Strengthening Real-Time Support in Wireless Networks

Mark Gleeson

A Thesis submitted to the University of Dublin, Trinity College

in fulfillment of the requirements for the degree of

Doctor of Philosophy (Computer Science)

May 2010

Declaration

I, the undersigned, declare that this work has not previously been submitted to this or any other University, and that unless otherwise stated, it is entirely my own work.

Mark Gleeson

Dated: May 10, 2010

Permission to Lend and/or Copy

I, the undersigned, agree that Trinity College Library may lend or copy this Thesis upon request.

Mark Gleeson

Dated: May 10, 2010

Acknowledgements

The journey toward submitting a Ph.D. thesis is a long one, of unknown duration and destination, many potential points where you can change direction and even more dead ends. Needless to say are no maps or timetables to guide you along the way, for that you must rely on your fellow passengers as we struggle together to find the best route towards our destination, the completed thesis.

Throughout this journey many passengers, both past and present members of the Distributed Systems Group provided never ending support and advice on finding my way. I must give particular thanks to Dr Raymond Cunningham and Dr Barbara Hughes whose work laid the foundations upon which this thesis was built. Thanks also to Shu Zhang whose MSc work provided a practical test of much of the implementation enabling many issues to be resolved. Prof Vinny Cahill and Dr Mélanie Bouroche for making me think about and explain what my work was about. Finally to everyone in the Distributed Systems Group for creating such a positive and supportive environment in which to carry out research. And to the many dozens of people, past and present members of the community that is Trinity College for their support (and timely distractions!) from the research.

The journey would have neither started or finished without my supervisor Dr Stefan Weber, who always had advice on how to move forward and most of all, always had time to discuss issues. A very special thanks to Dr Eavan O'Brien whose friendship and continuing understanding without which I very much doubt the journey would have finally found its destination. Many other ‘guides’ deserve special mention, Meriel Huggard whose practical advice opened my mind to new solutions and Dr Jacqueline Potter whose advice pointed me towards the light at the end of the tunnel.

Mark Gleeson
University of Dublin, Trinity College
May 2010

Abstract

Wireless networks exhibit unpredictable and varying connection reliability as a result of node mobility and resultant changes in wireless signal propagation. Wireless signal propagation not only depends on the receiver's mobility but also on unrelated objects both mobile and fixed in the environment. In such dynamic and unknown environments, the provision of real-time communication represents a significant challenge. Not only must transmissions be managed to avoid collisions but may also need to be rescheduled in response to communication failures to ensure success retransmission of frames following failures.

Previous work in this area has considered the support for real-time communication with significant constraints, by either relying on master/slave network architectures or through approri knowledge of the network structure. These approaches focus on the creation of collision-free transmission schedules determined off-line, distributed before the execution of the system or through a centralised contention-based resource allocation approach.

In this thesis, the need to provide real-time support in a wireless environment in the absence of a master node is discussed. In light of varying communication and node reliability, a distributed solution is chosen to provide resilience in the presence of failures.

Unlike other TDMA protocols, the impact of communication-related failures in the form of burst-and single-bit errors is addressed. The protocol reacts to overcome burst errors where possible by exploiting the statistical independence of communication between individual wireless nodes in conjunction with dynamic rescheduling in response to communication errors to enhance real-time support in the wireless domain.

Transmission reliability is addressed through an admissions control process based on clustering transmission resources of destinations with similar transmission probabilities through the application of binomial distribution. This process estimates the overall number of transmissions required to meet the real-time requirements of frames. The application of the clustering algorithm reduces the total number of transmissions required.

This thesis makes two key contributions: Firstly, a medium access control protocol, Hierarchical Distributed Time Division Multiple Access (HD-TDMA), is proposed, which incorporates in its design flexibility to permit multiple transmissions per TDMA slot. Secondly, a local, autonomous decision making process exploits the flexibility of HD-TDMA to schedule packets to meet real-time requirements while incorporating the ability to dynamically reschedule packets to overcome burst error conditions while, attempting to maintain real-time deadlines.

Publications Related to this Ph.D.

[1] Mark Gleeson and Stefan Weber. "Towards Optimised Retransmission Reservation in Real-Time Wireless TDMA", 8th International Workshop on Real-Time Networks (RTN), pages 1-6, July 2009, Dublin, Ireland.

[2] Mark Gleeson and Stefan Weber, "Strengthening Real-Time Communication Support In Wireless Networks". Work in Progress Session of the 19th Euromirco Conference on Real-Time Systems (ECRTS), pages 41-45, July 2007, Pisa, Italy.

[3] Mark Gleeson and Stefan Weber. "Fault Recovery and Redundancy in Real-time Wireless TDMA". Technical Report TCD-CS-2008-24, May 2008.

[4] Barbara Hughes, Mark Gleeson, Marcin Karpinski, Raymond Cunningham and Vinny Cahill, "Real-Time Communication in IEEE 802.11 Mobile Ad hoc Networks A Feasibility Study". Technical Report TCD-CS-2006-55, September 2006.

Contents

Chapter 2 Related Work 12

List of Tables

List of Figures

Chapter 1

Introduction

This thesis presents an approach to medium access control in the wireless communication domain, which addresses the problem of supporting real-time communication over communication links with time varying reliability.

The shared nature of wireless media allows transmissions to be interfered with by transmissions from other sources and/or general noise in the environment. The interference from other transmissions is generally addressed by employing protocols such as Time Division Multiple Access (TDMA) [Caccamo et al., 2002, Cunningham and Cahill, 2002, Sobral and Becker, 2008] with an aim to prevent concurrent transmission attempts. The remaining frame losses are results of dynamic propagation effects in the environment and take the form of burst errors of varying durations [Otani et al., 1981, Zorzi, 1998, Willig et al., 2002].

In the presence of burst errors, the attempt to immediately retransmit a frame following a failed transmission attempt has a high probability of failure, as the burst error condition which caused the failure, may still be present. Carrier Sense Multiple Access (CSMA) and Medium Access Collision Avoidance (MACA) protocols, which view all frame loss as being caused by contention, will attempt to resend a failed frame immediately or following a slight pause.

Basing the approach on a slotted time division multiple access (TDMA), each station is allocated one or more slots, permitting contention free medium access. As contention is eliminated between data transmissions, the focus switches as to how to perform, maintain and adapt this resource allocation, together with recovering from frame loss due to transient propagation effects through retransmissions.

This chapter introduces the background to this thesis, outlining the challenges presented by wireless networks where communication reliability is variable, defining what is meant by real-time requirements together with requirements for a real-time wireless communication systems. The probabilistic admissions control approach is introduced followed by the validation approach adopted considering the impact of channel modelling. Two contributions of this thesis are outlined, a protocol providing flexible TDMA style medium access, HD-TDMA, and a probabilistic admissions control and retransmission process.

Certain issues will not be addressed by this work, these are outlined within section 1.7. At the end of this chapter a road map is presented outlining how the rest of the thesis is structured.

1.1 Medium Access Control

The past two decades have seen a move away from the highly centralised single server mainframe, towards complex interconnected distributed systems. To ensure the successful operation of a distributed system, reliable communication among all hosts which make up such a system is essential. Within a fixed wired communications infrastructure, resources may be dedicated to provide redundancy to overcome failures [Kopetz and Grunsteidl, 1993].

Being tethered by wire, restricts the scope of distributed systems. But with the advent of digital wireless communication systems, a new class of distributed system, one which exploits mobility as a result of the use of a wireless communications system to provide functionality, evolved. The very characteristic of the wireless medium which makes such applications possible, mobility results in dynamic membership of the system not only as a result of the nature of the distributed system but also as a result of stations moving beyond wireless communications range.

Supporting the reliability requirements of a distributed system represent a significant challenge, when communication takes place over a wireless medium. The wireless medium is shared with other potentially unrelated stations and is heavily resource constrained in terms of both bandwidth and communications range.

To support many stations wishing to transmit concurrently, coordination is required. This coordination function is known as the medium access control (MAC) protocol. A MAC protocol is a protocol that is used to arbitrate access to the medium with the resulting transmission of data on that medium [Johnsson, 1999]. It defines how and when stations may access the medium [Carley et al., 2003]. The problem of wireless medium access control may be divided into two distinct areas: The manner in which access to the wireless medium is arbitrated and the nature of infrastructure or lack thereof available to support the communication.

In an infrastructure-based network, a portion of the network is fixed. This portion may contain one or more devices, such as base stations or access points, which can provide coordination services to all other hosts and centrally regulate access to the medium. Given the presence of a fixed infrastructure and potentially a fixed typology, a station can be designated as a master which establishes and maintains a transmission schedule by allocating resources in the form of time or frequency slots to participating hosts.

In the absence of infrastructure, there is no fixed, dependable resource to manage the communication, as such the network is termed ad-hoc. The stations must therefore coordinate there transmissions in an decentralised fashion. Random access approaches such as CSMA and MACA require no coordinator and no exchange of state information between stations. All frame losses and errors are interpreted as collisions. This results in unpredictable delays and poor resource utilisation. An alternative approach would be to dynamically create and maintain a transmission schedule through optimistic transmission attempts which become reservations if successful [João L Sobrinho and A. S. Krishnakumar, 1996, Lin and Gerla, 1999], thus avoiding collissions.

The medium access approach must efficiently and fairly share the available bandwidth. Additional control information, be it to handle allocation of resources or recovery from error must be minimised in order to gain maximum utility from the wireless channel. Each transmission increases contention for the medium and reduces the available bandwidth for data transmissions. In order to meet real-time requirements, management of slots should be bounded in duration, be fault tolerant and not in conflict with real-time traffic flows.

1.2 Real-Time Requirements

A real-time computer system is a computer system in which the correctness of the system behaviour depends not only on the logical results of the computations, but also on the physical instant at which these results are produced [Kopetz, 1997].

A hard real-time system is one in which the deadlines must always be satisfied. In the context of a communication system, messages must be delivered correctly and within given deadlines 100% of the time. Failure to meet a hard real-time requirement may result in a system failure.

Under soft real-time requirements, failure to satisfy deadlines 100% of the time will not result in a critical failure of the system, but may result in degraded performance. The message contents may have utility after its deadline. This utility is variable and specific to the application and cannot be determined or quantified by the communications system.

A special sub-class of real-time systems exists where failure to satisfy the deadline does not result in a critical failure, but distinct from soft real-time, the delayed message has no utility. Such systems are described as being firm real-time systems. This thesis addresses the provision of firm real-time support in wireless networks; therefore soft real-time concerns of message utility post deadline are not considered.

1.2.1 Wireless-Specific Real-Time Concerns

The use of wireless communication results in several additional concerns, particularly in terms of communication reliability and available bandwidth. Wireless communication is subject to significantly greater error rates than those found in wired systems [Nicopolitidis et al., 2004, Eckhardt and Steenkiste, 1998]. Communication may be subject to bursts of errors of random length [Willig et al., 2002].

When an error does occur, it is not possible for the sender to detect the error as it would be in the case of a collision on an Ethernet network for example. To verify successful reception of a packet, a positive acknowledgment scheme is required. These additional control packets further increase the use of the medium and competition for access to it as well as being expensive in terms of use of the available bandwidth.

Based on Koptez et al [Kopetz, 1997] a set of requirements for a real-time wireless communications system is produced:

- **Predictable Latency**: In order to meet real-time deadlines, the worst case delivery times of a message must be known before a request is made to send the message. Jitter must be minimised in order to meet application requirements.

- **Admissions Control**: Communications bandwidth is limited. To ensure real-time service is maintained the communications system must not be overloaded. Access should be permitted only when the request is not in conflict with all other real-time commitments.

- **Dynamic Reconfiguration**: The communication system should be able to support a variety of configurations and be able to adapt to changes in group membership at runtime.

- **Error Detection**: Given a challenging wireless environment, it is essential that errors be detected and resolved. When an error occurs, the recovery from that error should be possible without exceeding the latency and jitter requirements.

These requirements place a burden on the performance of both the medium access control and scheduling algorithms to sustain real-time guarantees in the face of a dynamic and unpredictable environment.

1.3 Probabilistic Admissions Control

The aim of an admissions control process is to admit only those tasks for which adequate resources exist to perform the task under worse case conditions. However, this results in a significant amount of resources allocated to cope with rare events leading to poor overall resource utilisation [Hamann et al., 2007].

The traditional approach in admissions control is to over-estimate the resource requirements in order to be able to provide guarantees under degraded conditions. This function is performed in isolation on a task by task, message by message basis. Reserving more resources than are required in order to ensure the correct delivery of messages within defined time-bounds, as each event, task, message is considered independent.

This thesis proposes to consolidate transmission resources between all messages in order to minimise the overall resource requirements while continuing to satisfy real-time requirements. Frames to each destination are modelled as a binomial process, a clustering process groups destinations of similar probability in order to minimise the transmission resources required while continuing to meet or exceed the real-time requirements.

1.4 Problem Statement

Communication over a wireless medium is subject to varying reliability owing to lack of coordination between stations, resulting in multiple concurrent conflicting transmissions combined with time varying propagation conditions as a result of both fixed and mobile elements of the environment.

In order to address these issues a new approach which combines medium access control, transmission scheduling and transmissions reliability is required in order to mitigate the impact of error-prone wireless communication links in order to support firm real-time communication.

1.5 Thesis

To address the above issues the Hierarchal Distributed Time Division Multiple Access Medium Access Control protocol (HD-TDMA) is proposed. HD-TDMA offers individual hosts the ability to transmit more than one frame within each TDMA transmission slot. As many frames may be transmitted within each TDMA slot, the order of transmissions may be dynamically changed. This scheduling flexibility offered by HD-TDMA is exploited by a probabilistic admissions control and retransmission process which relies on a clustered binomial process to determine the number of transmissions, including retransmissions required to meet the real-time reliability requirements. Frames are scheduled in order to maximise probability of successful reception in order to satisfy real-time requirements.

