

Project	IEEE 802.16 Broadband Wireless Access Working Group < http://ieee802.org/16 >
Title	Performance Analysis of a Distributed Approach to Interference Mitigation in LE 802.16 Systems
Date Submitted	2005-07-11
Source(s)	Seán Murphy, Omar Ashagi, Antonio Ruzzelli, Liam Murphy, John Murphy University College Dublin Belfield, Dublin 4, Ireland. Voice: +353-86-8578462 Fax: +353-86-58578462 mailto:sean.murphy@iname.com
Re:	Call for Contributions, IEEE 802.16h Task Group on License Exempt Co-existence, IEEE 802.16h-05/014, Item 4 – Other Topics Relevant for the PAR.
Abstract	This document describes a distributed approach to interference mitigation for 802.16 LE systems. The scheme applies to 802.16 systems operating on a single channel; each BS in the system remains silent for some period of time to facilitate access to the medium by other BSs in the area. The performance of this distributed approach is compared with the performance that can be obtained if full-knowledge of the system is known. The results show that the full system knowledge in conjunction with intelligent scheduling can result in quite good overall system performance; the distributed approach, however, does not achieve such good levels of performance. More work is required to improve the performance of this distributed approach.
Purpose	Information.
Notice	This document has been prepared to assist IEEE 802.16. It is offered as a basis for discussion and is not binding on the contributing individual(s) or organization(s). The material in this document is subject to change in form and content after further study. The contributor(s) reserve(s) the right to add, amend or withdraw material contained herein.
Release	The contributor grants a free, irrevocable license to the IEEE to incorporate material contained in this contribution, and any modifications thereof, in the creation of an IEEE Standards publication; to copyright in the IEEE's name any IEEE Standards publication even though it may include portions of this contribution; and at the IEEE's sole discretion to permit others to reproduce in whole or in part the resulting IEEE Standards publication. The contributor also acknowledges and accepts that this contribution may be made public by IEEE 802.16.
Patent Policy and Procedures	The contributor is familiar with the IEEE 802.16 Patent Policy and Procedures < http://ieee802.org/16/ipr/patents/policy.html >, including the statement "IEEE standards may include the known use of patent(s), including patent applications, provided the IEEE receives assurance from the patent holder or applicant with respect to patents essential for compliance with both mandatory and optional portions of the standard." Early disclosure to the Working Group of patent information that might be relevant to the standard is essential to reduce the possibility for delays in the development process and increase the likelihood that the draft publication will be approved for publication. Please notify the Chair < mailto:chair@wirelessman.org > as early as possible, in written or electronic form, if patented technology (or technology under patent application) might be incorporated into a draft standard being developed within the IEEE 802.16 Working Group. The Chair will disclose this notification via the IEEE 802.16 web site < http://ieee802.org/16/ipr/patents/notices >.

Performance Analysis of a Distributed Approach to Interference Mitigation in LE 802.16 Systems

Seán Murphy, Omar Ashagi, Antonio Ruzzelli, Liam Murphy, John Murphy
Dept of Computer Science, University College Dublin, Belfield, Dublin 4, Ireland

1. Introduction

IEEE 802.16 has been designed such that it can operate in license-exempt spectrum. The IEEE 802.16 system consists of a *Base Station* (BS) and one or more *Subscriber Stations* (SSs) distributed over a geographical area with a radius of typically up to a few kilometres. In the case in which there are a limited number of channels available, interference between different IEEE 802.16 systems can arise. In this paper we study the performance of a number of IEEE 802.16 systems operating in license-exempt mode of operation. More specifically, we wish to investigate the performance of a *distributed* approach which can be used to mitigate the impact of such interference. This approach is based on introducing sleep intervals for each BSs. Initially, these sleep intervals are created randomly.

A rudimentary Java simulator was implemented to determine the system performance. The simulator was designed to simulate a number of 802.16 systems operating in license-exempt mode in the same geographical area on the same channel. Furthermore, the interference mitigation approach described below is simulated. The simulator can determine the amount of interference between these systems and system performance metrics such as the throughput per SS.

Two basic models were then developed in order to determine how well the distributed system works. In these models, full knowledge of the system was assumed and using this knowledge, a schedule for the system was which results in small levels of interference. The first of the two basic models was a very conservative one in which no interference was permitted at any node in the system; the second was one in which interference in the system was controlled. Two variants of the latter model were developed – one in which the emphasis is on throughput and one in which the emphasis is on fairness. The results of this modeling exercise were then compared with the performance obtained by the distributed system.

The remainder of the paper is organized as follows: in section 2 related work is discussed. Section 3 gives a brief description of the distributed approach to interference mitigation which is the focus of this work. Section 4 describes the models that were developed to compare against the performance of the distributed scheme. Section 5 describes some results which were obtained. First, the simulator which was developed is described; this is followed by a discussion of some results obtained from simulating the distributed scheme. Following this is a comparison of results obtained from simulating the distributed scheme and the results predicted by the models developed. Finally, conclusions and future work is presented in section 6.

2. Related Work

Hoymann et al were the first to investigate the performance of IEEE 802.16a and ETSI HiperMAN in [4,5]. Their results showed that the MAC functions introduced an overhead of approximately 10%. Also, they have shown how the optional 802.16 packing and fragmentation features can optimise the throughput especially in the case where several active connections are sharing the MAC frame. In [6] the authors looked at the performance issues in 802.16 using different modulation and coding schemes and looked at their impact on delay and throughput.

