136 system. In this section, our objective is to develop a simple model whereby we can observe the effect of SIR estimation errors in our simulation.

3.5.1 Signal strength measurements in IS-136

In IS-136, the mobile tunes to and makes signal strength measurements on forward RF channels specified by the base. In SIR estimation, the averaging interval is a vital factor. It has been shown that the mean square estimation error is inversely proportional to the averaging interval [19, 20]. We get a better estimate by averaging over more samples. When signal strength measurements are done on forward traffic channels other than the current traffic channel, the standards specify that the minimum number of samples per channel is computed as

$$S_{min} = \frac{50}{N} \tag{3.5}$$

where N is the number of channels in the channel list sent by the base to the mobile.

Each sample is taken during an idle half-rate slot(6.67ms). Thus the averaging interval is given by

$$T_{avg} = 6.67 * S_{min} = \frac{333.50}{N}\text{ms} \tag{3.6}$$

parameter [units]	**Value**
Call Arrival Rate [mobiles/sec]	3.2
Uplink target SIR [dB]	20
Downlink target SIR [dB]	20
Uplink error std dev [dB]	0.01
Downlink error std dev(Current Channel) [dB]	0.5
Downlink error std dev(Other Channels) [dB]	0.89

Table 3.2: SIR estimation simulation parameters

3.5.2 System Model

Since our simulation model does not employ a signal model, we merely simulate SIR estimation errors by adding an error to the actual SIR. We assume that the estimation errors are uncorrelated samples from a lognormal distribution with zero mean and a standard deviation which depends on the averaging interval. The number of channels in the channel list sent to the mobile by the base is taken to be 3 which works out to an averaging interval of approximately 100ms from equations 3.5 and 3.6. For this averaging interval, we look up the corresponding error variance from the results obtained for an IS-136 system in [19]. For measurements on the current traffic channel, the averaging interval is taken to be 300ms. In IS-136, the number of signal strength samples taken on the current channel is more than that on other channels. Consequently, the error variance on the current channel is smaller than that on other channels and we would expect to get a better estimate on the current channel.

3.5.3 Simulation Parameters

3.5.4 Simulation Results

The distributions of the actual and estimated downlink SIRs for the MAX scheme are plotted in Figs 3.12 and 3.11 considering perfect as well as imperfect measurements for error deviations of 2.5dB and 0.5dB respectively. The measurements are made for all mobiles in the system during their assigned time slot. Perfect measurements imply that we use the actual measured SIR in our radio resource management algorithms.