Dynamic allocation of TDMA slots is provided for by employing a modified version of the TBMAC [Cunningham and Cahill, 2002] membership protocol which provides time-bounded slot allocation, ensuring each station has a consistent view of the slot allocations. Membership management messages are passed in band using dedicated frames in such a manner as to be independent of real-time traffic.

Scheduling based on awareness of the bursty error phase reduces the possibility of a transmission failure and introduces independence between transmission attempts to the same destination. A probabilistic admission control process, based on a clustered binomial distribution is employed to estimate the total number of transmissions required to ensure real-time requirements are met, while minimising the overall resources required.

1.6 Validation

The combination of the HD-TDMA protocol and the probabilistic admission control and retransmission process implement a firm real-time wireless communication sub-system. An evaluation of this sub-system must verify that the proposed model provides an effective response to dynamically varying wireless conditions.

A series of simulations are performed comparing the performance of the proposed model against other approaches. This evaluation seeks to experimentally verify that the proposed communications architecture provides a robust approach when faced with random single bit errors as well as a combination of single bit errors and bursty error conditions.

1.6.1 Channel Modeling

In order to react effectively to packet loss, it is essential to understand the characteristics of the wireless channel. Much previous work in the field of medium access control has ignored the channel, assuming packet loss to be a result of contention, a station failing or moving out of communication range. The success of frame reception was based on overly simplistic models based purely on distance, assuming a flat unobstructed path.

Combined the two-ray model and Rayleigh fading model provide a more realistic propagation model. However, they fail to account for the layout of the local environment. Alternatives based on the ray optical model which employ ray tracing [Stepanov and Rothermel, 2008], the scale of the undertaking to accurately model the signal paths in a dense urban environment is mammoth. Ray tracing performed off-line fails to acknowledge that in addition to the mobile wireless devices the urban environment is dynamic. Mobile entities such as buses and trams can significantly alter the propagation paths.

While accurate knowledge of signal propagation path assists in simulation, it provides no insight into the design of medium access control approaches. It is infeasible for a mobile device to have global knowledge of signal propagation paths, such that it could determine the success or failure of reception of a message before attempting to transmit.

Previous experimental work [Otani et al., 1981, Willig et al., 2002] suggests that the burst length of an error follows a long-tailed distribution, therefore a Pareto distribution is introduced; to model the length of error bursts. A normally distributed random process models the arrival rate of bursts. The burst model is then combined with the single bit error random process, in order to model both single and burst errors. This approach may be thought of as a constrained Fritchman model [Fritchman, 1967] as it combines the normally distributed single bit error (BER) model with the Gilbert & Elliot [Gilbert, 1960] two state channel model.

1.7 Scope

The goal of this work is to provide a wireless real-time communication sub-system which can maintain a high level of real-time communication guarantees in the presence of channel errors, specifically bursty channel conditions as a result of the dynamic changes in propagation caused by mobility and the environment.

The work in this thesis is limited to single hop communication. As typified by the OSI model [ISO/IEC, 1994], a fully featured communication system is made up of many layers. This work concentrates on the data link layer.

HD-TDMA provides functionality to the layers above. Knowledge of the exact location and number of slots allocated for instance, allow higher layers to make decisions as to obtaining the shortest end to end delay. Inter-cell communication and real-time routing of traffic between cells will not be addressed in this work. The HD-TDMA protocol is compatible in key aspects with TBMAC [Cunningham, 2003, Cunningham and Cahill, 2002]. Work such as the Space Elastic Routing protocol, SEAR [Hughes, 2007,Zhang, 2008] may be used in conjunction with this work to provide routing services with real-time guarantees.

This thesis does not address the issue of clock synchronisation. A number of solutions such as GPS or direct reference to an atomic clock are available to provide this. The issue of the application mapping its periodic transmission requirements to match the underlying transmission resources, i.e the data transmission slots is also not considered in this work.

1.8 Road Map

The remainder of this thesis is structured as follows:

Chapter 2 presents a review of existing literature with respect to real-time communication in the wireless domain, medium access control approaches and the nature and impact of communication errors in wireless communication.

Chapter 3 presents a description of the Hierarchal Distributed Time Division Multiple Access (HD-TDMA) protocol.

Chapter 4 describes the probabilistic admissions control and retransmission process, including a mathematical model to calaculate the transmission resources required.

Chapter 5 describes the implementation of the communication architecture elements in the OPNET modeler descrete event based simulator.

Chapter 6 demonstrates the performance analysis of HD-TDMA and the probabilistic admissions control process through simulation, as both individual units and as a combined communications sub-system.

Chapter 7 presents conclusions and outlines possible future works.

1.9 Summary

This chapter outlined the subject and goals of this thesis. Issues which are not directly addressed by this work were highlighted. Background information was presented relating to wireless medium access control and the issue of real-time requirements in order to outline the challenges in supporting real-time communication over unreliable wireless links.

This chapter also outlined the contributions of this work, the Hierarchal Distributed Time Division Multiple Access protocol, HD-TDMA, the related admissions and retransmission process and an approach to modelling bursty error conditions on unreliable wireless links.

Chapter 2

Related Work

This chapter will investigate and discuss a range of issues with respect to provision of real-time communications support in a dynamic ad-hoc wireless environment. A number of medium access control protocols are introduced and their suitability for the provision of timely message delivery in a potentially error prone communications channel is considered.

The key concern is not only that of timeliness, but also reliability of message delivery. Not only must the message be received correctly but, it must also be received before its real-time deadline. If a packet transmission fails, the recovery process in the form of one or more retransmissions must also complete within the real-time deadline of this packet, without compromising the real-time service offered to other packets throughout the communications system.

This chapter focuses on two key wireless communication concerns, firstly the broad issue of medium access control is discussed in the form of contention-based and agreement protocols, expanding to discuss the similarities and differences between real-time communication in the wired and wireless domain in the context of seeking protocol support to overcome the variables of the wireless domain in terms not only of potential mobility, but also in minimising the impact of transmission failures.

Secondly the aspects of wireless signal propagation and the constraints under which wireless communication takes place, must be analysed and discussed to understand the context in which errors occur. The nature of these errors must be further analysed to determine their distribution of duration and frequency of occurrence. The environmental contribution in the form of propagation effects and background noise and the impact on the reliability of communication are discussed together with scheduling approaches to overcome or to minimised the impact of errors and to reduce instances of repeated errors.

Structured protocols such as TDMA require agreement as to slot allocations or transmissions times in advance. In fixed wired environments the schedule may be fixed and delivered off-line in advance, however with dynamic network topologies found in the wireless domain, agreement of schedule updates and delivery to participating stations without impact on ongoing real-time communication is essential.

Related work commences with section 2.1 which provides a general overview of medium access control concepts and protocols. Section 2.2 discusses the nature and impact of errors as a result of wireless signal propagation. A discussion of real-time communication and how it made by achieved is contained in section 2.3. Section 2.4 expands on section 2.1 through the investigation of the real-time support and capabilities of a number of wireless medium access control protocols. Based on the discussion contained within the previous sections, section 2.5 presents a list of issues which must be addressed to support real-time communication over a wireless medium. The chapter then concludes with a summary.

2.1 Medium Access Control

A radio channel cannot be accessed simultaneously by two or more stations which are within radio interference range [Rhee et al., 2005], while using the same frequency or spreading code. To support many stations wishing to transmit concurrently, a coordination function is required as packet collisions are unavoidable otherwise due to the fact that traffic arrivals are random [Hassanein et al., 2005]. This co-ordination function is known as the medium access protocol or MAC. The MAC protocol is the protocol used to arbitrate access to the medium, the radio link in a wireless network with the resulting transmission of data on that medium [Johnsson, 1999]. It defines how and when stations may access the medium [Carley et al., 2003]. Medium access control is an important technique that enables the successful operation of the network [Ye et al., 2004].

Ideally in the wireless domain, a protocol should be flexible to adapt to conditions as a result of mobility and time varying propagation and medium reliability. Available bandwidth is constrained, with communication links typically half duplex requiring complex control mechanisms to ensure transmissions do not interfere with each other. As stations may not be in communication range with all other stations, ensuring that both hidden and exposed terminal problems are eliminated or minimised is essential. Guaranteed delivery and bounded delay are fundamental requirements for implementing real-time services in any networked environment [Basagni and Bruschi, 1999].

The medium access control protocol implements two distinct functions, first that of access arbitration and secondly that of transmission control [Malcolm and Zhao, 1995]. The first goal is achieveable by allowing each node on the network to schedule its transmissions such that collisions are avoided [Basagni and Bruschi, 1999]. The second goal requires the ability to detect failed transmissions and to manage retransmissions in such a way as not to impact on other transmissions [Sudhaakar et al., 2009].

Murthy et al [Murthy and Manoj, 2004] present eight factors which they consider any medium access protocol for wireless networks should support;

- The operation of the protocol should be distributed.
- The protocol should provide QoS support for real-time traffic.
- The access delay, which refers to the average delay experience by any packet to get transmitted, must be kept low.
- The available bandwidth must be used efficiently.

- The protocol should ensure fair allocation (either equal or weighted allocation) of bandwidth to nodes.
- Control overhead must be kept as low as possible.
- The protocol should minimize the effects of the hidden and exposed terminal problems.
- The protocol must be scalable to large networks.

Approaches to medium access control may be classified into two distinct groups:

- Contention-Based Protocols
- Agreement Following Contention

Contention-based protocols, such as CSMA [Kleinrock and Tobagi, 1975] and MACA [Karn, 1990] were derived from ALOHA [Abramson, 1970] and are based on traditional computing assumptions of low average usage, with infrequent bursts of heavy usage. Such assumptions may no longer hold, particularly in a sensor environment where many sensors may potentially observe the same event and attempt to transmit at the same time resulting in poor throughput but high medium utilisation or in the case of multimedia where a constant stream of data must be transmitted.

Agreement following contention approaches aim to set up a transmission schedule in an ad-hoc fashion. Such an approach may be particularly valid for periodic communication with either constant bit rate sources or on/off sources. Station contend for access typically using an adapted form of the MACA request to send, clear to send exchange, if the exchange is successful a reservation exists.

2.1.1 Contention-Based Protocols

An early packetised wireless communication protocol was Aloha [Abramson, 1970]. Under the Aloha protocol, whenever a station had a packet to transmit, the station simply transmitted the packet without regard to any other ongoing transmissions. The destination station issued an acknowledgment in response to a correctly received packet.

While this approach relied on assumptions of low average usage, with short bursts of intense use, with asymmetric traffic flow, the throughput peaked at 18% of available bandwidth. Increasing the load beyond this level resulting in overload resulting in a collapse in throughput.

A refinement of Aloha, Slotted Aloha [Roberts, 1975] utlises a centralised station to send out periodically, a small clock tick packet to the outlying stations to signify the start of a new transmission slot. The introduction of a slotted structure almost doubles the performance compared to original Aloha [Roberts, 1975]. Namislo et al [Namislo, 1984] through the application of a Markov model indicates a maximum efficiency of 0.5 to 0.65 packets per aloha slot, which aligns with Roberts et al [Roberts, 1975] suggestion of up to 0.6 being achievable using an FM receiver compared to satellite owing to capture effects.

All variants of Aloha suffer from the lack of medium sensing before transmission which severely impacts on performance. The problem of the hidden terminal [Tobagi and Kleinrock, 1975] also arises as stations have no knowledge of other ongoing transmissions.

In response, a listen before send family of protocols evolved. Implementations differ in the manner in which the wireless mediums current state is determined. The virtual approach employed by Multiple Access with Collision Avoidance (MACA) [Karn, 1990] and Multiple Access with Collision Avoidance for Wireless Networks (MACAW) [Bharghavan et al., 1994] rely on the reception of control packets sent by both the sender and receiver of the data packet. The CSMA [Kleinrock and Tobagi, 1975], carrier sense multiple access family actively listen to the medium, with the medium viewed as active if a packet is being received or the background noise level exceeds a fixed threshold.

As all stations contend for medium access in both MACA and CSMA a possibility that transmissions will overlap exists, resulting in packet losses. Retransmissions are required and a noticeable delay appears [Hassanein et al., 2005].

Shepard et al [Shepard, 1996] and Tobagi et al [Tobagi and Kleinrock, 1975] identify the nature of collisions in wireless communication, which occur due to lack of coordination between stations, collectively this issues are known as the hidden terminal problem, shown in figure 2.1 where the transmission of a station (A) conflicts with the reception or sending of a packet somewhere else (from station C), the sending station could not know aprori of the potential for collision.

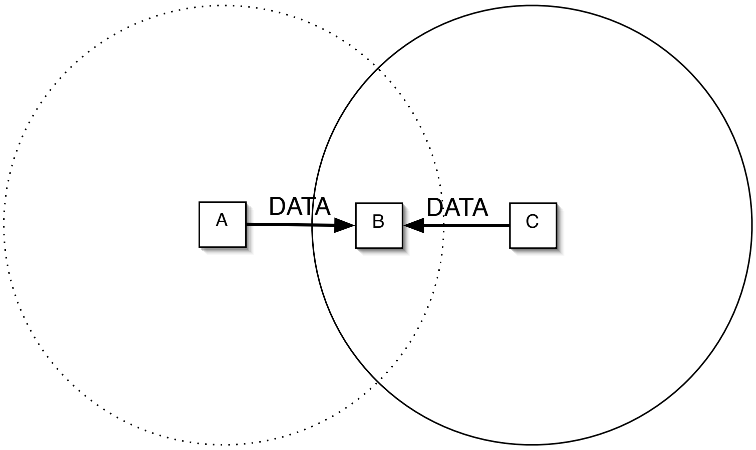

Fig. 2.1: Hidden Terminal

In a large scale network, where not all stations are within single hop communication range, despite the medium being viewed as idle two stations sufficiently far apart not to receive each others transmissions could choose to transmit at approximately the same time to a station which is in range of both transmitters, the transmission will fail despite both stations observing the medium as idle. This phenomena is known as the hidden terminal problem [Tobagi and Kleinrock, 1975] as collisions occur at the receiver, of which the transmitter has no knowledge.