The authors in [7] focused on the performance of the 802.16 PHY and their results showed that the performance of early systems will fall considerably short of that which many had envisaged. They proposed a number of standards-compliant approaches to improving the overall system throughput, including use of *Multiple-Input Multiple-Output* (MIMO) antenna systems and implementing an interference cancellation mechanism along with *Hybrid Automatic Repeat Request* (HARQ). Interestingly, the authors point out that interference cancellation for users who are located on the edge of a cell is a bigger problem than in conventional cellular systems such as GSM due to low mobility where users are likely to be there for ever. In GSM concept the problem has been recently has been addressed in [8,9]. They have suggested a possible solution by developing a low-complexity interference-cancelling receiver for the SS.

The authors in [3] investigated interference issues in IEEE802.16 mesh networks. They proposed an interference-aware framework comprising of routing and scheduling schemes to deliver good throughput and high spectral efficiency in the mesh system.

Interference issues in different radio systems have been studied in a number of contexts. In [1], the authors address the interference problem between two different *Wireless Local Area Network* (WLAN) technologies - IEEE 802.11a and ETSI HiperLAN/2. They proposed a central coordination device which is capable of operating in and controlling both systems, but requires changes to both standards to facilitate use of the medium by both systems. In [2] the authors looked at the impact of interference on Bluetooth when it is operating in close proximity to a WLAN system. The authors used a probability analysis approach to measure the packet error for Bluetooth. Their results showed the significant impact of interference on Bluetooth performance. Furthermore, they have showed that the longer Bluetooth packets are more prone to interference than the shorter packets.

None of the contributions to the above areas addresses the problem under study here. The studies on performance of 802.16 systems discussed above consider the licensed case. The WLAN related interference work is not applicable in an 802.16 context as there are very significant differences between WLAN and WMAN systems (system architecture, radio interfaces, range issues, etc)

3. A Distributed Approach to Interference Mitigation

The objective of this scheme is to have a simple distributed approach which can result in moderately efficient use of the available resources. More specifically, it does not require explicit co-ordination between the different systems operating in close proximity to each other but yet delivers reasonable (and stable) use of the medium.

In the proposed scheme, each BS must remain silent for some periods of time in order to facilitate access to the medium by other users. A random process controlled the periods of time during which a BS was active or inactive. Each BS had an associated *activity factor*, $\alpha \in [0,1]$. For each BS, then, the probability that it was active for any individual frame was α . Indeed, α also reflected the fraction of time that the BS was active over the long run.

Round robin scheduling was used to map the SSs to the BS frames under the assumption that there were a fixed number of SSs scheduled in each frame and each SS had the same amount of time scheduled for uplink/downlink transmission.

The above scheme operates such that a number of BSs can transmit simultaneously and interference can occur. However, the idea behind the scheme is that with suitably chosen activity factors, the amount of interference is not so large and a simple distributed algorithm can result in reasonably efficient used of medium.

A number of important assumptions were made in the design of this system. These were as follows:

- *that there is no explicit co-ordination between the different BSs*
- *there was synchronisation between the BSs: all of the frames have the same duration and commence at exactly the same time. While this is a constraining assumption, it was thought reasonable for a first exploration of these ideas – further, it may be possible to realize this using GPS-based synchronization, for example;*
- *that there was always data ready to send and receive at the SS: hence the focus of the work is on the saturation performance of the system;*
- *the time was divided equally between the uplink and the downlink: hence the system can support the same amount of traffic in the uplink as the downlink;*
- *a rudimentary power control scheme was used: the power required for uplink and downlink transmission was determined using the Friis free space propagation model in conjunction with the antenna gains*
- *it was assumed that all of the SSs communicate with the BS using 16-QAM*

A simulator was developed to determine how the system performs and this was used to generate some of the results which are presented below. More details on the simulator are provided in section 5 below.

4. Performance Modelling of Interference Mitigation Schemes

Two approaches to resource allocation and scheduling of a number of 802.16 BSs operating in LE mode are presented. In both cases, full knowledge of the system is assumed a priori – including, in particular, all BS and SS locations – and hence they differ quite fundamentally from the distributed approach described above. However, they can be used to determine an upper bound on the performance that can be attained by the distributed system.

Both approaches can be used to determine a way of scheduling DL and UL transmissions to/from the SSs. In the first approach a schedule is determined in which there is no interference between any transmissions in the system. This is the *interference-free* (IF) approach. In the second approach, there is a limited amount of interference, but it only occurs at nodes which are not affected by the transmitted data, i.e. interference only occurs at nodes which would not benefit from receiving either of the interfering transmissions. This is the *controlled-interference* approach. Indeed, two variants of the controlled interference approach are described here – one which is optimised to maximise throughput and one which is optimised to maximise fairness. These are termed the CI-T and CI-F schemes respectively.

In order to describe the two approaches, it is first useful to make the following definitions:

- K = Maximum number of SSs that can be allocated in one frame;
- N = Number of different providers in an area network; From now on, we will refer to different provider as different base stations BSs in the area network ;
- $BS_1 \dots BS_N$ = Base stations in an area network;
- $SS_1 \dots SS_N$ = Number of SS of each base station (e.g. SS_3 is the number of subscribers of the 3-rd base station);

Given B_t as a transmitting BS, we define a *colliding base station of B_t* as a base station that has at least one of its SSs within the transmitting range of the B_t .

Given St as a transmitting subscriber, we define a *colliding SS of St* as a SS associated with a different BS that has at least one of its SS within the transmitting range of St .

4.1 The IF approach

The IF approach is based on ensuring that there are no interfering transmissions sent throughout the system. As the FCH header is transmitted at full power and all the BSs are assumed to be synchronised, interference occurs during transmission of the FCH. Hence, the essence of this approach is to ensure that there is no interference during transmission of FCHs.