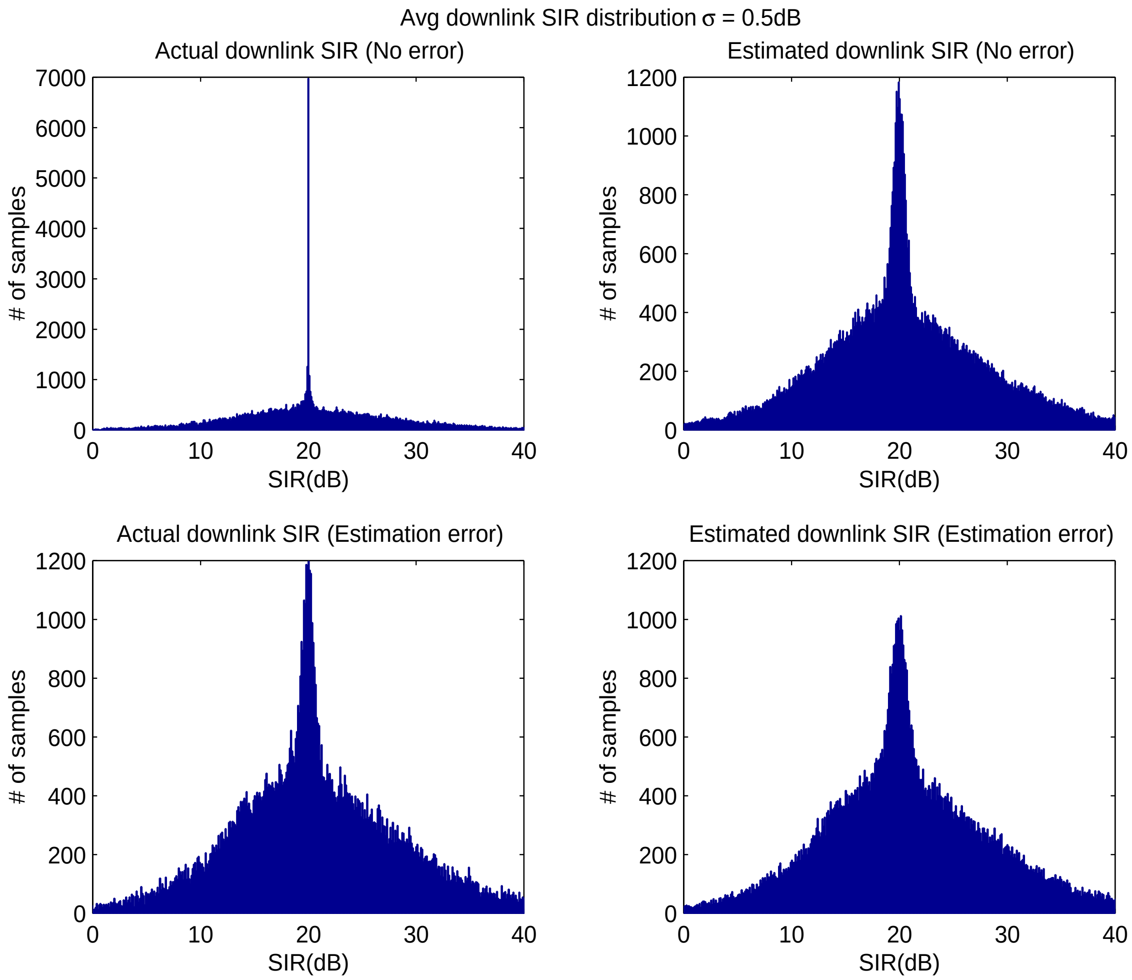

Figure 3.11: Distribution of the downlink SIR for perfect and imperfect measurements of the MAX scheme, $\sigma = 0.5$dB