The virtual channel sensing approach in the form of MACA [Karn, 1990] and MACAW [Bharghavan et al., 1994] requires additional small control packets to setup each communication, which in doing so aims to eliminate the hidden terminal problem. The sending station firstly checks if it aware of any ongoing transmissions by checking its network allocation vector, a time based structure which maintains a record of when the medium will become idle following the current transmission transaction.

A station intending to send a data frame first sends a 'request to send' frame (RTS), containing the duration of the proposed full transmission i.e. the time to transmit the RTS, CTS, DATA and acknowledgement frames. The receiving station responds with 'clear to send' (CTS) if it is free to transmit. Again the CTS frame contains the expected duration information less the duration already expended by the RTS transmission. In doing so all stations in communication range of both sender and receiver are aware of the imminent transmission and its duration and defer all transmission attempts until the duration indicated in the RTS/CTS exchange has elapsed.

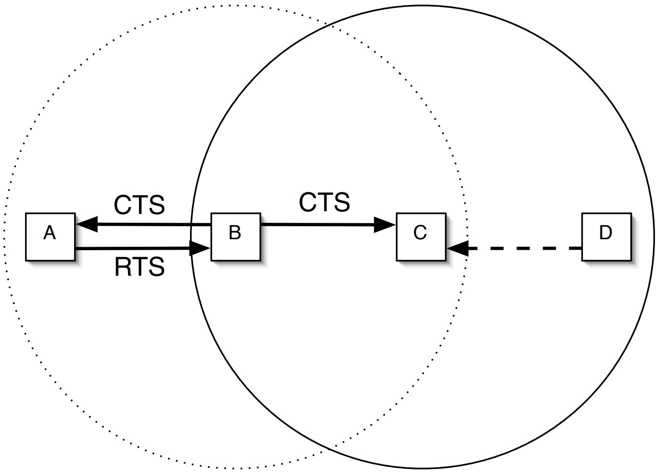

Fig. 2.2: Exposed Terminal

While the RTS-CTS exchange resolves the hidden terminal problem it does introduce the exposed terminal problem, where non interfering transmissions become blocked as a result of the RTS-CTS exchange. Figure 2.2 presents an example of this. Station A is transmitting to station B and the same time as D attempts to transmit to C but the control RTS-CTS exchange sent by A and B will cause C and D to be unable to communicate, as station C having received a CTS from station B will be unable to send an RTS frame itself or respond to an RTS frame from station D until the time indicated in the CTS frame from station B has elapsed.

In some cases an arbitration process is performed utilising a second dedicated control channel. BTMA, busy tone multiple access [Tobagi and Kleinrock, 1975] utlises this control channel to broadcast a busy tone whenever a carrier is detected on the data channel. Split channel reservation multiple access (SRMA) [Tobagi and Kleinrock, 1976] employs CSMA or ALOHA to determine when to send a request to send, RTS packet over a control channel. While the use of a control channel separates channel access requests from data transmissions, avoiding potential collisions, the control channel results in extra cost and complexity at device level, as well as additional bandwidth requirements.

Floor acquisition multiple access (FAMA) in single-channel wireless networks [Garcia Luna Aceves and Fullmer, 1999] the CTS returned is such in duration to jam any RTS packet sent by other stations, ensuring that all stations in range view the medium as being in use, the 'floor' being granted to a station upon correct receipt of a correctly addressed CTS packet.

MACAW [Bharghavan et al., 1994] extends MACA by the inclusion of a data sending packet, DS to indicate that the RTS-CTS exchange was successful and that the data packet will be sent. An acknowledgment packet is sent to confirm reception of the data packet. MACAW attempts to overcome capture effects which occur as a result of the binary exponential back-off feature by allowing stations to exchange information concerning their back-off counter state.

However the RTS/CTS approach as employed in various forms by MACA, MACAW, FAMA and SRMA cannot be employed in broadcast communication frequently used for the transmission of routing, membership and service discovery inquiries in the wireless domain.

An improvement is the use of collision detection, in the form of carrier sense multiple access collision detection (CSMA/CD) as used by Ethernet [Metcalfe and Boggs, 1976]. Owing to the wired nature of Ethernet the presence of a collision results in a voltage increase on the shared medium, enabling positive detection of a collision, eliminating the need for acknowledgments. However such an approach cannot be successfully employed in the wireless domain. To address this challenge, carrier sense multiple access collision avoidance (CSMA/CA) was proposed, which merges features of both the MACA and CSMA protocols, employing a MACA like RTS-CTS exchange while incorporating CSMA like channel sensing.

MACA and CSMA feature the use of a random back-off function to randomise the interval between medium access attempts. The exact number of retries and the growth function of the back-off delay varying between implementation. The back-off process is predicated on the assumption that transmission failure was a result of another ongoing transmission e.g. a collision. Packets lost due to dynamic error behavour of the wireless medium are viewed as collisions.

Further work combined the features of MACA and CSMA protocols leading to the development of CSMA/CA which incorporates both virtual and physical carrier sensing, allowing the RTS-CTS to be optional. IEEE 802.11 [IEEE, 2005a, Crow et al., 1997] allows specification of thresholds for the use of the RTS-CTS exchange to maxmise throughput of data. For small data payloads the overhead of the RTS-CTS exchange is high, while the probabilty of packet collision is low.

While IEEE 802.11 has been extended to include a quality of service specification in terms of IEEE 802.11e [IEEE, 2005a, Garg et al., 2003], where multiple priority queues are combined with a revised back-off procedure. IEEE 802.11e is still a random access protocol with a random back-off process. Management functions within IEEE 802.11 such as the periodic beacon frame, which are transmitted every 102.4 ms are transmitted in priority to application traffic resulting in transmission delays, the impact of which has been recorded, [Gleeson, 2004].

While random access protocols based on the Aloha, CSMA and MACA families offer simple implementation without need for knowledge of topology or network size, these protocols do not provide users with a deterministic or guaranteed performance [Chlamtac et al., 1997]. Performance is also sensitive to the configuration parameters of back-off algorithm employed and the number of retries permitted.

2.1.2 Agreement Following Contention

While offering low access delays in lightly loaded networks as load increases contention-based protocols may become unstable. A hybrid protocol borrowing on contention based access and on deterministic approaches allows stations to contend for transmission resources by making an optimistic transmission or attempting to set-up a reservation using an adaptation of the RTS-CTS exchange [Lin and Gerla, 1999, Manoj and Siva Ram Murthy, 2002], success of the request being signified by successful transmission of the data.

The aim is to dynamically create a non-overlapping schedule which avoids both the hidden and exposed terminal problem by contending for unused time slots on the medium, either through in band or out of band resource allocation through a contention-based approach.

Sobrinho et al [João L Sobrinho and A. S. Krishnakumar, 1996] propose a solution based on modifying the inter-frame spacing used in CSMA/CA and combining this with a contention resolution solution based on generating a jamming signal of a varying duration. The technique is described with reference to how it would integrate with the normal IEEE 802.11 channel access scheme. By altering the inter-frame spacing traffic is prioitised such that high and low priority, real-time and non real-time traffic are separated. IEEE 802.11, defines four inter-frame spaces, of which only three (SIFS, DIFS & EIFS) are used when running in ad-hoc mode. Sobrinho et al proposes that the unused inter-frame space Point Coordination IFS (PIFS) could be used to prioritise traffic.

Despite this alteration, contention is still non-deterministic, so in addition a contention protocol is used to achieve agreement. The contention protocol relies on stations transmitting a jamming signal. The length of this jamming signal known as a black burst is proportional to the length of time the station has been waiting to transmit. By sending a longer burst than any other station contending at the same time a station may acquire the medium.

The intervals used are critical as to ensure that other stations do not observe the medium idle for sufficiently long as to consider it free, the interval must be less than the slot time and also the minimum inter-frame space, in the context of IEEE 802.11b this interval is 10μsec.

A negative acknowledgement scheme is employed to maximise throughput. However a key oversight is how to handle a lost packet. Retransmission requests are treated are assigned a low priority. The usefulness of any retransmission, in the multimedia domain the traffic is of a firm real-time nature, the loss of number of packets should not lead to failure, but delayed arrival of a packet is of no use to the application. If two stations each transmit pulses of the same duration, while unlikely they will both observe the medium free and given a negative acknowledgement scheme there is no end to end confirmation.

MACA/PR [Lin and Gerla, 1999] adopts a fast reservation setup based on a RTS/CTS exchange to form a primitive TDMA structure based on a fixed duration cycle. Reservation information is piggybacked on data packets, with the receiving station issuing an acknowledgment. Each station maintains a table which records the receive and transmit slots in use. Best effort data may also be supported by listening at the beginning of each slot and identifying an unused slot.

Both Amouris [Amouris, 2001] and Katragadda et al [Katragadda et al., 2003] seek to allocate transmission resources based on the physical location of the station, thus allocating the resources in a deterministic manner. Both divide the environment into fixed sized cells. Amouris divides each cell into an N by N matrix, each entry having a unique transmission time within the TDMA cycle, the time slot used by a station depends on the station instantaneous location.

Katragadda et al [Katragadda et al., 2003] however assumes that each cell will have only a single member and thus sends on this cells frequency and listens on the frequencies of all neighboring cells. An assumption of full duplex communication is made, if two neighboring stations transmit at the same time can both messages be received?

Unlike Katragadda et al, under Amouris's STDMA protocol it is possible for more than one station to share a transmission slot on a round robin basis. When a station moves from one space slot to another, a space slot update message (SSU) is sent. If a station is present in the slot already, it replies with a sorted list of a stations in that space slot including the new station. However how this process is undertaken without interfering with ongoing transmission is not discussed. While both approaches result in deterministic and instantaneous allocation of a transmission slot, a number of deficiencies exist primarily the potentially inefficient usage of the medium if many stations are located in the same physical space slot while the other space slots are unused.

D-TDMA [Tadokoro et al., 2008] is a decentralised TDMA protocol, incorporating a data structure within each frame to describe the slot allocation state. This data structure stores status information for all slots, ACK where a frame was received correctly, RTC request to change in response to a packet collision, NACK, for when a packet was not received and FREE for no packet. This represents a significant improvement over the learning approach required by RTMAC and MACA/PR.

Through reception of messages from other stations, a list of free slots is constructed and one of those empty slots is chosen at random for transmission similar to MACA/PR [Lin and Gerla, 1999] and HCT-MAC [Sobral and Becker, 2008]. Successful transmission results in a reservation until either the station fails or it explicitly releases the slot by sending a 'Request to Change' similar to the [Manoj and Siva Ram Murthy, 2002] ResRelsRTS.

Similar to many other approaches, MACA/PR [Lin and Gerla, 1999] RTMAC [Manoj and Siva Ram Murthy, 2002] D-TDMA [Tadokoro et al., 2008], the Hybrid Contention TDMA - HCT-MAC [Sobral and Becker, 2008] protocol listens to the medium to determine the existing slot allocation state based on overhead transmissions. Again an optimistic approach is adopted, dynamically self organized formed clusters and transmitting in slots deemed to be unallocated.

HCT-MAC is arranged as follows, each TDMA cycle or round contains a known number of super frames, a super frame containing several transmissions. Somewhat similar to D-TDMA [Tadokoro et al., 2008], a snapshot of current allocations is sent by the cluster head in the form of a broadcast known as the start beacon which contains a map structure to indicate which of the following data slots are unreserved. A station wishing to transmit, selects an unreserved slot at random. The super frame ends with a finish beacon from the cluster head with an acknowledge map listing in which data slots the cluster head received a valid transmission.

Under HCT-MAC it is possible for two or more stations to believe they are both allocated the same slot within a super frame due to a failure to consider the impact of radio reception characteristics. Circumstances may exist where two stations transmit in the same time slot, but only one transmission is received by the cluster head, due to differences in wireless signal propagation, the protocol assumes all collisions will occur where two or more transmissions are made in the same slot. At the super frame end the finish beacon frame slot acknowledge map will record the slot in use, which is interpreted by both stations as successful slot allocation, resulting in continuing frame collisions.

RTMAC [Manoj and Siva Ram Murthy, 2002] has been proposed to address some of the issues raised with MACA/PR [Lin and Gerla, 1999] to support both real and non real-time traffic. The initial reservation exchange, ResvRTS-ResvCTS-ResvACK is sent with a higher priority to these compared to there best effort equivalents e.g. RTS-CTS-ACK. A reservation may be released by sending a ResRelRTS to the destination of the reservation which replies with a ResRelCTS. If a reservation goes unused for several cycles it is released.

RTMAC transmits the ResvRTS-ResvCTS-ResvACK within a slot, the data frame only sent in the next occurrence of the slot, if the reservation was successful, in the case of MACA/PR the RTS-CTS-DATA-ACK exchange subsequently only required a DATA-ACK thus leaving a portion of the transmission time previously occupied on a once off basis unused and by inverse prevented reservations occurring as the available slot size while large enough for the DATA-ACK frame exchange could not accommodate the initial RTS-CTS-DATA-ACK exchange.