The approach to determining a scheduling for the BS frames operates as follows:

- 1) Determine *Transmission Allocation Table* (TAT): this table consists of the list of all of the BSs with their associated colliding BSs;
- 2) Order the TAT in order of descending colliding BS such that the BS with most colliding stations comes first;
- 3) Determine BS_i , the unallocated BS with the most colliding BSs – allocate a frame for BS_i ;
- 4) Determine BS_j , the unallocated BS with the most colliding BSs which does not collide with the other BSs that have been allocated a frame at this time.
- 5) If there is no such BS, then move on to the next frame time, else go back to step 4.
- 6) If all BSs have not been allocated a frame go to step 3

The SSs scheduling is performed using the given BS scheduling and applying round robin scheduling. Once this is determined, parameters such as the level of activity of each BS and the throughput for each node in the system can be determined without difficulty.

4.2 The Controlled Interference approach

The second approach seeks to maximize simultaneous transmission of all BSs. To this end, the area occupied by the nodes is divided into different so-called *interference zones*. For each interference zone, the activity status of each BS is defined; the nodes that fall into a single interference zone can all successfully receive a DL transmission from a specific BS without interference from other BSs. The (presumed circular) coverage area of each BS is thus divided into a number of interference zones and each SS is in a single zone.




Figure I: Example of interference zones in network with 3 BSs

Figure I shows a simple example of three interfering BSs. In this example there are 6 zones. 3 of them where the three of the BS are interfering which are the overlapping areas between these BSs. The other 3 which presents the interference free zones. Figure II shows some interfering zones for a system of 5 BSs that operate in the same area; Although the example is relatively general, it is interesting to notice that is very difficult to draw an example that includes all the possible interaction between BSs. For instance, the example does not show a stand-alone interaction either between BS1/BS3 and BS4/BS5.

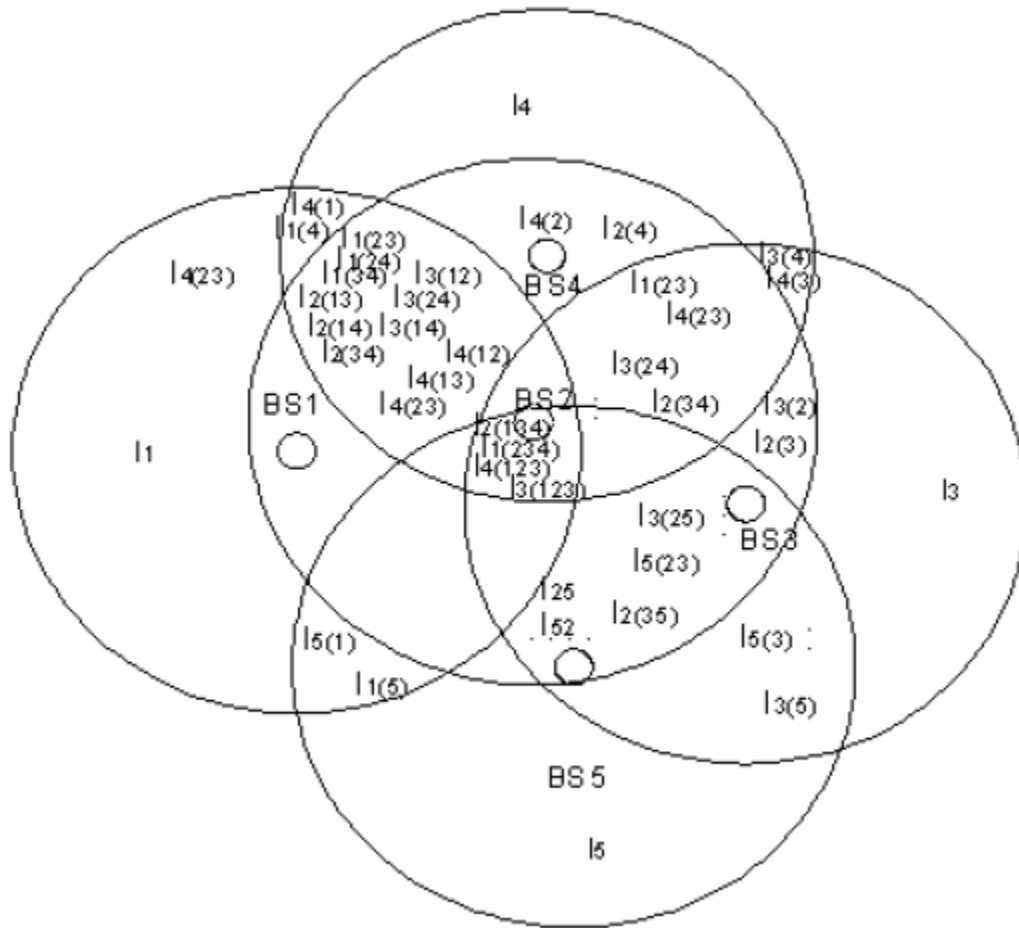


Figure II: More complex example showing interference zones concept

When all the interfering zones are identified then the number of SSs that belongs to each zone can be organized in a set of matrices generated according to the following considerations:

- The number of matrices is equal to the number of BSs in the system; we call them $M_0 \dots M_{N-1}$.
- The number of rows of each matrix is equal to the number of BSs in the system;
- The number of columns of each matrix is calculated by using the combination function $C(N,x)^1$ where $N = \text{Number of BSs in the system}$, $x = 1 \dots N$.