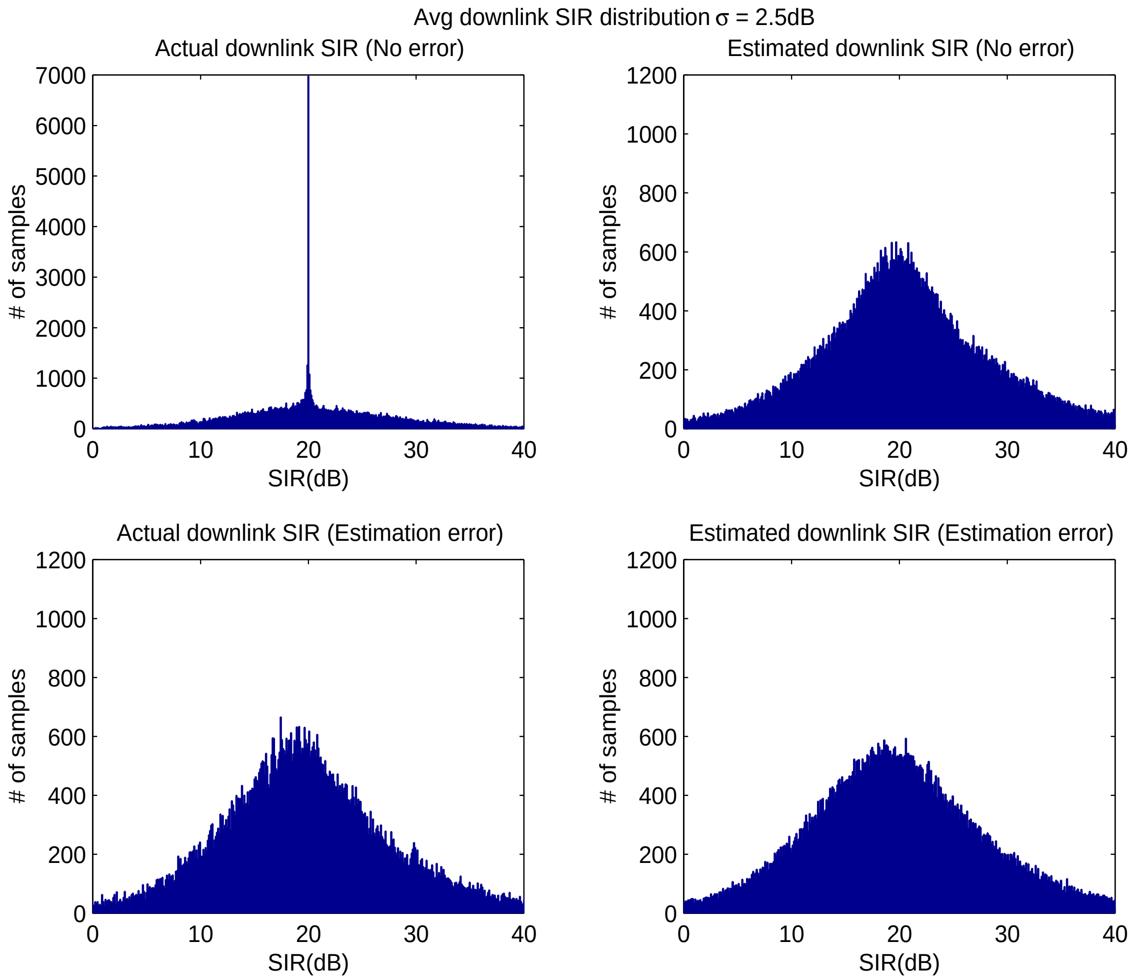

Figure 3.12: Distribution of the downlink SIR for perfect and imperfect measurements of the MAX scheme, $\sigma = 2.5$dB