2.2 Communication Errors

The wireless communication medium is subject to error rates many orders of magnitude greater than those experienced in the wired domain [Nicopolitidis et al., 2004,Eckhardt and Steenkiste, 1998,Nguyen et al., 1996]. Communication reliability varies constantly as a result of changes in the wireless signal propagation paths of signals as a result of mobility of transmitters and/or changes. Errors encountered in most of the real channels are not independent but appear in clusters like in wireless communication channels where fading, shadowing and other impairments exist [Berber, 2003,Swarts and Ferreira, 1999], Channel fading results in bursts of bit errors [Zhu et al., 1991].

As highlighted by Decotignie [Decotignie, 2008], many basic assumptions made in protocol design and evaluation fail to recognise the realities of wireless communication. Communication is neither symmetric nor can communication range be assumed to be a perfect circle.

Background noise on the Wi-Fi spectrum from other devices sharing the same frequency may have a significant impact, Fu et al [Fu et al., 2008] highlights the time varying nature of channel noise, composed of other stations and background noise observed by stationary receivers in an office like environment.

Communication involves implicitly the transmission of information from one point to another through a succession of processes [Haykin, 2001]. Communication encompasses the entire process from the point of generation of the information, to point where the original message is recreated. These processes involve the encoding and conversion of the original message signal into a format suitable for transmission using the chosen physical medium.

The received transmission, the reproduced message is an approximation of the original message, not the actual original message. Therefore each received frame must be verified, to show that what was received was in fact what was sent. Either a simplistic parity schema or an error checking code in the form of a cyclic redundancy check (CRC) [Peterson and Brown, 1961] are required to enable the receiving station to determine with high confidence that the received frame is error free.

If a frame contains an error, the calculated CRC will with a high level of confidence, differ from that contained in the received frame. If a frame is found to contain an error, it must be discarded as it provides no information, as no confidence exists as to any of the frames contents. In the worst case the received frame may have simply been channel noise.

2.2.1 Wireless Signal Propagation Restrictions

A wide range of concerns must be considered when analysing the reliability of packet delivery, not only are there limiting factors such as channel capacity, signal to noise ratio and thermal noise, but also losses incurred during transmission and impacts from the physical world in terms of propagation effects and mobility.

A list of important concerns is presented:

- Bandwidth
- Signal to Noise Ratio
- Thermal Noise
- Path Loss
- Power Link Budget
- Propagation Effects
- Doppler Shift

2.2.1.1 Bandwidth

Shannon's theorem on information capacity [Shannon, 1948] defines the relationship between channel capacity, bandwidth and the signal to noise ratio. Bandwidth is not infinite, it is a commodity, typically regulated by state or international bodies such as the International Telecommunication Union (ITU). Several small pieces of spectrum collectively known as the Industrial Medical and Scientific, ISM band[1] are available for license free use. Bluetooth, Hyperlan, Zigbee and IEEE 802.11 all utilise portions of the ISM band.

The fundamental basis of any digital communication system, is the physical layer where the digital bit stream is converted by encoding into a format suitable for transmission over the chosen communications medium. The information capacity theorem stated by Shannon [Shannon, 1948] provides a fundamental upper bound on channel capacity, presented in equation 2.1.

$$C = B \log_2(1 + SNR) \; bits/s \tag{2.1}$$

The capacity of any communication system expressed in bits per second, is bounded by the available bandwidth, B and is related to the prevailing signal to noise ratio, SNR of the channel. Shannon's equation provides a theoretical bound on channel capacity, C for given B and SNR. The efficiency of the channel with respect to Shannon's bound is termed η which represents a measure of efficiency where the closer to unity η is the more efficient the system is relative to the theoretical maximum. For a given transmission rate R, the efficiency is given by equation 2.2.

$$\eta = \frac{R}{C} \tag{2.2}$$

As η approaches unity, the complexity of the encoding scheme required increases, which in turn increases the potential for the erroneous decoding of the packet.

---

[1] ISM bands are defined by the ITU in Radio Regulations 5.138 and 5.150

2.2.1.2 Signal to Noise Ratio

Signal-to-noise ratio (SNR) is the power ratio between a signal (meaningful information) and the background noise [Haykin, 2001]. Shannon's theorem [Shannon, 1948] provides a relationship between bandwidth capacity and SNR as the data transmitted is reduced below the Shannon limit the bit error rate (BER) approaches zero which represents the ideal case. However the Shannon limit, remains a theoretical limit which cannot be achieved in practice owing to the noise contribution from the receiver and the efficiency of the modulation scheme chosen. Equations 2.3 and 2.4 show the deterimation of the SNR, where either the signal and noise power (P) or the signal and noise amplitude (A) are known.

$$SNR = \frac{P_{signal}}{P_{noise}} = \left(\frac{A_{signal}}{A_{noise}}\right)^2 \tag{2.3}$$

More commonly it is expressed in dB:

$$SNR(db) = 10\log\left(\frac{P_{signal}}{P_{noise}}\right) = 20\log\left(\frac{A_{signal}}{A_{noise}}\right) \tag{2.4}$$

2.2.1.3 Thermal Noise

Thermal noise, is noise emitted from every object in the universe [Nyquist, 1928,Johnson, 1928]. All objects have heat caused by Brownian motion where that object is at a temperature greater than absolute zero, 0 Kelvin / -273 Celsius and as a result emit RF noise which is Gaussian in nature.

$$N = k_b T B \tag{2.5}$$

Where:

k_b Boltzmann's constant $= 1.3803 \times 10^{-23}$ J/K

T Temperature, in Kelvins, 290K is taken as room temperature

B Bandwidth, Hz

2.2.1.4 Path Loss

The received power at the receiver reduces with the square of the distance, this occurs as the transmitted wave spreads out, which can be thought of as an ever expanding sphere. The signal is further attenuated if it passes through walls. Ali-Rantala et al [Ali-Rantala et al., 2003] presents a review of the attenuation effects of various common building materials. Path loss in free space [Parsons, 2000] is presented in equation 2.6, where λ is the wavelength and d the distance from the transmitter source and L_{fs} the path loss.

$$L_{fs} = 20 \log \left(\frac{4\pi d}{\lambda} \right) \tag{2.6}$$

2.2.1.5 Power Link Budget

The relationship between the probability of error and the (E_b/N_0) ratio is known as a waterfall curve, operating point one refers the upper limit set on error however to ensure reliable operation of the link in light changing conditions a significant margin must be built in, this second operating point refers to the $(E_b/N_0)_{rec}$ for this lower probability of error, clearly the margin can be made arbitrarily large by increasing $(E_b/N_0)_{rec}$ but this comes at the cost of an increased transmission power requirement with knock on effects on power consumption. However the ISM bands impose power output limits restricting such solutions. The margin can be defined in dB as:

$$M(dB) = \left(\frac{E_b}{N_0} \right)_{rec} (dB) - \left(\frac{E_b}{N_0} \right)_{req} (dB) \tag{2.7}$$

Employing Erceg et al's [Erceg et al., 1999] expanded path model shown in equation 2.8, the received power for a given fade margin is given by equation 2.9 as shown by Proakis [Proakis, 2001]. Andersen et al [Andersen et al., 1995] presents path loss exponents (n) for various building types.

$$L_p = A + 10nLog_{10}(d/d_0) + s \tag{2.8}$$

Where A is the standard path loss as given previously in equation 2.6, where $d_0 = 100m$ and s is a zero-mean Gaussian variable to represent shadow fading.

$$P_{rx} = P_{tx} + G_{tx} + G_{rx} - Lc_{tx} - Lc_{rx} - L_p - FadeMargin \tag{2.9}$$

Where G_{rx}, G_{tx} are the gains of the receiver and transmitter antennas, Lc_{tx}, Lc_{rx} are the losses due to the cable joining the antennas to the transmitter and receiver respectively while L_p is a measure of the path loss for the distance between the receiver and transmitter.

Thus the required power at the transmitter, P_{tx} must be greater than or equal to the required power, P_{rx} at the receiver less the losses, owing to transmission effects such as path loss, while still allowing an adequate margin to allow for the dynamic effects such as fading.

2.2.1.6 Propagation Effects

Faced with no line of sight or even with line of sight radio waves reflect, scatter and diffract when faced with obstacles, buildings, vehicles, lamp posts, rain and so on. In doing so the signal strength and quality varies dramatically making reliable signal reception difficult particularly in a mobile urban or industrial environment where a high density of such objects exist.

Practical experiments conducted by Willig et al [Willig et al., 2002] show a correlation between movement in the environment and successful frame reception. Stepanov et al [Stepanov and Rothermel, 2008] employ ray tracing to determine the received signal state accounting for reflection and diffraction. However this approach does not account for mobility within the environment.

Three distinct phenomena lead to degradation of the transmitted signal at the receiving station:

- **Reflection**: When a wave comes across a very large object compared to its wavelength it reflects off the surface. Typically this class of behaviour would be seen in the presence of walls, buildings and the earth itself.

- **Diffraction**: Secondary waves are formed by the incident wave striking a sharp edge, e.g. corner of a building this leads to a effect where the wave is seen to 'bend' around obstacles.
- **Scattering**: When a wave comes in contact with small objects with respect to the wavelength of the signal, signs, lamp posts and similar this leads to the generation of wavelets.

2.2.1.7 Doppler Shift

Where either the transmitter, receiver or both are in motion a frequency distortion effect is present, Doppler shift. Clearly if the mobile host is moving the path length to the transmitter is changing. This results in a minor change in the frequency as seen by the receiver, frequency increases as the path length reduces. This may be expressed simply as:

$$\Delta f = \frac{\nu}{\lambda} cos\theta \tag{2.10}$$

The frequency change, Δf is related to the velocity $\nu cos\theta$, where θ represents the angle with respect to the source and the wavelength λ. Zhou et al [Zhou et al., 2005] analyses the performance of IEEE 802.11b in various scenarios, showing reductions in throughput as speed increases.

2.2.1.8 Summary of Propagation Effects

Owing to a combination of diverse factors which are beyond the control of the designer, electromagnetic waves are subject to a very unpredictable environment, which result in challenging propagation conditions, particularly in an industrial or urban environment where a high density of obstacles exist, where many of the hosts and obstacles are potentially mobile.

The energy in the electromagnetic wave arrives at the receiver having travelled by more than one path, each path is a different length thus each signal is slightly out of phase compared to a imaginary signal which travelled by line of sight. Signals sum together and based on the phases of the signals, either constructive or destructive interference occurs and results in the creation of a composite signal whose strength varies dynamically as the mobile host moves. As the individual path lengths vary, so does the signal received at the mobile host.

Given the significant variance in propagation caused by environment in terms of building construction as discussed by [Andersen et al., 1995, Ali-Rantala et al., 2003] combined with the impact of mobility resulting in signal path changes resulting in dynamically varying channel conditions, stations must dynamically adapt however little data is available upon which to basis such adaptation, limited to signal to noise ratio and bit error rate statistics obtained from the wireless physical layer.

2.2.2 Time Varying Channel Conditions

The quality of wireless links can change significantly and unpredictably over time and space [Eckhardt and Steenkiste, 1996]. Errors are burst like in nature due to attenuation, fading or interference radiating sources [Eckhardt and Steenkiste, 1998]. In the presence of burst errors, immediate attempts to retransmit a failed packet will result in a another failure [Willig et al., 2002]. CSMA style channel access protocols attempt to resend the packet immediately or after a slight pause based on the exponential back-off process, FIFO ordering of the packet queue is maintained. Typically an IEEE 802.11 transceiver will attempt seven retries before dropping the packet. In doing so the queue of packets for transmission is blocked leading to jitter combined with slow and variable transfer rates resulting in poor use of the channel capacity.

Work presented by [Bhagwat et al., 1997, Balasubramanian et al., 2006] suggest a reactionary protocol which faced with burst errors to dynamically alter the transmission schedule in an effort to maximise throughput of packets. Each link may be considered statistically independent, since in a mobile wireless multi user environment propagation effects seriously impact on the quality of the signal received thus the signal received at any two stations is likely to be different even if they are in close proximity. Multi-path fading may result in zero signal at one location only for a different station to have a good quality signal a matter of meters away thus the definition that each point to point link is independent is reasonable. To overcome a bursty error situation the transmission queue may be reorded dynamically producing a new schedule at runtime [Balasubramanian et al., 2006] or implement a queuing strategy based on service requirements [Bhagwat et al., 1997].

As a TDMA channel access system is employed, contention and collisions can no longer be considered as causes of packet loss. When a packet loss is detected, its loss must be attributed to propagation conditions of the medium.

2.2.2.1 Characterisation of the Bursty Channel

Basic telecommunication theory is developed on the assumption that noise on communication channels is additive white Gaussian noise [Haykin, 2001,Berber, 2003], bit errors are independent and the channel has no memory of previous errors. The consequence of this assumption is that there is no statistical dependence between succeeding transmitted bits [Berber, 2003].

However Kanal et al [Kanal and Sastry, 1978] states that the channel has memory, due to the impact of time varying propagation effects and suggests the use of Markov chains to model the wireless network. A Markov chain determines the state transition probabilities, based on the current state of the channel. The channel therefore has memory and the success or failure of a transmission is dependent on prior attempt outcomes. Error bursts are infrequent occurrences, a bursty channel is constant and reliable with relatively little bad bursts [Wang et al., 2007].

The use of a Markov chain to model the wireless channel and in particular the two state Markov chain model known as the Gilbert and Elliot model [Elliott, 1963] is widely accepted as modeling the wireless channel with sufficient accuracy. Others propose the use of more complex analysis with large numbers of states such as the N state Fritchman model [Fritchman, 1967], or Aldridge et al [Aldridge and Ghanbari, 1995] who adopt a three state model.