Matrix M_0 is vector of length N (an $N \times 1$ matrix). Each row of M_0 contains the number of SSs that can successfully receive the DL transmission from their associated BS when all other BSs are active. The Matrix M_1 is of dimension $N \times C(N,1)$. Each field of M_1 contains the number of SSs of each BS that fall into an interference zone in which exactly two BSs interfere with each other. For each of these SS to successfully receive DL transmissions, one or other BS must be silent at any time; this is taken into account when determining the system schedule. In general, each matrix, M_i , contains the number of SSs that fall into interference zones which require i silent BSs for successful DL transmission. The number of interference zones is dependent on i and hence the number of columns in M_i is also dependent on i .

¹ The notation $C(N,x)$ has been used here rather than ${}^N C_x$ to denote the combination function.

Table 4 shows the matrices that would be generated for a general case of N BSs. Note that:

- 1) $C_y(N,z)$ = y -th combination of the set of combinations of N elements taken z at time; (e.g. $C(4,2) = 12, 13, 14, 23, 24, 34 \rightarrow 14$ is the 3rd element).
- 2) $I_{x,C_y(N,z)}$ = Number of SSs that belongs to the BS x that interfere with BSs $C_y(N,z)$. (e.g. Number of SSs of BS y that require BSs 1 and 4 to be silent in order to receive data successfully from BS y are in $I_{1,C_y(4,2)} = I_{1,(1,4)}$)

$[M1]$
2 BS silent at time

BS 1	$I_{1,C_1(N,1)}$	$I_{1,C_2(N,1)}$...	$I_{1,C_N(N,1)}$
BS 2	$I_{2,C_1(N,1)}$	$I_{2,C_2(N,1)}$...	$I_{2,C_N(N,1)}$
...
BS N	$I_{N,C_1(N,1)}$	$I_{N,C_2(N,1)}$...	$I_{N,C_N(N,1)}$

$[M2]$
3 BSs silent at time

BS 1	$I_{1,C_1(N,2)}$	$I_{1,C_2(N,2)}$...	$I_{1,C_{N-1}(N,2)}$	$I_{1,C_N(N,2)}$
BS 2	$I_{2,C_1(N,2)}$	$I_{2,C_2(N,2)}$...	$I_{2,C_{N-1}(N,2)}$	$I_{2,C_N(N,2)}$
...
BS N	$I_{N,C_1(N,2)}$	$I_{N,C_2(N,2)}$...	$I_{N,C_{N-1}(N,2)}$	$I_{N,C_N(N,2)}$

...

$[M_{N-1}]$
 $N-1$ BSs silent at time

BS 1	$I_{1,C_1(N,N-1)}$	$I_{1,C_2(N,N-1)}$..	$I_{1,C_{N-1}(N,N-1)}$	$I_{1,C_N(N,N-1)}$
BS 2	$I_{2,C_1(N,N-1)}$	$I_{2,C_2(N,N-1)}$..	$I_{2,C_{N-1}(N,N-1)}$	$I_{2,C_N(N,N-1)}$
...
BS N	$I_{N,C_1(N,N-1)}$	$I_{N,C_2(N,N-1)}$..	$I_{N,C_{N-1}(N,N-1)}$	$I_{N,C_N(N,N-1)}$

Table 4

It is interesting to notice that the matrix M_0 is a vector while M_{N-1} is a diagonal matrix. In fact, M_0 consists of the SSs located out of range of all other BSs in the system; all of the SSs in this matrix can receive data from their respective BS simultaneously. On the other hand, M_{N-1} consists of the SSs in the area which is in range of all BSs. Each BS must transmit to the SSs in this area in sequence; all other BSs must remain silent in order for successful transmission to take place.

The above set of matrices can be used to determine a scheduling for the entire system. The objective, then, is to determine a periodic schedule which permits each BS to transmit to each SS at least once during the period. The period of the schedule can be kept as low as possible by allowing appropriate simultaneous transmissions to different interference zones.

A key observation in this process is that all of the SSs in a single column of a matrix can receive DL transmissions simultaneously without suffering from interference.

When considering how to determine the schedule a common trade-off arises, viz that of fairness against throughput. There are two natural approaches to optimizing the schedule – maximizing the fairness of the system by ensuring that each BS gets approximately fair access to the medium or maximizing the throughput of the system by, for example, frequently scheduling medium access to those BSs that interfere least with others. This trade-off gives rise to two variants of the CI approach: CI-T, which maximizes throughput and CI-F which maximizes fairness. Each of these are discussed in the separate sections below.

4.2.1 The CI-T approach

In this case, the main objective is to maximize the BS throughput. In this approach, a new set of matrices, $F_0 \dots F_{N-1}$, are generated based on the M matrices using the least common multiple function for each matrix entry as follows:

$$F_{(i,C(j,k))} = \text{LCM}(I_{i,C(j,k)} N_s), \quad I_{j,k} = 1 \dots N$$

The new set of matrices contains the smallest number of frames required to construct a ‘mini-schedule’ for a specific interference zone in which each SS receives the same amount of data from the BS.

The throughput of the system can be maximised by constructing a schedule in which the largest number of frames required to transmit to all interference zones characterised by those silent BSs is allocated. Any zones which require the same BSs to be silent but require a smaller number of frames are allocated the larger number of frames and can have a greater number of transmissions.