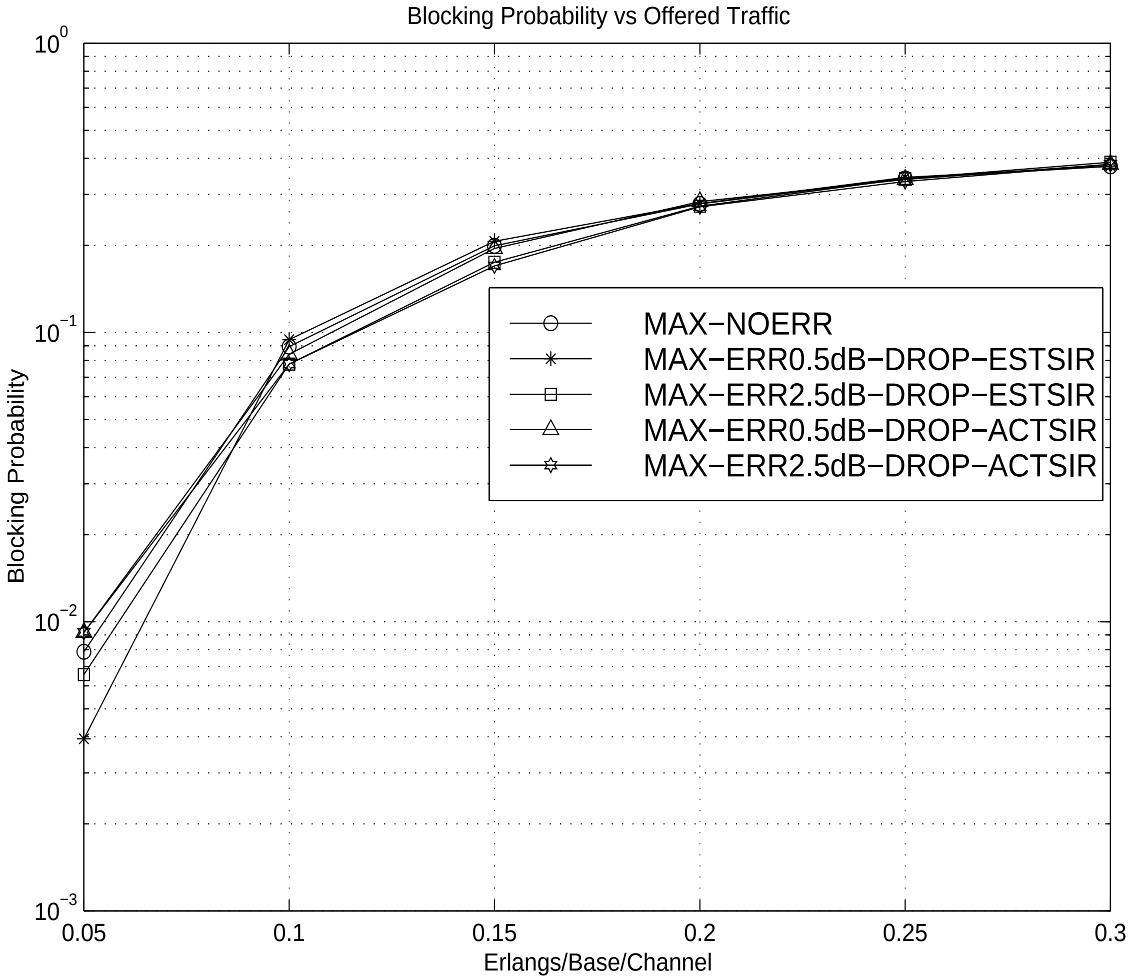

Figure 3.13: Blocking probability vs offered traffic with SIR estimation errors. The error deviation is given in dB. 'EST' and 'ACT' indicate if the dropping criterion is based on the estimated or the actual SIR.

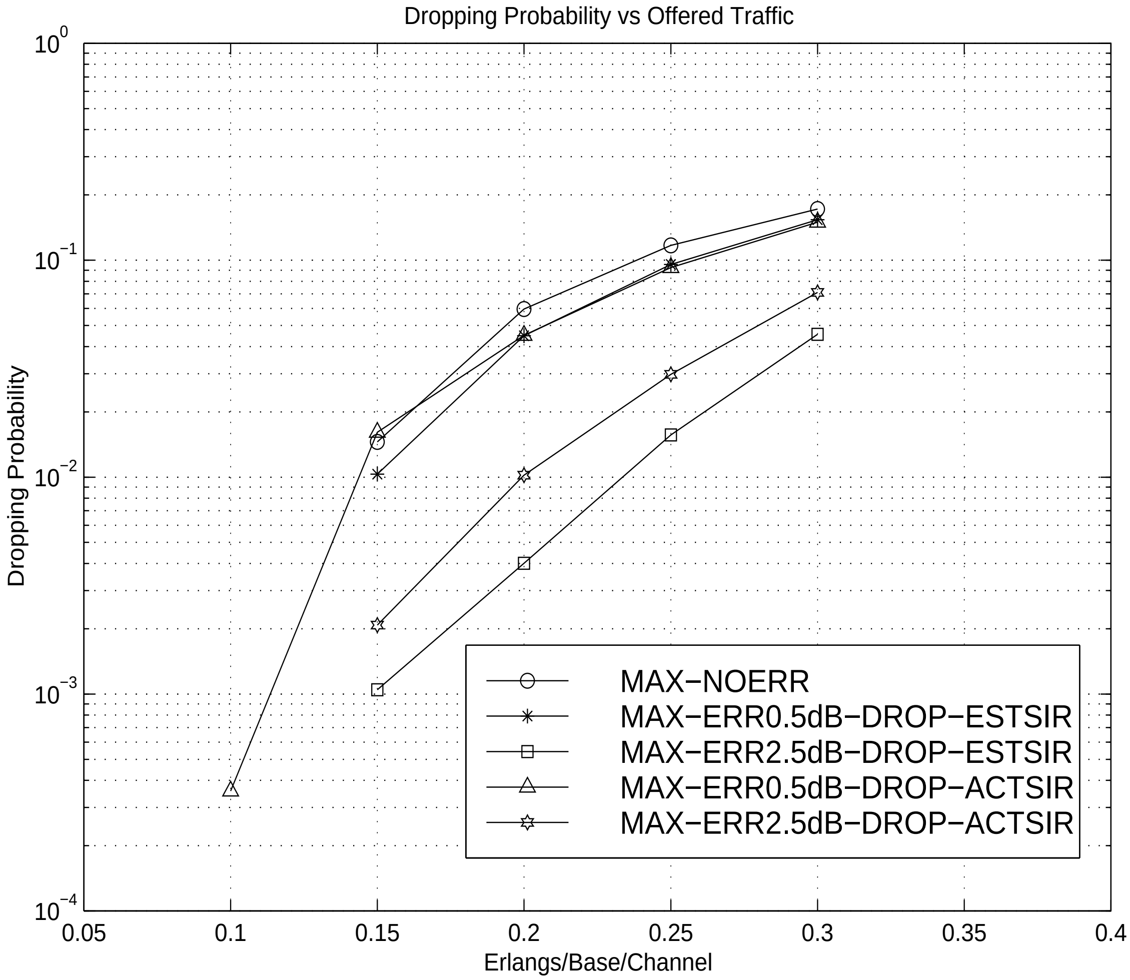

Figure 3.14: Dropping probability vs offered traffic with SIR estimation errors. The error deviation is given in dB. 'EST' and 'ACT' indicate if the dropping criterion is based on the estimated or the actual SIR.