2.2.2.2 Gilbert/Elliot Channel Model

The Gilbert/Elliot Channel Model [Elliott, 1963] is a simple two state model, in which the channel is either in a good or bad state, where the bad state indicates a packet upon reception will have one or more erroneous bits. Probabilities classify the probability of moving from one state to another given the state of the previous packet e.g. probability of transition to bad given last packet was good, P_{gb}. Immediately before each packet is sent the channel state is determined.

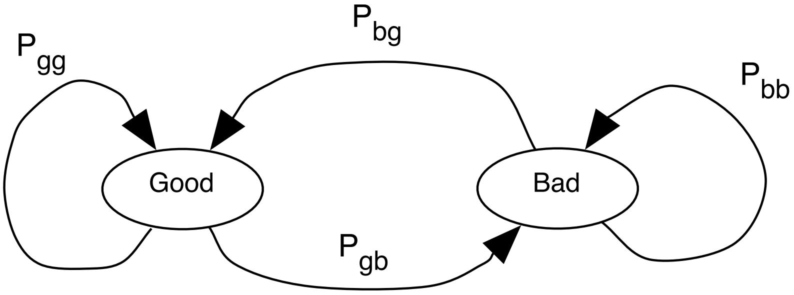

Fig. 2.3: Gilbert/Elliot Channel Model

This simplistic but useful model hides some of the detail such as the number of erroneous bits and the point at which they occur within a frame, Willig et al [Willig et al., 2002]. Figure 2.3 shows the state transition diagram of the Gilbert/Elliot Channel Model, the probabilities of transitions of the model can be described as follows:

P_{gg} Probability given last packet was not effected by error that the next packet received will also be error free, good good.

P_{bb} Given the last packet was effected by error, probability that the channel will remain erroneous, bad bad.

P_{gb} Probability given last packet was not effected by error that the next packet received will be, good to bad.

P_{bg} Probability given last packet was by effected error that the next packet received will not, bad to good.

Numerous works, Jiao et al [Jiao et al., 2002], Wang et al [Wang et al., 2007] Zorzi et al [Zorzi, 1998] and others, have expanded on the works of [Fritchman, 1967] and [Elliott, 1963]. Errors occur in bursts, errors occur due to phenomena related to dynamic real-world events, channel fading due to stations and other objects moving causing the signal strength to vary.

Aldridge et al [Aldridge and Ghanbari, 1995] identifies the need to model single bit errors in a different manner to which burst errors are modeled. A three state Markov model is proposed, this being a simple Fritchman model with error states for both the single bit error case as well as the burst error condition. A Poisson distribution is suggested to model the bursty channel condition, which for a low average λ e.g. 1 or less has a long tailed distribution which has been shown to represent the real world experience as noted by others [Otani et al., 1981, Willig et al., 2002]. An analysis is presented which requires knowledge of the average overall mean bit error rate, the mean gap between errors and the mean error burst length.

Communication is a two way process, in the case of acknowledged transmissions the loss of the data packet or the acknowledgment are treated the same as a result of lack of global knowledge, Wang et al [Wang et al., 2007] discusses the issues which arise when a transmission requires handshaking e.g. the 4 way DCF handshake in IEEE 802.11, using a two state Markov model to represent channel conditions. While a detailed analysis is presented the values chosen to represent the channel probabilities and burst lengths are not linked to any real world experience.

2.2.2.3 Channel State Dependent Scheduling

Ci et al [Ci and Sharif, 2000] employ a simple adaptive approach to overcome packet loss by halving the fragmentation threshold upon failure to receive an acknowledgment frame. While this reduces the probability of a single bit error by sending smaller packets. The process does not overcome bursty errors, which are transient effects possibly effecting several packet transmissions. While the papers title suggests it addresses the bursty channel issue the approach adopted appears to target an environment with a high bit error rate.

Willig et al [Willig et al., 2002] performs a low level evaluation of packet loss in an industrial context. A very useful outcome is the validation of the non independence of transmissions under failure conditions, for the conditional probability of a packet loss given the previous packet was lost, $P_r[packet_{n+1}\ lost|packet_n\ lost]$ results in value of 0.7179, which highlights that the use of instant retransmission of failed packets is suboptimal. The converse probability, $P_r[packet_{n+1}\ success|packet_n\ success]$ results in a value of 0.9804.

Willig [Willig et al., 2002] proposes that a better understanding of the statistical properties of the errors encountered at the bit level can provide a better MAC layer. Due to the detailed analysis of the errors encountered at the physical layer, the MAC layer can be conditioned to react appropriately to channel conditions, since the presence of certain error types or patterns informs us of conditions that allow appropriate reaction - as to opposed to the typical single assumption of a collision.

Bhagwat [Bhagwat et al., 1997] argues that the typical approach of multiple retries is a poor solution when faced with bursty channel error conditions, since further transmissions to that station will continue to fail in the short term. This results in the blocking of other transmissions and in poor utilisation of the bandwidth. Transmissions to other hosts are not effected, by virtue of the statistical independence of the communications link to individual hosts. A transmission queue can be implemented for each host and if a transmission fails, that queue can be marked as error prone for an interval, and packets from other queues can be sent instead. As a result, head of queue blocking is avoided.

Balasubramanian et al [Balasubramanian et al., 2006] propose a combination of three techniques to provide a real-time communication system with high reliability. A TDMA structure is combined with the ability for nodes to exchange groups of slots as to allow a host to make a retransmission attempt in the future to allow for any busrty error condition to subside and the ability to retransmit small percentages of packets if needed. This results in a robust approach with several levels of recovery, however, the schedule is managed by a master node. Stations submit requests through a contention-based access approach to obtain slots. Some flexibility and performance is lost, since a precondition exists that all messages will be available at the start of the TDMA cycle and have a deadline that expires after the end of the cycle to eliminate scheduling concerns and allow the slot exchange between stations to be undertaken without real-time deadline concerns.

2.3 Real-Time Communication

While the works presented in section 2.1 focused on throughput and fairness offering best effort data communication services, the needs of real-time communication, differ significantly. The three key concerns are presented:

- Reliability
- End to end delay
- Low jitter

Traditional approaches to communication favour maximisation of throughput or minimisation of average message delay. Real-time communication focuses on the real-time constraints of individual messages [Malcolm and Zhao, 1995]. The correctness of the systems depending on both the logical results and the time at which those results appears [Kopetz, 1997]. A communication system which provides timely message delivery is therefore an essential component of a distributed real-time system.

Each message has unique set of characteristics, a deadline, priority and a reliability requirement. As each message is unique, each message is considered separately to ensure each messages requirements are met [Hamann et al., 2007]. Special cases may allow for more flexible approaches. Redundancy may exist in the data, which may allow selective dropping of frames e.g. MPEG video [Zou et al., 2003] without resulting in a failure or the transmission schedule is known in advance or in a closed system where the packets to be transmitted are known of in advance enabling an optimised schedule to be prepared off-line.

2.3.1 Medium Access Control

Controller-area network (CAN) bus [Robert Bosch GmbH, 1991] implements a priority based addressing scheme to reach a deterministic outcome for each contention event. This approach titled, Carrier Sense Multiple Access/Bitwise Arbitration (CSMA/BA) relies on synchronisation and the use of a wired medium allowing one station to dominate that medium.

While CAN employs priority based arbitration [Davis et al., 2007] that priority is encoded in its address, and as each station has a unique address, arbitration by station address will always result in a deterministic outcome, the station with the highest priority. Thus a single global priority based queue is said to exist [Davis et al., 2007].

Binary 1 and 0 are described as being recessive and dominant respectively. If one or more stations transmit a recessive address bit and a single station transmits a dominant address bit, all stations will see a 0, as the medium is tied to ground overriding the recessive binary 1. The stations transmitting the recessive bit will know another station with a higher priority is trying to access the medium and will cease its attempt to transmitt. This process continues until the address transmission is completed and one station remains.

CSMA/DCR [Prodromides and Sanders, 1993] retains the CSMA/CD slotted structure, while introducing bounded and deterministic resolution of collisions. Following each transmission collision, each host on the network is allocated a slot in the deterministic collision resolution (DCR) phase. Each station transmits in its allocated slot within the DCR phase, if a station was not involved in the initial collisions it sends no data, after the elapse of a slot time, the remaining scheduled transmissions are moved up to reduce the overall time thus speeding up the resolution. While the outcome is a deterministic resolution in bounded time, each station must know the total number of peers and the schedule to adopt upon a collision, similar to CAN a wired broadcast medium is required.

The time-triggered protocol (TTP) proposed by Koptez et al [Kopetz and Grunsteidl, 1993, Kopetz, 1997] is a TDMA channel access method with statically assigned slots, supporting both a highly distributed fault tolerant service, including both communication redundancy and replication known as TTP/C and a low cost, non fault tolerant version, TTP/A which can be built from generic components.

Each TTP network consists of a number of fault tolerant units (FTU), each unit potentially containing several hosts and communication controllers to provide redundancy. A key feature of TTP/C is that of the redundancy management layer which is provided to manage the cold start procedure, replacement and reintegration and switching in a shadow node following a failure.

The use of multiple channels is provided as a means to overcome a communications link failure. However the application of multiple channels may not provide redundancy in the wireless domain. Channels of a similar frequency share similar propagation effects. If the path is obstructed in some way communication may not be possible despite the use of multiple channels.

2.4 Existing Protocols

The CSMA, MACA and Aloha family and their adaptations are not deterministic by design, contention and collisions are random as is the response in terms of random back-off process used to resolve contention and collisions. While section 2.3.1 descibed a number of deterministic approaches, these were applicable to the wired domain only. The restrictions of TDMA are well known, under high load TDMA is the optimal MAC protocol to be used [Gobriel et al., 2008]. Under low load, the fixed slot structure results in efficient use of bandwidth but with variable transmission jitter owing to the need to transmit in preallocated slots. Naturally scalability is restricted by the number of TDMA slots [Decotignie, 2008].

In an attempt to overcome these restrictions several protocols merge both TDMA and CSMA characteristics in an attempt to combine the strengths of TDMA and CSMA while minimising the impact of their weaknesses, Z-MAC Zebra-MAC [Rhee et al., 2005]. Z-MAC employs a two tiered structure, the slot owner is given a smaller initial contention window than non owners, allowing it priority access to the medium, if the slot owner has nothing to transmit the slot is sensed as idle and CSMA access is employed by other stations. In the case of clock synchronisation errors, the fact each host always employs carrier sense regardless of being slot owner prevents transmissions colliding.

Z-MAC overlays a TDMA like structure, which gives priority to scheduled channel access but falls back on CSMA operation in order to provide resiliency against clock synchronisation errors between stations. A slow cold start sequence is followed which includes neighbour discovery, slot assignment, local frame exchange and global time synchronisation. The discovery process, which lasts 30 seconds results in each node becoming aware of its two hop neighbours. CSMA/CA channel access applies, however the owner of each slot has a higher access priority through the manipulation of the initial contention window. If the slot owner has no data to transmit, other stations will observe the medium as idle and will contend for access. Similarly should a timing failure occur medium access reverts to CSMA like access.

Developing a TDMA slot schedule to allow maximum concurrent transmission and channel reuse is an NP hard problem [Ramanathan, 1999]. Each station contains unique schedule as the peers within its two hop neighbourhood are potentially different. The goal is to produce a non-conflicting schedule, such that two or more stations transmitting at the same time, on the same frequency do not interfere with each other as they are sufficiently geographically far apart.

Ramanathan [Ramanathan, 1999] assumes an ideal environment being both noisless and immobile. In a realistic mobile environment, the two hop neighbours of each host are in a constant state of change, requiring constant revision of the slot allocation table. Given the unreliable nature of wireless communication, determining all two hop neighbours, developing a globally valid schedule and delivery of such a schedule represents a significant challenge, with the dual challenges of schedule generation and schedule propagation.

The use of a cellular environment avoids the NP hard problem, as the cellular structure and frequency assignment ensures no interference between transmissions in adjacent cells. A centralised or decentralised slot allocation process at cell level can be applied to allocate slots within the cell. As the typical cell has six sides, a static allocation of slots may also be made to allow inter-cell communication. As a station moves it would require to perform a handover between cells.

RT-MAC proposed by Baldwin et al [Baldwin et al., 1999] focuses on achieving deadlines at the cost of reliability by dropping packets as soon it is known that their deadline cannot be met. As a single FIFO queue is used, no priority ordering is implemented, unlike the multi queue scheme adopted by IEEE 802.11e [Garg et al., 2003,IEEE, 2005a]. The focus is on discarding a packet as soon as it is known that its deadline cannot be achieved, thereby releasing resources to enable the real-time deadlines of other packets to be achieved. Firstly when a packet is taken from the queue its deadline is checked, after the back-off timer expires, the deadline is checked again. The process repeats if retransmissions are required. If at any point the deadline has expired the packet is dropped.

Similar to MACAW [Bharghavan et al., 1994] the back-off counter value is exchanged between stations in order to provide an enhanced collision avoidance scheme through knowledge of when adjacent stations will transmit.

2.4.1 Prioritisation of Traffic

If a station it wishes to gain access to the medium under CSMA/CA protocol it must first sense the medium for a short period. If a station was to vary the length of this sensing interval in proportion to the urgency of its traffic, it results in priorities medium access, real-time stations still have to contend with each other. The result is a segregation of high priority from low priority traffic, however in a network where all traffic is high priority this offers no advantage in the area of contention resolution.

Nelson et al [Nelson and Kleinrock, 1985] proposed a collision free multi-hop channel access protocol based on TDMA incorporating spatial reuse. The protocol uses a compatibility matrix which contains an entry for each directed arc a 1 indicates that both i and j arcs can be in simultaneous use. Spatial reuse is maximised by the use of directed arcs so the arc between stations A to B is treated differently to B to A thus the exposed terminal problem is eliminated.