Using this approach, the total number of frames to allocate all SSs – the period of the schedule in frames – is given by:

$$\sum_{i=0..N-1} \{ \sum_{j=0..C(N,i)} \max F_i(\text{column } j) \}$$

4.2.2 The CI-F approach

In this case, the main objective is to achieve fairness between the different BSs in the system. In this approach, a set of matrices, G_0 - G_{N-1} are constructed using:

$$G_{(ijk)} = \text{ceil}(\max M_i(\text{column } j) / N_s) \quad i=0..N-1, \quad j=1..N$$

The G matrices contain the minimum number of frames required to enable transmission to each SSs in each interference zone. As with the CI-T approach, these matrices can then be used to determine the period of the schedule required to facilitate transmission to all SSs in the system at least once. This is given by:

$$\sum_{i=0..N-1} \{ \sum_{j=0..C(N,i)} \text{ceil}(\max M_i(\text{column } j) / N_s) \}$$

5. Simulation and Results

A simulator was developed throughout this work to model a number of different IEEE 802.16 systems operating in close proximity to each other on the same channel. In the following subsection an overview of the simulator is given followed by results and discussion subsection.

5.1 The IEEE 802.16 License-exempt Mode Simulator

The simulator that was developed can model a number of different 802.16 systems operating in point-to-multipoint mode. Specific aspects of the 802.16 MAC and PHY such as scheduling, power control and interference are modelled in the simulator.

The scheduling mechanism is designed to distribute resources equally between the SSs in the system. In the simulator, it was assumed that each frame can accommodate a fixed number of SSs. Indeed, it was assumed that each subscriber has access to the medium for 50ms at a time (50ms uplink and 50ms downlink), resulting in 5 subscribers being scheduled in 0.5ms frame, 10 in a 1ms frame and 20 in a 2ms frame.

The TDD frame duration is divided between the downlink and the uplink subframes; the duration of the downlink preamble is 2 OFDM symbols and the FCH is transmitted using one OFDM symbol. In the uplink, the preamble is 1 OFDM symbol. There is a bandwidth contention period of 1 mini-slot and an initial ranging contention period of 4 mini-slots. The downlink and uplink preambles and the contention periods are called the TDD frame overhead; the parameters defined in Table 1 can be used to determine the amount of time consumed by this overhead. The downlink overhead consists of 3 OFDM symbols, 2 for preamble and 1 for FCH, where the uplink overhead consists of 2 contention periods of the same duration as 10 physical-slots, and 10 OFDM symbols divided between the SSs as uplink preambles.

The BS capability to become inactive for one or more frame durations is useful in the context of the license-exempt environment as it provides an opportunity for others on the same channel to transmit. If all BSs were to transmit continuously, then there could be very substantial interference for all users in the system, rendering it quite useless for all users. To avoid this, a probabilistic approach in which each BS remains inactive for some period of time is used. More specifically, each BS is configured with an *activity factor* which controls what fraction of the time the BS is active for. For each frame, the probability of the BS being active is equal to the activity factor.

The BSs schedule their traffic using a rudimentary scheduling approach, where the BS looks at how many SSs it has and what frame duration it uses. Then it finds the *Lowest Common Multiple* (LCM) between the number of SSs and the number of SSs per-frame. After that, the LCM is divided by the number of subscribers. The result is the number of different schedules required. The schedules are then constructed by placing the SSs consecutively into frames as shown in Figure III.

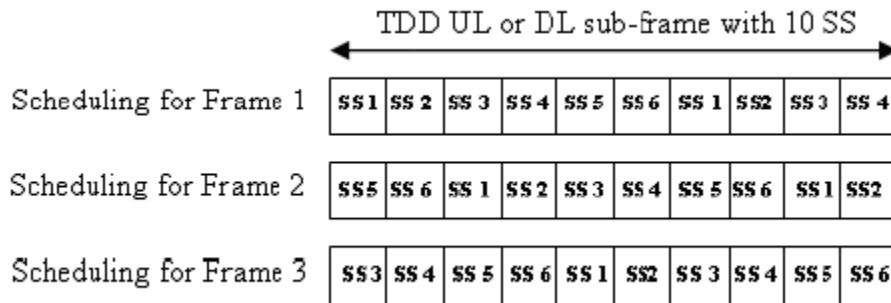


Figure III: Scheduling for 6 SSs

The downlink and the uplink transmission powers are calculated at the system start-up using a power control algorithm, which is based on the well-known Friis power transmission equation. Determination of interference is also based on power calculations. This is done by comparing the receiver power of signals which are

transmitted simultaneously. If the difference between these signals exceeds 5dB, then the signal can be received correctly; otherwise interference is deemed to have taken place. In the simulations below, each BS is active for a defined period of time. A random process controls which specific frames a station is active for, so, while the fraction of time a BS is active is defined, the particular active and idle frames for a BS vary from simulation to simulation. It is worth noting that an active time of 100% corresponds to the system being active all time; this is how the system would behave if there was no support for this sleep mode.

Parameter	Value
Wavelength	5.1238cm
Transmitter Gain	15 dBm
Receiver Gain	15 dBm
System Loss	1
Max Transmit Power	1mW
Bit rate	46.08 Mb/s
OFDM Symbol Time	12.5us
TTG and RTG	100 us
Minislot duration	0.347us
Channel BW	20 MHz

Table 1: Simulation Parameters

5.2 Results and Observations from Simulating the Distributed Scheme

To study the performance of IEEE 802.16 licensed exempt systems, we used the simulation parameters specified in Table 1. Random network topologies consisting of 4 BSs and 30 SSs per BS were used in these experiments. The network was located in a 100km² area as illustrated in Figure IV.




Figure IV: Topology of network 1

Several simulations were performed each lasting 60 sec of simulation time. The results were aggregated and analysed to provide graphical representation as shown in the following figures.