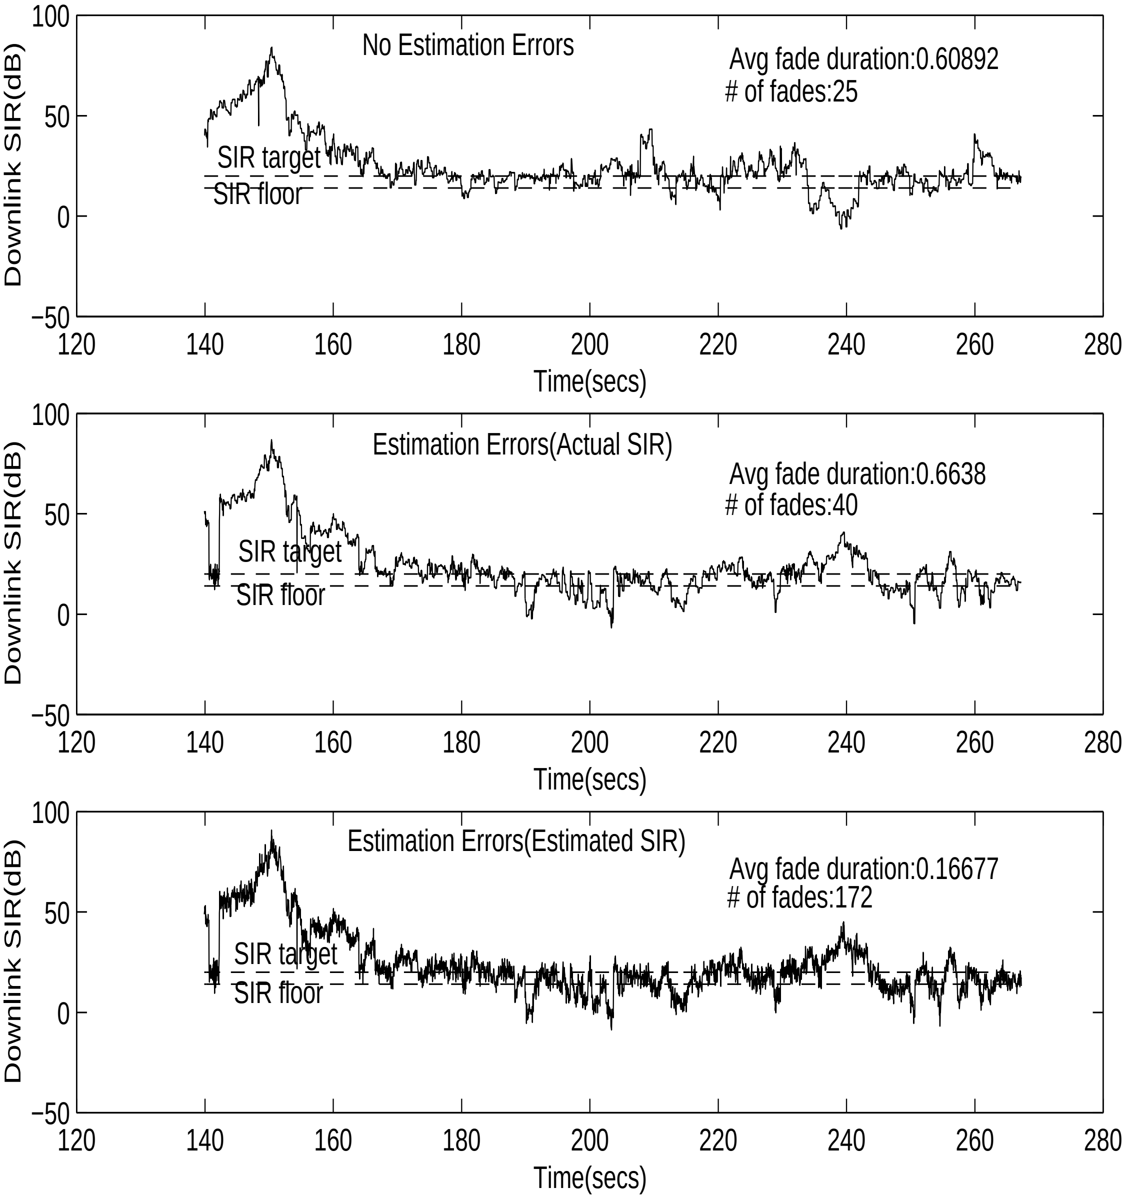

Figure 3.15: Actual and estimated downlink SIR profiles of a mobile with perfect and imperfect measurements