While promising, the requirement for the network to be static and the position of each station known to calculate the adjacency matrix to determine which arcs of the network can contain non-interfering transmissions is a significant restriction. While in a sensor network such a layout may be possible given a fixed network the inherent flexibility obtained from a wireless network is lost. In such an application a wired solution would offer a much higher bandwidth, which in turn would significantly reduce queuing delays.

Zhu et al [Zhu et al., 1991] proposed a decentralised channel oriented media access control protocol (DCAP) based on a TDMA frame structure for use in highly dynamic vehicular environments. Handovers between adjacent clusters as a result of the mobility of station are also discussed.

Each station maintains a bit-map of which channels are in use, a station logically or's each bit-map it receives to create a local bit-map which represents the free channels which do not conflict with any of their neighbours.

Knowledge of the slot allocations in use is provided by each host maintaining a vector quantity known as a slot bit-map a 1 indicating the slot is in use or channel noise was sufficiently high to consider it in use and 0 where the slot is unused. Each slot begins with the slot bit-map, by xor'ing the local bit-map with the bit-map received, thus generating a view of the network around that host and minimises collisions by not transmitting in slots it knows to be in use.

The channel access approach is similar to those employed in [Lin and Gerla, 1999, Baldwin et al., 1999, Manoj and Siva Ram Murthy, 2002] though opportunistic use of transmission resources which have been determined to be unused. To achieve detection of which channels are unused a significant assumption is required, that hardware exists to continuously monitor all channels and store information as to the received signal quality of those channels.

Flexible TDMA (FTDMA) [Willig, 1997] addresses a number of issues by providing MAC level acknowledgments as well as permitting the use of unallocated slots on a random access basis. A common shared channel is divided into slots of equal length to form a TDMA system. Each slot contains exactly one data transmission and corresponding acknowledgment. As the number of transmission slots is finite a bound exists on the number of real-time clients that each picocell can support. A complex set of interactions takes place between real-time clients and a central base station. Stations are polled in order and respond in order after the polling transmission, a reservations stage commences in which stations request or release transmission slots.

A dedicated slot exists to allow new stations to register with the base station such that they will be polled in subsequent cycles. Unused slots may be used by non real-time data, using a random access protocol such as Aloha. This forms a robust approach featuring MAC level acknowledgments and dedicated resources for registration, however the polling and registration phase is complex and time consuming. The schedule is transmitted to all stations, similar to [Huber and Elmenreich, 2004].

Caccamo et al [Caccamo et al., 2002] and Cunningham et al, [Cunningham and Cahill, 2002, Cunningham, 2003] adopt a slotted TDMA approach assigning particular slots or groups of consecutive slots to specific stations. In doing so the end to end delay of messages is known. Caccamo el al [Caccamo et al., 2002] focus on the formation of a schedule based on modifying EDF whereas Cunningham et al focus on the issue of achieving consensus in both real-time and in the presence of failures leaving scheduling to a higher layer. Caccamo proposes a method to reuse slots a station fails to use, however this supports only aperiodic messages with soft deadlines, where in many cases an aperiodic message is a panic message.

2.4.2 TBMAC

The Time Bound Medium Access Control Protocol as proposed by Cunningham et al [Cunningham and Cahill, 2002,Cunningham, 2003] and implemented by Gleeson [Gleeson, 2004] and Hughes [Hughes, 2007] proposes a fault tolerant membership service in order to implement a dynamic TDMA cycle membership service. The implementation of a time bounded fault tolerant membership service enables time bounded access to the wireless medium.

To reduce the probability of collisions during the transmissions of mobile stations, the geographical coverage area occupied by the mobile stations is statically divided into a number of geographical cells. Each cell is allocated a distinct radio channel to use, maximizing the total overall bandwidth available in the ad-hoc network. It is further assumed that a higher level layer exists above TBMAC to continually monitor the current position of the mobile station and to notify the TBMAC protocol of a change of cell, and thus a change of the radio channel to use.

TBMAC implements a TDMA approach where each TDMA cycle is divided into a contention free (CFP) and contention period (CP). In order to allow stations to join, e.g. to obtain a transmission slot a portion of the TDMA cycle is reserved to explicitly accommodate requests from new mobile stations to join the cycle. Both the CFP and the CP periods are divided into slots of a well-known duration. Dividing access to the wireless medium into these two well-known time periods requires synchronization amongst all clocks of the mobile stations in a cell and furthermore, in the ad-hoc network.

While a station may have time bounded access to the communications medium, no reliability guarantees are offered. The TBMAC protocol provides no support for acknowledgments or failure detection at MAC level instead relying on an unspecified upper layer to implement end-to-end reliability.

TBMAC overcomes many of the restrictions found in HCT-MAC proposed by Sobral et al [Sobral and Becker, 2008] as clusters are formed based on proximity, as a cellular approach is adopted. Slot allocation is distributed and fault tolerant with the MAC address of each slot owner known.

2.4.2.1 Contention Free Period (CFP)

During this portion of the cycle contention and collisions are prevented as each mobile station transmits only in the specific slot(s) which they have been allocated. To ensure slot management atomic broadcast messages are propagated and also to ensure timely detection of mobile station failure(s) a mobile station will always transmit a packet in each slot it is currently allocated. In the absence of data the message contents are left empty, such a packet is termed a ‘null packet’ where as a packet with a data payload is termed a ‘data packet’.

2.4.2.2 Contention Period (CP)

For a mobile station to gain its first CFP slot, it must first contend with other mobile stations seeking a CFP slot. A portion of the TDMA cycle known as the contention period is used exclusively for this purpose. Mobile stations which currently has no slots allocated to them must transmit a request, a CP request in at least one CP slot to initiate the slot allocation procedure.

2.4.3 Wireless TTP

Huber et al. [Huber and Elmenreich, 2004] describe an approach to use wireless communication to facilitate communication between a number of possibly mobile field bus clusters through the use of time-triggered paradigm. However the implementation is highly inflexible as it is based on modifying an existing wired real-time protocol to operate in the wireless domain. As a result of its origin as a wired protocol the round structure must be programmed in advance, while it is possible for the master host to choose from a number of different round structures by indicating the structure to be used at the start of each round in the round indicator inside the master frame, it is not possible to alter these predetermined round structures in an on line fashion.

The protocol is unable to match the real-time performance and characteristics of TTP, in particular the use of multiple communication channels, at best the communication redundancy offered is equivalent to TTP/A. A master is required to maintain clock synchronisation between all members. A comprehensive fault hypothesis based primarily on failures not continuing for longer than the maximum synchronisation internal is presented. If a station which become desynchronised it will cease transmission.

2.4.4 Hyperlan2

Hyperlan2 [Rappaport, 2001, Johnsson, 1999, Doufexi et al., 2002] adopts an approach similar to that which underlies GSM [Rahnema, 1993] by providing a connection oriented service to ensure the bandwidth, delay and jitter requirements of the traffic are satisfied. A centralised scheduler is employed, with a dynamic TDMA channel access method. Additional reliability is provided through link adaptation. A convergence layer is provided as a middleware to integrate Hyperlan2 with existing network protocols such as ATM [Minzer, 1989].

2.4.5 Bluetooth

Bluetooth [IEEE, 2005b, Tanenbaum, 2002], is a centralized Time Division Multiplex (TDM) system, where a designated master controls access to the wireless medium and determines which host may transmit within a fixed 625μsec time slot. Similarly, in IEEE 802.16 (WiMax) [Tanenbaum, 2002], a base station (BS), controls access to the wireless medium for both downstream (BS to host) and upstream (host to BS) communication, where medium access is divided into time slots or related to QoS specifications, respectively. Both of these protocols achieve predictable medium access latency by adopting a centralized scheduling solution. However, given the potential for a highly dynamic network topology and mobility of participating hosts a centralized solution is not suitable for the ad-hoc domain.

2.4.6 IEEE 802.15.4

IEEE 802.15.4 and its related ZigBee application stack [IEEE, 2003, Egan, 2005] implements a low power wireless personal area network (WPAN). To meet the requirements of cost and energy saving three classes of Zigbee device exist.

Each network contains exactly one personal area network (PAN) coordinator, which stores details of the network configuration and enables communication between other networks. A network may contain several ZigBee routers which act as intermediaries allowing devices to exchange data, IEEE 802.15.4 defines such devices as being a full-function device, FFD. The ZigBee end device is the most basic device known as a reduced-function device or RFD, which has the minimum of resources required to communicate with either a router or coordinator with no support for routing.

Each cycle consists of a beacon and a super frame containing the active period, the inactive period. The super frame is further divided to include both a contention access period (CAP) and a contention free period (CFP), not dissimilar to TBMAC [Cunningham and Cahill, 2002]. Channel access during the CAP is based on CSMA/CA.

Within the active period of a IEEE 802.15-4 cycle, following the CAP a maximum of 15 guaranteed time slots (GTS) [Koubaa et al., 2006] are available in the contention free period (CFP). These slots offer a means for a station to obtain guaranteed contention free medium access at least once during each beacon interval, however the CFP is optional and only activated through a request to the PAN. However, only seven stations may hold a GTS slot at any time, though each station may hold several GTS slots. While [Koubaa et al., 2006] presents an upper bound on data transmission delay in a GTS slot. A node must make an explicit request to obtain a GTS slot and cannot share that slot with other stations, similar to other works [Gobriel et al., 2008] Koubaa et al [Koubaa et al., 2006] mechanisms are proposed to allow stations to share slots based on run time requirements in order to make effective use of the GTS slot time. However, the i-GAME approach proposed by Koubaa et al relies on the PAN to provide admissions control services. A concern given IEEE 802.15-4 support for reduced-function devices, a scenario exists where no other station within the network has the ability to act as a coordinator should the current one fail or become uncontactable.

2.5 Requirements for a Real-Time MAC

Drawing on the review of medium access control protocols and recognising the potentially negative and varying impacts of signal propagation, this section presents a list of functional requirements to describe a firm real-time fault tolerant wireless communication sub system.

- **Decentralised**: The protocol must operate in a distributed manner and continue to operate in the presence of station failures.

- **Dynamic Membership**: Stations must be able to join and leave the networking in a real-time fashion. A key factor is the time from power up until a station can transmit its first application data packet.

- **Transmission Delay**: A stable and known a priori delay is essential to support real-time traffic. This delay should be bounded regardless of the volume of traffic on the medium.

- **Admissions Control**: The scheduler must implement admissions control. Such is essential to ensure existing stations service is not compromised by the scheduling of further transmissions. The request to gain admission should have real-time properties.

- **Flexibility**: The protocol must not be application specific, and must not rely on the specific traits of a particular application, e.g. MPEG video [Zou et al., 2003]. The protocol must acknowledge that different applications require different levels of real-time support and have different data transmission requirements. The protocol must support a wide range of configurations.

- **Timing**: All packets admitted will be transmitted and received within the defined real-time requirements provided that the expected rate of packet loss on the medium does not exceed the predicted value.

- **Reliability**: The level of reliability achieved by the communication system must satisfy or exceed the requirements of the data transmitted. Account must be taken of the unreliable wireless medium to insure packets are received correctly within the defined deadline. Retransmissions may be required.

- **Prioritisation**: Messages should be transmitted in priority order while maintaining real-time deadlines as to insure under worst case conditions the most urgent messages are prioritised for transmission.

- **Performance**: The protocol must maximise its utilisation of the available bandwidth. Performance should not degrade under extremes of load and the defined level of reliability should be maintained at all times for admitted packets.

- **Adaption**: Channel conditions may vary dynamically, the protocol must be able to adapt to conditions while still maintaining service in compliance of the real-time needs of all admitted traffic.

2.6 Summary of Related Work

This chapter introduced and discussed both medium access control and the nature of communication errors in the wireless communication domain. Several wireless medium access control protocols where introduced to provide an overview of the problem domain. An overview was provided of various aspects of the wireless medium and propagation effects which may negatively impact on the reliability of communication.

In the context of these propagation effects a number of wireless real-time medium access protocols where discussed with reference to existing wired real-time approaches. Implementing real-time communication in the wireless domain represent a significant challenge owing to the considerable unknowns, in terms of propagation characteristics as result of both mobility of both sender and/or receiver as well as objects in the environment.

A trade-off exists between the flexibility offered by unscheduled random access self resolving protocols such as the CSMA and MACA family and the strict slotted approach offered by TDMA. Under low load conditions CSMA and MACA offer low access delays and address errors through a random back-off approach which may counteract the issue of burst channel errors, by introducing a delay between transmission attempts. However schedulability, priority and deterministic behavour can only be obtained by adopting coordinated scheduled approaches such as TDMA, which provide timeliness in exchange for a lack of flexibility to adapt to channel conditions and to provide support for retransmissions.

A number of protocols where discussed in this chapter where flexible scheduling was combined with a TDMA slot allocation structure, which permitted dynamic swapping of slots between stations to overcome communications failures, however significant scheduling restrictions and communication overheads where identified.

The chapter concluded with a list of desirable properties which should be supported by a real-time wireless communications system. Many of these requirements demand precise timing and coordination between stations, features which are not available within a large portion of classic wireless medium access control protocols such as the MACA and CSMA family approaches.

Chapter 3

Hierarchal Distributed TDMA Protocol

This chapter describes the Hierarchal Distributed Time Division Medium Access Control (HD-TDMA) protocol. The goal of HD-TDMA is to sustain time-bounded communication in a distributed wireless system in the face of frame loss as a result of variable wireless propagation conditions. The provision of a hierarchal division of slot allocations permits the use of local scheduling to dynamically reorder and retransmit frames in the face of frame errors and loss.