Figure V and Figure VI show the overall system throughput and the throughput per BS respectively. It can be seen from these figures that the performance of the system varies significantly with the activity factor. When the activity factor is low, the overall system throughput is quite low as each BS is inactive much of the time. When the activity factor is high, the overall throughput is also low, as there is much interference and few successful transmissions are made. There is a region around 25% to 40% during which time the throughput is maximised. At this point, every BS gets approximately 13% of the throughput that could be obtained if a BS had no interference issues. This can be compared to a scheme in which there is co-ordination between the BSs and the time could be divided such that each BS is active for 25% of the time. In this case, each BS could obtain 25% of that which it would obtain if it had exclusive access to the medium.

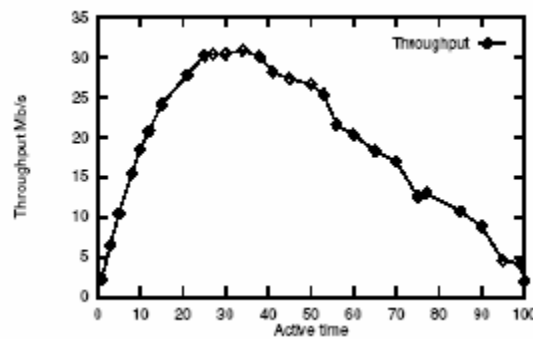


Figure V: Overall system throughput




Figure VI: Throughput per BS




Figure VII: Min, Ave, Max Throughput




Figure VIII: Average UL and DL throughput




Figure IX: Overall System Collisions




Figure X: Collision per BS




Figure XI: Min, Ave, Max Collisions




Figure XII: Average DL and UL Collisions




Figure XIII: Throughput at 50% Active Time

In Figure VII, the variation in the SSs throughput curve with the activity factor can be seen. This graph exhibits similar behaviour to that of the previous graphs - the throughput is low for both low and high activity factors, and is highest for some intermediate values. It can also be seen from this graph that there is a considerable variation in the throughput achieved by each SS, as evident from the significant difference between the minimum and maximum. Figure VIII shows that the average system downlink and uplink throughput coincides with Figure V and Figure VI and also shows the expected difference between the downlink and the uplink throughput due to the uplink overheads.

Figure IX and Figure X show the overall numbers of collisions on the system. From these results, it can be seen that there is a very linear relationship between the activity factor and the number of collisions in the system. Further, it can be seen that the collisions are divided pretty equally between all the BSs in the system. The numbers of collisions experienced by the SSs is shown in Figure XI. As with the previous graph, there is a quite linear increase in the mean number SS collisions with the activity factor. Also, as with the SS throughput, there is a significant variance in the amount of collisions that can be experienced by a SS. Figure XII shows the average uplink and downlink collision rate. It is clear from the figure that most of the collisions occur in the downlink. It is worth noting, however, that in many cases a collision in the downlink can result in a SS missing

an opportunity to transmit: if the SS does not receive the UL Map correctly, it does not know when to transmit and hence misses a transmission opportunity. The much greater number of collisions in the downlink can be attributed to the fact that some of the downlink information is transmitted at full power.

In Figure XIII the nodes in the system have been classified into those that obtain high throughput, medium throughput and low throughput. The results depicted in Figure 14 were generated using an activity factor of 50%. This classification is performed based on the difference from the overall mean throughput: nodes that obtain throughput of less than 50% of the mean throughput are deemed low throughput and nodes that obtain throughput of 50% greater than the mean are considered high throughput. It is clear from the results that the nodes which are located in the centre of the area obtain lower throughput and those at the extremities obtain significantly higher throughput. This is not surprising as those at the centre are more likely to experience interference.

5.3 Comparison between the Models and the Distributed Scheme

This section describes initial results obtained from comparing the performance of the distributed approach described in section 3 above and the performance predicted by the models described in section 4 above. For the evaluation we have chosen a scenario that consisted of a random distribution of 4 BSs of different providers with 25 SSs in each that operate in the same area network, as illustrated in Figure XIV.




Figure XIV: Network 2 Topology

Results for the distributed approach were obtained by simulating variants of the scenario for 10 minutes with different BS activity factors. As was shown previously, the overall system performance is heavily dependent on the activity factor and, Figure XV shows how it impacts the minimum throughput per SS. For comparison with the models, it was necessary to chose a single activity factor; the activity factor that resulted in the best overall system performance was chosen for this comparison. This was determined by averaging the four curves in Figure XV and choosing the maximum. The distributed approach achieves highest uplink and downlink overall throughput with a 30% activity factor per BS.

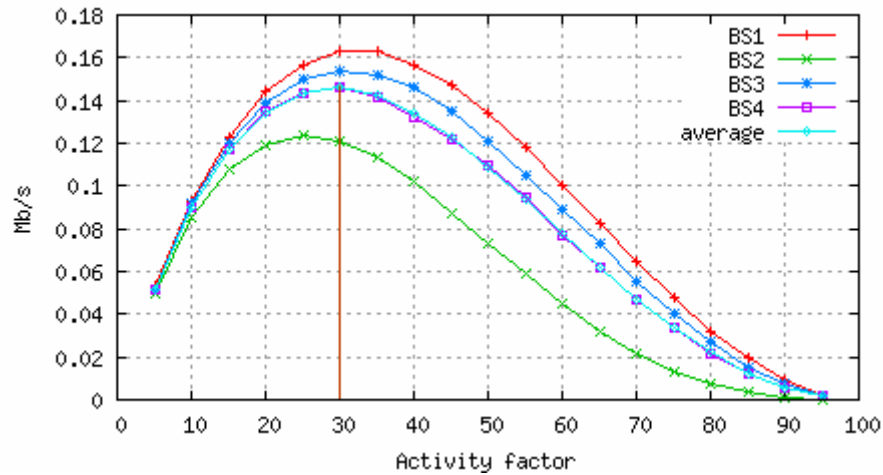


Figure XV: *the Overall Min Throughput Obtained using the Distributed Approach*

Once this was performed, the IF and CI approaches were applied to the chosen topology. In each case a schedule was obtained for the entire system. This could be used to determine parameters such as the overall throughput of the system and the level of activity of each of the BSs which could then be compared against results obtained from the distributed approach.