With imperfect measurements, we add a lognormal error to the actual measured SIR and use this estimate in our algorithms. However we track the actual and the estimated SIR in both cases.

We observe that in both cases, the distribution of the actual downlink SIR is concentrated around the target SIR in the absence of SIR estimation errors. As expected, in the presence of estimation errors, the distributions are more dispersed about the target SIR. Comparing the distributions for deviations of 0.5dB and 2.5dB in Figs 3.11 and 3.12 respectively, we observe that the distribution with the smaller deviation has a pronounced peak around the target SIR. This is because with a smaller error variance, we would expect more samples to be clustered around the target SIR. For imperfect measurements, the distributions of the actual and estimated downlink SIRs do not differ significantly. This result is not fully understood.

The blocking and dropping curves with SIR estimation errors are plotted in Fig 3.13 and 3.14. The curves plotted with the blocking and dropping criteria are based on the estimated as well as the actual SIR. It is interesting to note that the system does better with SIR estimation errors. Dropping rates decrease significantly as we increase the error variance. This could be a result of our assumption that the estimation errors are uncorrelated. Consider a scenario when the actual SIR is a dB below the drop threshold and the outage time counter is close to the maximum outage time, implying that the call is on the verge of being dropped. A positive deviation of around 2.0dB at this point would increase the estimated SIR above the drop threshold, resetting the outage time counter and hence prolonging the call.

We experiment by basing the dropping criterion on the actual SIR but that does not change things significantly. We plot the downlink SIR trace of a single mobile with perfect measurements and the actual and estimated SIR traces with imperfect measurements in Fig 3.15. A drop in the SIR below the drop threshold(γ_{floor}) constitutes a fade. We observe from the SIR profile with no estimation errors that there is one long fade which lasts long enough for the call to have been dropped and a few small fades. With estimation errors, the actual SIR profile reveals more fades but we observe that the fades are much shorter than the maximum outage time. A fade is often followed by a positive excursion which resets the outage time counter as discussed earlier. The estimated SIR profile has a greater number of fades and SIR fluctuations and the fade durations get even smaller.

The conclusion seems to be that the uncorrelated lognormal estimation error system model results in the SIRs fluctuating randomly from frame to frame and this leads to lower dropping rates because of the dropping policy we have adopted.

Chapter 4

Summary

The principal motivation behind the study was to evaluate the penalty in employing carrier level power control on the downlink in IS-136 and also to explore various combinations of channel allocation, power control and handoff algorithms to try mitigate this effect.

A brief overview of IS-136 was presented, followed by a description of the key problem of not being able to vary the base station transmit power on a slot by slot basis on the IS-136 downlink. We then discussed the following radio resource management algorithms

- Power Control: The MAX and MIN schemes on the downlink
- Channel Allocation: The conventional Least Interference DCA scheme(LI-DCA) and the Power Differential scheme which incorporates carrier packing and tries to reduce the number of partial slots.
- Handoffs: The RSSI based inter-cell handoff scheme and intra-cell handoffs which occur both due to poor link quality as well as for slot repacking. The slot repacking algorithm was also disussed in this context.

The implementation details of the simulator, WiPPET$_{\mathrm{IS-136}}$ which is written in TeD and is built up on the existing WiPPET [15] were then discussed. Event flow and interaction diagrams for the various protocols it supports were also presented.