The HD-TDMA medium access control protocol is supported by two independent modules: a dedicated membership-management module which manages the allocation of transmission slots and an admissions control and retransmission module which manages the order in which packets are transmitted within each slot. The probabilistic admissions control and retransmission protocol are discussed in detail later in chapter 4.

The remainder of this chapter is structured as follows: section 3.1 introduces HD-TDMA by providing a brief overview of the protocol. Section 3.2 details the medium access control approach used by HD-TDMA and section 3.3 describes the slot states. Section 3.4 describes the frame types and layout and section 3.5 describes the communication modes that are supported. Finally, in section 3.6 the distributed membership service is described.

3.1 Overview

HD-TDMA presents a new approach to the problem of wireless medium access control. Distinct from traditional TDMA approaches, multiple frames may be transmitted within each TDMA slot. This overcomes the restrictions and efficiency issues inherent to fixed frame sizes. Furthermore, the HD-TDMA MAC provides support to exchange membership and control information to facilitate a dynamic membership service.

Figure 3.1 shows the basic layout of the communication architecture, consisting of the MAC layer upon which the membership service and probabilistic admissions control and retransmission protocol sit, the layer which sits above these services implements an external API to allow an application to interact with HD-TDMA.

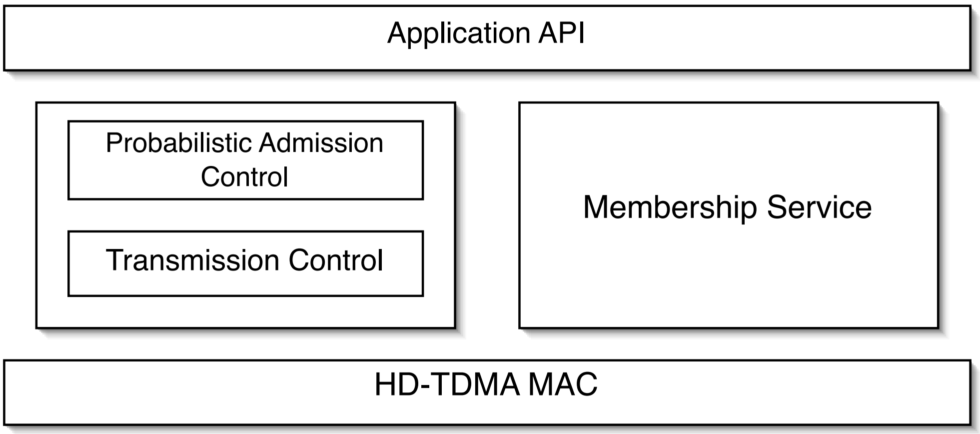

Fig. 3.1: High Level View of System

HD-TDMA module provides the medium access control interface, managing the sending and receiving process. The HD-TDMA module also provides the time base for all other modules. Each HD-TDMA slot consists of a beacon-frame containing membership and control information followed by a number of data frames provided by the Queue module.

The Queue module includes both the admissions control and the transmission control processes, relying on the membership module for knowledge of current slot allocations and proposed changes.

The membership service manages the allocation of slots which are utilised by the HD-TDMA module for transmission and by the queue module to queue frames for transmission. The process ensures that time-bounded fault-tolerant slot allocation and de-allocation occurs, to ensure each station has a consistent view of slot allocations.

3.1.1 Assumptions

In order for HD-TDMA to operate and to meet its goals a number of restrictions must be applied to the environment and the capabilities of the participants. These restrictions relate to the clock synchronisation between nodes which impacts on both TDMA slot timing and membership services and the need for location awareness to implement a cellular structure.

3.1.1.1 Clock Synchronisation

In order to implement a TDMA-like transmission slot structure, tight clock synchronisation is required between all participating stations. Stations which are passive, which do not have allocated to them a transmission slot, do not require clock synchronisation as they respond only to polling frames sent by other stations with slots.

Having knowledge of time and cycle duration a station can determine the start time of the next slot and next cycle and thus is synchronised with the TDMA cycle. Synchronisation accuracy of the order of μs is envisaged, the accuracy of synchronisation impacts on the duration of the inter-slot space. Tighter synchronisation will allow the inter-slot gap to be reduced and result in more efficient use of the available bandwidth.

3.1.1.2 Cellular Structure

The topology of the network that is considered consists of a cellular structure in which a station is member of a cell. Cell membership is maintained by a membership service and is dynamic in order to account for both mobility and failure of stations. All stations within a cell are within communication range of each other.

The choice of a cellular structure is motivated to provide a structured and constrained environment to maximise the reliability of communication. The cellular approach ensures all stations are within communication range of each other, thus avoiding the hidden terminal problem [Tobagi and Kleinrock, 1975]. Spatial reuse is maximised as each cell has a unique frequency which does not conflict with adjacent cells.

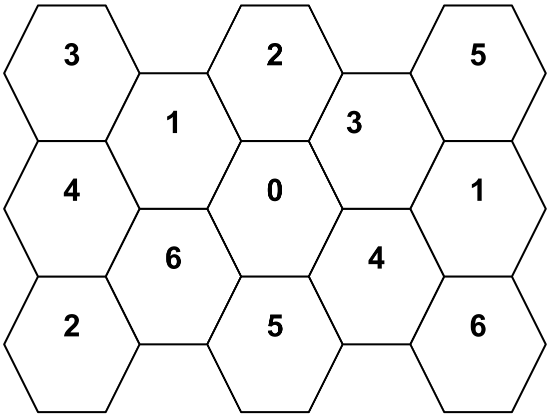

Fig. 3.2: Cellular Layout

Figure 3.2 presents the common seven-cell cluster cellular layout. The cell width is a function of the physical layer technology in use and could conceivably be tens or hundreds of metres to several kilometres in diameter. As detailed by Parsons [Parsons, 2000] each cell may be hexagon shaped and be assigned a unique frequency. For the common, seven cell cluster, each cell has six direct neighbours, each with a unique frequency, the distance between the closest edges of two cells sharing the same frequency is greater than twice the cell width. This significantly reduces the potential for interference from transmissions in cells sharing the same frequency or code and avoids the exposed terminal problem [Tobagi and Kleinrock, 1975].

3.1.1.3 Frame Loss

The HD-TDMA protocol is designed to tolerate the loss of frames. As a TDMA channel access scheme is utilised for all real-time data traffic, it is assumed that where frames are lost or damaged, this is as a result of failure of a station under fail-silent conditions or due to dynamic wireless propagation impacts resulting in frame corruption.

3.2 HD-TDMA MAC

The design of the HD-TDMA medium access approach is motivated by three goals that have been derived from available research: 1. the approach needs to avoid competition for the medium and provide time-bounded access to the medium. 2. the approach needs to be flexible, such that sporadic messages and retransmissions may be accommodated. 3. the medium access approach must recognise the non-ideal characteristics of the wireless medium in terms of variable propagation and packet delivery reliability. These goals are the basis for the adoption of a TDMA-style approach which incorporates flexibility in order to support localised autonomous scheduling decisions.

3.2.1 Structure of a HD-TDMA Cycle

HD-TDMA consists of a TDMA channel access method in which each transmitting station is allocated at least one slot, within which it can make a number of transmissions. The slot allocation process occurs in band, but is separate from data transmissions. Slot allocation is a function of the membership service discussed later in section 3.6. The use of TDMA with allocated slots eliminates contention between data transmissions, eliminates the hidden terminal problem and thus is able to provide time-bounded medium access.

Each slot in HD-TDMA begins with a beacon transmission to indicate the start of the slot. This enables other stations to synchronise with slot boundaries. The beacon is sent as a broadcast, acting as an announcement of the fact that the slot is in use by a station and contains information concerning the allocation of all slots in the system as perceived by the transmitting station as well as information used in the allocation and de-allocation of slots. Thus reception of this beacon fully informs all stations as to the status of the full HD-TDMA cycle.

Within each slot, a number of separate transmissions may occur. The exact order of transmissions is decided by the local station according to a local scheduling strategy as will be detailed later in chapter 4. This grants considerable autonomy, allowing each station to dynamically adapt the schedule without the burden of achieving agreement with other stations. Within the slot, frames can be of differing sizes as determined suitable by the application requirements of the sending station.

Previous work [Hughes et al., 2006, Gleeson, 2004] involved the implementation of a TDMA-based MAC layer known as TBMAC [Cunningham and Cahill, 2002] in Real-Time Linux using commonly available IEEE 802.11b wireless cards. This implementation was the basis of a real-world evaluation [Gleeson et al., 2009]. The outcome of this work highlighted a number of practical issues involved in implementing a TDMA-like MAC in a resource constrained environment. Latency in sending packets as well as the practical difficulties in achieving synchronisation between mobile stations were the most obvious of these. This experience motivated the design of the HD-TDMA protocol, two of the results of which are the flexible slot structure and acknowledged data service.

The basic structure of the HD-TDMA MAC protocol consists of a TDMA slotted structure. Each slot itself may contain an arbitrary number of transmissions from the same station thus facilitating variable packet sizes and the ability to implement acknowledged and unacknowledged packet transmission.

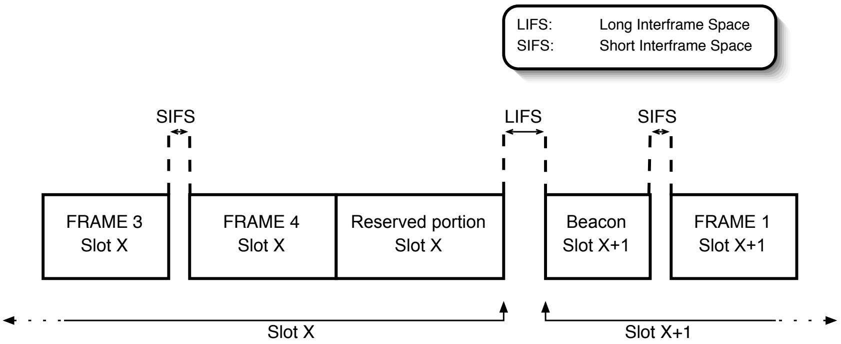

Fig. 3.3: Inter-Frame Gaps

The technical motivation in the design of HD-TDMA is to minimise the amount of idle time on the medium as doing so will maximise the potential throughput. In random access protocols such as MACAW [Bharghavan et al., 1994] and CSMA [Tanenbaum, 2002] gaps are required between each packet sent. The duration of these gaps must be sufficient to ensure that other stations did not detect the medium as being idle during a communication exchange ongoing between other stations, but long enough to allow for timing variations and propagation delays.

For example, IEEE 802.11 offers a fragment burst mode [IEEE, 1999]. A station may make several back-to-back transmissions in this mode without the need to wait for a complete distributed inter-frame space (DIFS) interval, by only waiting the short inter-frame space (SIFS), giving it uninterrupted access to the medium. In a distributed TDMA system, the length of an inter-frame space is directly related to the quality of the clock synchronisation that can be achieved [Huber and Elmenreich, 2004, O'Connor, 2003]. It is desirable to reduce the number of instances where a full inter-frame gap is required in order to minimise the idle time of the medium. In order to achieve this the transmissions from the same station are grouped together. As figure 3.3 shows in HD-TDMA a long inter-frame space (LIFS) is only used between slot boundaries since only there does the clock synchronisation issue arise as all transmissions within a slot originate from the same station.

It is common to group transmissions together [Balasubramanian et al., 2006], as this enables more flexible scheduling and recovery. Having grouped transmissions together from each station, a station of course could have several groups of transmissions. The approach is formalised by introducing a slotted structure with each group of transmissions residing within one or more slots.

The local station is delegated full responsibility for transmissions within these slots and thus can add/remove/reorder/retransmit frames within each slot as it deems appropriate without reference to its peers. This enables stations to send sporadic messages and retransmissions without first achieving consensus and thus makes it possible to achieve the time-bounds required. This flexibility will later be exploited by the probabilistic admissions control and retransmission control protocol as presented in chapter 4.

3.2.2 Slots

Typically TDMA systems have fixed slot sizes. This may prove inefficient where the payload varies introducing a considerable waste of bandwidth where packets are much smaller than the allocated time for there transmission. HD-TDMA permits variable packet sizes up to a defined maximum transmission unit (MTU) within each slot, thus minimising overhead while maximising throughput. Both unicast packets with or without acknowledgement, broadcast and polled communication are supported by HD-TDMA. In unicast a station can send a packet and receive an acknowledgement similar to the MACA [Karn, 1990] DATA-ACK thus allowing positive confirmation of reception.

Figure 3.4 illustrates a possible configuration of a HD-TDMA cycle with an N slot TDMA cycle, the contents of one slot having been expanded. The slot begins with a mandatory beacon-frame containing membership information. Subsequently a number of frames of varying sizes are sent followed by a reserved space managed by the admissions control process. A portion of the slot is reserved to allow for retransmissions and to accommodate sporadic messages.

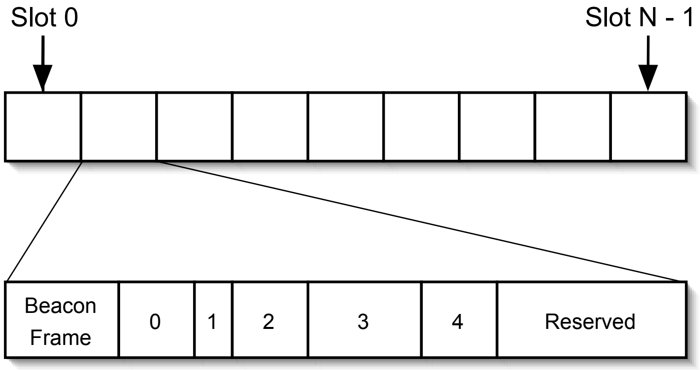

Fig. 3.4: Relationship Between The Slot, Its Contents And The TDMA cycle

3.2.3 Beacon Frame

The start of each slot consists of a beacon-frame sent as a broadcast. This contains sufficient information to uniquely identify the slot and its transmitter. The beacon-frame, will thus as a minimum include the slot number and the owners identity. Owing to the need to support a dynamic membership, as well as quick detection of failures, each beacon contains a snapshot of the allocated state of each slot. The beacon messages provide a heart-beat to indicate that a station is on-line and information about the organisation of the TDMA cycle.