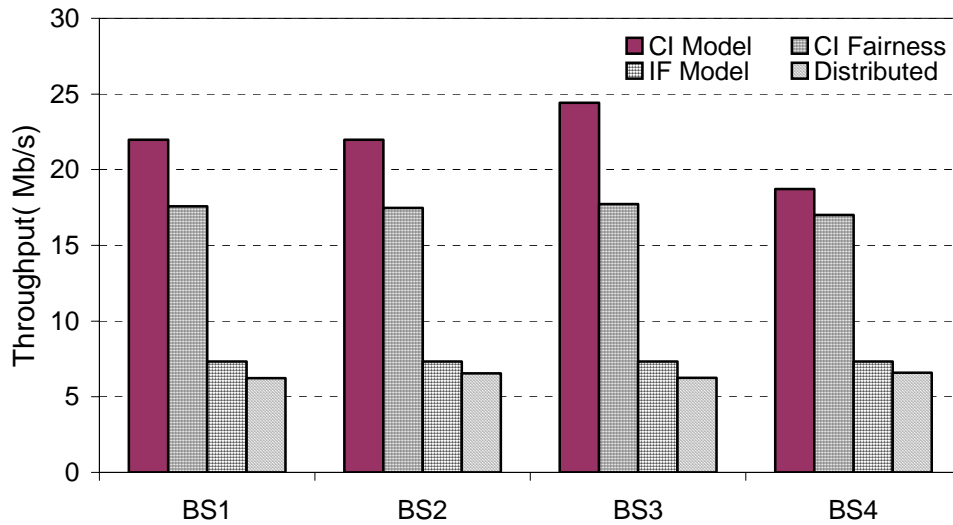


Figure XVI: *Overall System Throughput Per BS with 30% Activity Factor for the Distributed Approach*



Figure XVII: Average Throughput per SS at 30% Activity Factor for the Distributed Approach

Figure XVI and Figure XVII show a comparison between the overall and the average system throughput per BS obtained using the different approaches. The figures demonstrate that the CI approaches make most efficient use of the medium and result in greatest throughput. The difference between the throughput of the CI-T system and the CI-F system is clearly visible for all the BSs. Also, the equalisation of resources allocated to each BS is clear from these figures. Not surprisingly, the IF approach achieves much poorer performance than the CI approaches. Interestingly, however, the distributed approach performs worse even than the very conservative IF scheme. This indicates that there is a significant inefficiency associated with the lack of co-ordination and the collisions that arise within the distributed approach.

Another comparison between the four approaches is depicted in Figure XVIII: where the average system throughput per SS is shown for each approach. The figure shows the same characteristics as those discussed above. Once more, the distributed approach gives the lowest throughput of the four approaches.




Figure XVIII: *the Average System Throughput per SS with 30% Activity factor for the Distributed Approach*

Figure XIX shows the variations of the minimum throughput per SS for each BS by using the four approaches. In the IF approach the bandwidth is evenly divided between the SSs. Therefore, the same throughput is given for the average and the minimum throughput per SS. The difference between the minimum throughput for the CI-T and CI-F schemes is apparent in from the figure. The CI-T approach can starve some SSs in order to maximise throughput and hence it is possible that in some cases the minimum SS throughput is lower for the CI-T scheme. However, as the CI-T scheme makes more efficient use of the medium, for the case studied, it turns out that the CI-T scheme does not significantly starve SSs in order to maximise throughput. This would seem to indicate that the CI-T scheme is best overall for this topology.

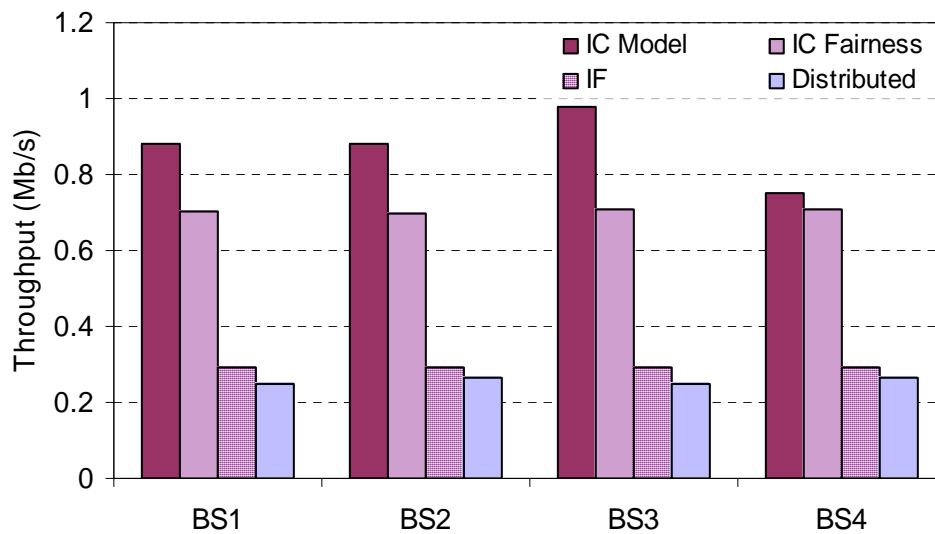


Figure XIX: *The Minimum Throughput per SS with 30% Activity Factor for the Distributed Approach*

It is also interesting to examine the level of BS activity obtained using these models. This is shown in *Table 2*. The activity factor chosen for all the BSs in the distributed approach is 30%, which is considerably less than that obtained by both the CI models. It is, however, greater than that obtained by the conservative IF model which only permits each BS to transmit while none of the others are transmitting.