Finally, we analyzed the simulation results we obtained in terms of blocking and dropping curves, SIR profiles, average transmit powers, slot occupancy statistics and average number of handoffs per call for various combinations of the above radio resource management algorithms. Some of the conclusions which followed are given below:

- Dynamic channel allocation requires a system with intra-cell handoffs. Specifically, when we perform intra-cell handoffs due to poor link quality, blocking and dropping are reduced. Performance improves further when we also make intra-cell handoffs for slot repacking.
- We pay a considerable penalty for carrier level power control on the downlink as opposed to varying it on a slot by slot basis and turning off during idle slots.
- Slot packing improves performance by reducing the number of partial slots and consequently the additional downlink interference.
- MIN outperforms MAX but at the expense of a large number of intra-cell handoffs per call and poor link quality. The performance improvement is solely in terms of capacity.
- The target SIR needs to be a few dB below the desired target SIR for MAX and a few dB above the target SIR for MIN.
- Independent SIR estimation errors introduce random SIR fluctuations which reduce dropping
- Channel allocation, power control and handoff algorithms are tightly coupled

References

[1] TIA. *TIA/EIA Interim Standard IS-136.2*, Dec 1994.

[2] E. Sourour. Time Slot Assignment Techniques for TDMA Digital Cellular Systems. *IEEE Transactions on Vehicular Technology*, 43(1):121–127, 1994.

[3] Y. Akaiwa and H. Andoh. Channel segregation - a self-organized dynamic channel allocation method:application to TDMA/FDMA microcellular system. *IEEE Journal on Selected Areas in Communications*, 11(6):949–954, 1993.

[4] M. Almgren, H. Andersson, and K. Walltstedt. Capacity enhancements in a TDMA system. In *Proc. IEEE Vehicular Technology Conference, VTC-93*, pages 277–280, 1993.

[5] J. Zander. Distributed cochannel control in cellular radio systems. *IEEE Transactions on Vehicular Technology*, 41(3):305–311, August 1992.

[6] G. J. Foschini and Z. Miljanic. A simple distributed autonomous power control algorithm and its convergence. *IEEE Transactions on Vehicular Technology*, 42(4), 1993.

[7] W. Tschirks. Effects of transmission power control on the cochannel interference in cellular radio networks. *Electrotechnik und Informationstechnik*, 106(5), 1989.

[8] T. Fuji and M. Sakamato. Reduction of co-channel interference in cellular systems by intra-zone channel reassignment and adaptive transmitter power control. In *Proc. IEEE Vehicular Technology Conference, VTC-88*, pages 668–672, 1988.

[9] J. Zander. Performance of optimum transmitter power control in cellular radio systems. *IEEE Transactions on Vehicular Technology*, 41:57–62, February 1992.

[10] S. Grandhi, R. Vijayan, D. J. Goodman, and J. Zander. Centralized power control for cellular radio systems. *IEEE Transactions on Vehicular Technology*, 42(4):466–468, November 1993.

[11] S. Grandhi, J. Zander, and R. Yates. Constrained power control. *International Journal of Wireless Personal Communications*, 1(4), 1995.

[12] R. W. Nettleton. A high capacity assignment method for cellular mobile telephone systems. In *Proc. IEEE Vehicular Technology Conference, VTC-89*, pages 359–367, 1989.

[13] L. J. Cimini, G. J. Foschini, C. Lin, and Z. Miljanic. Call blocking performance of distributed algorithms for dca in microcells. *IEEE Transactions on Vehicular Technology*, 42(8):2600–2607, 1994.

[14] D. J. Goodman, S. Grandhi, and R. Vijayan. Distributed dynamic channel assignments schemes. In *Proc. IEEE Vehicular Technology Conference, VTC-93*, pages 532–535, 1993.

[15] J. Panchal. Parallel simulator of wireless networks. Masters thesis, Rutgers, The State University of New Jersey, New Brunswick, NJ, Jul 1998.

[16] K. Perumalla, A. Ogielski, and R. Fujimoto. MetaTeD: A meta language for modeling telecommunication networks. TR GIT-CC-96-32, College of Computing, Georgia Institute of Technology, 1996.

[17] K. Perumalla and R. Fujimoto. A C++ instance of TeD. TR GIT-CC-96-33, College of Computing, Georgia Institute of Technology, 1996.

[18] TIA. *TIA/EIA Interim Standard IS-136.1*, Dec 1994.

[19] M. Turkboylari and G. L. Stuber. An efficient algorithm for estimating the signal-to-interfernce ratio in TDMA cellular systems. *IEEE Transactions on Communications*, 46(6):728–731, 1998.

[20] M. Andersin, N. Mandayam, and R. Yates. Subspace based estimation of the signal to interference ratio for TDMA cellular systems. WINLAB-TR 115, Wireless Information Networks Laboratory, Rutgers University, February 1996.