	BS1	B2	B3	B4
IF Model	25%	25%	25%	25%
CI-T model	75%	75%	83%	63%
CI-F model	60%	60%	60%	60%

Table 2: The Calculated Activity Factor for the Three Models

The models produce results which clearly indicate that the performance of rudimentary distributed described above has much potential to be improved – the throughput is wanting and the fraction of time the BS is active can also be improved considerably. However, the distributed approach requires the addition of considerable intelligence to be able to determine – to some degree – where the interference zones are and how to perform much more efficient scheduling. This will be the subject of future work.

6. Conclusion

This paper describes some early work on a distributed approach to mitigate interferences issues that may arise in 802.16 systems operating in license-exempt operation. The approach is a natural one; it operates by enabling some BSs (and their associated SSs) to remain inactive, or asleep, for some periods of time, thereby permitting others in the same geographical area to use the limited available spectrum.

Three approaches to determining the potential performance of the system were also developed. All of these approaches assumed that the BSs in the system could be co-ordinated and the approaches focused on developing suitable scheduling for the entire system. In the IF approach, no interference between any nodes in the system was permitted, which in the CI-T and CI-F approaches, controlled interference was permitted in the system. The CI-T and CI-F approaches differed in terms of their objective: CI-T was optimised to maximise throughput, while CI-F was optimised to maximise fairness.

Simulation of the distributed approach showed that the throughput of the system is greatly increased by limiting the amount of time a BS is active, compared with allowing all the BSs to transmit all the time. For the case studied, the best system performance is obtained when each BS is active for quite a small fraction of the time (<40%). Another interesting finding in this work is that there is a significant discrepancy between the uplink and downlink performance: the downlink delivers better throughput due to the greater amount of overhead introduced by uplink overheads.

One issue that had a significant impact on the system performance was that of transmission of the FCH in the downlink. As all the BSs were synchronised and this information is transmitted at full power at the same point in a frame, this was frequently the cause of collisions.

The three approaches to modeling the system were then compared with the distributed approach in terms of the throughput of the system, the throughput per BS and level of activity of each BS in the system. The results showed that the performance of the distributed approach is poor, even compared to the conservative IF approach; the distributed approach delivers lower throughput, even though the BSs are active for more of the time. This can be attributed to collisions in the distributed approach. The CI-T and CI-F approaches show how much more efficiency can be achieved within the system. For the cases studied, they obtain over twice the

throughput of either of the other two approaches. Also, they have BS activity factors of up to 70% which is much higher than the 25% of the IF approach or the 30% of the distributed approach.

The results indicate that there is considerable scope for improving the performance of the distributed approach. Two obvious approaches to improving the system performance are to consider an unsynchronized variant of the system and to implement more sophisticated power control mechanisms.

Acknowledgments

The support of the Informatics Research initiative of Enterprise Ireland is gratefully acknowledged.

References

- [1] S. Mangold, J. Habetha, S. Choi and C. Ngo “Co-existence and Interworking of IEEE802.11a and ETSI BRAN HiperLAN/2 in MultiHop Scenario”. In Proc. IEEE Workshop on Wireless Local Area Network. Sep. 2001
- [2] N. Golmie and F. Mouveaux “Interference in the 2.4 GHz ISM Band: Impact on the Bluetooth Access Control Performance”. In Proc. IEEE International Conference on Communication. June. 2001
- [3] H. Wei, S. Ganguly, R. Izmailov and Z. Haas “ Interference-aware IEEE 802.16 WiMax Mesh Networks”. In Proc. IEEE VTC. Spring. 2005
- [4] C. Hoymann, M. Putter, I. Forkel “Initial Performance Evaluation and Analysis of the global OFDM Metropolitan Area Network Standard IEEE 802.16a/ETSIHiperMAN”. In Proc. European Wireless conf. Feb. 2004
- [5] C. Hoymann, M. Putter, I. Forkel “The HIPERMAN Standard Performance and Analysis”. In IST. Mobile and Wireless Communication Summit. June. 2003
- [6] S. Ramachandran and C. Bostian, S. Midkiff “Performance Evaluation of IEEE 802.16 for Broadband Wireless Access”. In Proc. OPENNETWORK. Aug. 2002
- [7] A. Ghosh, D. R. Wolter, J. G. Andrews and R. Chen “Broadband Wireless Access with WiMax/802.16: Current Performance Benchmarks and Future Potential”. IEEE Communication Magazine Feb 2005
- [8] A. Mostafa, R. Kobylinski, I. Kostanic and M. Austin “Single Antenna Interference Cancellation (SAIC) for GSM Networks” In Proc. IEEE VTC. 2003
- [9] J. G. Andrews “Interference Cancellation in Cellular Systems: A Contemporary Overview” IEEE Wireless Communication. April 2005