The beacon also contains an indicator to inform other stations as to whether the rest of the slot is idle. If the rest of the slot is indicated as idle, other stations may transmit sporadic messages in the remainder of the slot using a CSMA channel access approach. However, no guarantees can be provided to the reliability of transmissions sent in this manner.

3.2.4 Data Transmission Rates

A long standing goal and evaluation metric for medium access control protocols has been that of maximising throughput of application data. While such a goal is compatible with best-effort data communication the provision of real-time data services demands both timeliness and reliability in priority to throughput.

Probability of error in wireless communication increases rapidly [Willig et al., 2002, Zhou et al., 2005] as the data rate is increased. As the transmission coding scheme becomes increasingly complex to transmit at the increased data rate the space between encoded symbols reduces resulting in a higher probability of erroneous decoding.

As a result HD-TDMA adopts a flexible approach allowing data to be transmitted at up to eight different data rates, all stations must support a basic transmission rate which is the lowest data rate. All control frames and frames carrying cell membership information are transmitted at the basic rate. This includes all management type frames: beacon frames, acknowledgments, both positive and negative, as well as polling requests.

In the case of data type frames, a three bit field in the management portion of the frame encodes the transmission rate of the data which follows. This field `tx_speed` is always present. Thus, the transmission speed of data frames may be varied as a mechanism to improve reliability and as a result the timeliness of message delivery.

3.2.5 Failure Detection

In any real-time system it is critical that the failure of a station is identified promptly and that suitable steps are taken to resolve the failure. To facilitate this, HD-TDMA begins each slot with a beacon-frame transmission. Thus any station that does not receive a beacon-frame in a slot in which it knows to be in use infers that the communications link between itself and that station may have failed. It may be the case that either the station has failed silently or its transmissions can no longer be received. If a number of beacons are not received a query phase is initiated and through consensus the membership service can then decide if the station has failed e.g. failure to receive the beacon at one station may be due to a bursty channel error where the other station received the beacon correctly.

3.2.5.1 DATA-ACK

An acknowledgement of a data frame, or DATA-ACK, is transmitted by the receiving station if the data has been received correctly. This provides positive feedback, assuming that the data frame is received correctly by the station and that a ACK frame is sent back, leaving three separate points of failure-frame, failure outbound, station failure at receiver and frame failure return.

Any frame which is received and fails to pass the CRC test is dropped, thus it is not possible to determine if the transmission failured occur to the DATA or ACK portion.

3.2.5.2 DATA-NEGACK

An extension of the DATA-ACK approach, by providing feedback as to a failed transmission only partially received. The frame design of HD-TDMA has been chosen to permit the operation of negative acknowledgments, not on the basis of expected traffic but on the partial correct reception of a frame.

Negative acknowledgments, are possible by providing a separate CRC field for both the management data and the application data, when the management portion of a frame is received correctly, but the data portion fails its CRC check a negative acknowledgment is sent. Typically the data portion will be substantially larger than the management data thus is more likely to suffer from a single-bit error. Detection of such errors and the use of negative acknowledgments provides a source of channel information.

3.3 Possible Slot States

In order to make the most effective use of slots, each station must be aware of the state of each TDMA slot regardless of its current ownership. In order to achieve this each beacon-frame propagates its view of the current slot configuration.

Three possible slot states exist under the HD-TDMA protocol- allocated, unallocated and unknown. Within the allocated state two sub-states exist. Each station views the state of allocated slots locally. Either the slot is owned by this station, ‘owner’ or is owned by another station, ‘other’. Further to the owner or other state, a slot may be idle or in use as indicated by the beacon-frame. Figure 3.5 presents an overview of the relationship between states.

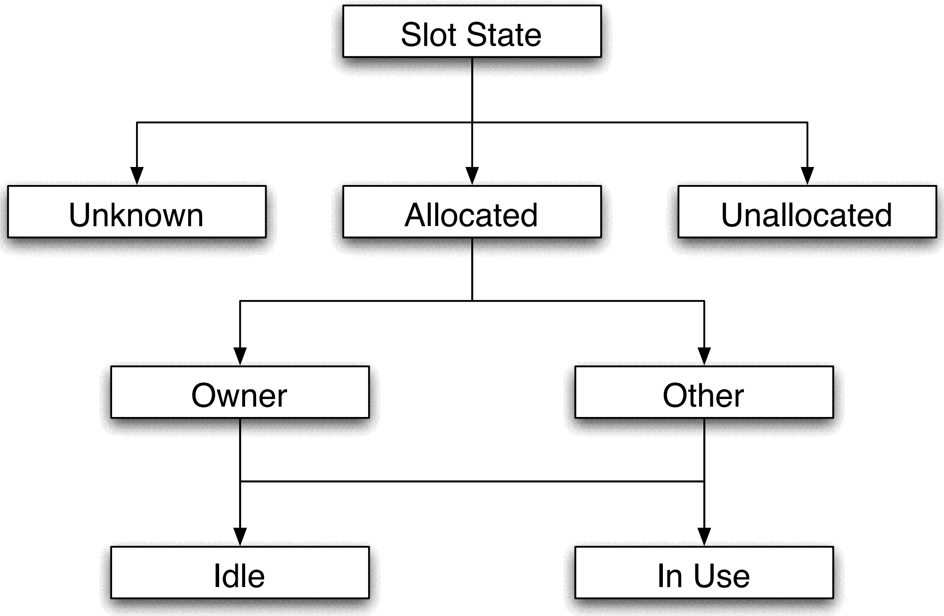

Fig. 3.5: View Of Slot States

- **Unallocated**: The default state of a HD-TDMA slot is unallocated. When a slot is in the unallocated state it cannot be used for any mode of data transmission. In this state, slots marked as unallocated are available for use by the membership-management service as a dedicated slot in which to broadcast a request for an initial slot allocation, doing so without impacting on ongoing real-time data transmissions.
- **Allocated**: When a slot is allocated it is no longer available to the membership service. The station allocated the slot will transmit a beacon-frame at the begin-

ning of the slot. Two sub states exist with respect to the view of ownership, owner and other.

- **Other**: Stations within the same cell which are not the allocated owner of a slot will mark the slot as allocated to another station, for short, ‘other’.
- **Owner**: The station which is the allocated owner of a slot will mark the slot as ‘owner’, as each slot can have only one owner, only one station within a cell will mark each slot as owner.
- **Idle**: The idle state is a subset of the allocated state. HD-TDMA slots which are allocated consist of a mandatory beacon-frame followed by zero or more real-time data transmissions. A special case exists where no data transmissions are scheduled for the slot. If this is the case, the idle flag is set in the beacon-frame. Stations receiving a beacon-frame with the idle flag set may use the remainder of the slot to transmit sporadic messages and best-effort data using a CSMA channel access method.
- **In Use**: The in-use state is the inverse of the idle state, where the owner of the slot has indicated in that slots beacon-frame that the station has real-time data queued for transmission in the slot.
- **Unknown**: Each station within its beacon-frame broadcasts its view of the slot allocation. If a station has received differing views of the allocation state of a slot, a failure has occurred within the membership-management process to allocate slots in a consistent manner. This is most likely to occur where a partition occurs within a cell or during the initial empty cell bootstrapping case.

  The slot is therefore marked as unknown and cannot be used for data or membership transmissions of any sort. The membership-management service may then resolve the conflict when an unknown state arises to recover the slot for normal use.

3.4 Frame Layout

The basic frame format in HD-TDMA is shown in figure 3.6, consisting of five distinct sections. The PHY component is specific to the physical layer employed by the underlying wireless technology, typically the PHY consists of a known sequence of bits to enable the receiving station to synchronise to the data stream that follows.

The management portion contains necessary control information to identify the type of frame this transmission contains, as well as the frames origin and destination if applicable. A CRC field is provided so the validity of the management data may be confirmed independently from the data portion.

The data field and related CRC are optional and the contents of the data field may be application data or control information as specified by information within the management frame.

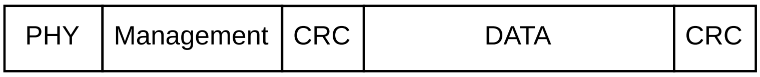

Fig. 3.6: Frame Layout

3.4.1 Beacon Frame

Each HD-TDMA slot commences with a beacon-frame of fixed length which contains control information concerning the TDMA cycle. Presented in figure 3.7, the additional information portion of the beacon-frame contains information from the membership service. The slot bit-map provides a quick overview of the slot allocation state of each slot.

To support clock and cycle synchronisation each beacon-frame contains a time stamp and a slot ID is provided within the beacon-frame. From time to time a station may have no real-time data scheduled for transmission, the idle flag within the frame will indicate this state.

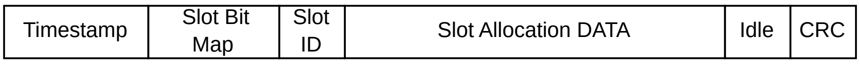

Timestamp	Slot Bit Map	Slot ID	Slot Allocation DATA	Idle	CRC

Fig. 3.7: Beacon Frame layout

3.4.2 Management Frame

Each HD-TDMA frame contains a management sub-frame which describes the purpose of the frame. Each management frame provides a source and destination address in the form of an IEEE 48bit MAC address, a frame type field and for transmissions including a data sub-frame the tx speed and length fields are populated with the parameters of the data sub-frame.

Tx Speed	Frame Type	Frame Seq	SRC MAC	DEST MAC	Length	CRC

Fig. 3.8: Management Frame Layout

An initial slot allocation request, data acknowledgment, negative data acknowledgment and polling request are all examples of where no data frame is required, the management sub-frame is sufficient in these cases.

3.4.3 Data Frame

In some cases, the management frame is insufficient and a data sub-frame is required. The management frame indicates the frame type, its length and the transmission speed. Two basic variants of data frame exist, broadcast and acknowledged. In the case of a broadcast frame the destination address will be set to all bits 1 in the management frame.

Fig. 3.9: DATA Frame Layout

3.5 Communication Modes

Different applications have differing data transmission requirements, to this end HD-TDMA supports six distinct communication modes. It is possible to mix the use of these modes within any slot to provide maximum flexibility. Transmission modes include unicast, broadcast, acknowledged and polled together with support for sporadic and best-effort data messages.

3.5.1 Unicast

Unicast point-to-point communication where no acknowledgments are provided is supported. The management frame specifies the destination address. The data sub-frame may be sent at a higher data rate as specified by the tx speed field in the management sub-frame. No real-time delivery guarantees are provided.

3.5.2 Broadcast

In order to support group communication to all stations within the cell, a broadcast mode is provided by HD-TDMA. The frame is sent to the Ethernet broadcast address, with the data portion sent at the basic transmission rate to ensure maximum probability of successful transmission by all cell members. No real-time delivery guarantees are provided.

3.5.3 Acknowledged

As described in section 3.4.3 each data packet in HD-TDMA is composed of two independent portions, the management sub-frame and the data sub-frame. As both sub-frames have independent CRC checks, it is possible to distinguish between a frame which is corrupt from one where the management portion is not corrupt but the data portion is corrupt.

If both the management and data sub-frames are received correctly, both CRC checks will pass and the complete frame is valid. In this case a management frame is generated by the receiver of type positive acknowledgment. If however the CRC check of the data sub-frame fails a negative acknowledgment is created.

Real-time support is provided for acknowledged frames through the use of a probabilistic admissions control and related retransmission protocol which will be discussed in chapter 4.

3.5.4 Sporadic Messages

Sporadic messages are supported through the ability of stations to make local decisions concerning the transmission schedule of each slot. This empowers stations to reorder transmissions, subject to the flexibility in the message deadlines and available recovery space at the end of the frame in order to facilitate transmission of sporadic messages.

If all scheduled transmissions have successfully completed and sufficient time remains in the slot to transmit sporadic messages, such messages may be sent. This occurs where all real-time frames have successfully been retransmitted. Sporadic messages may also be sent in slots marked as idle in that slots beacon-frame transmission. Pending sporadic messages have priority over best-effort data on the transmitting station.

3.5.5 Best Effort Data

Where no real-time data is allocated for transmission in a slot or where all the real-time data scheduled for transmission in a slot has been transmitted successfully, best-effort data may be transmitted.

Best effort data will only be sent in data slots marked as idle. Priority is given to any sporadic messages.

3.5.6 Polled

The number of transmission slots in a TDMA-based system is fixed, resulting in issues of scalability. The number of potential transmitting stations cannot exceed the number of slots for a given duration. Clearly it is not possible, nor is it desirable, to allocate transmission resources to all possible transmitting stations in advance.

To overcome this constraint HD-TDMA supports a polled transmission mode. A special management frame type allows a station to poll other stations to see if they have data they wish to transmit, this allows stations which do not have transmission slots to transmit data.

To poll a station, the polling initiator generates a management frame of type polling request. The destination mac address set to the intended target. The length field specifies the maximum data size with the transmission speed field indicating the transmission speed to respond at.

Upon transmission the polling initiator starts a timer, set to expire after sufficient time has elapsed to allow for the transmission time of the management frame, the duration of the response assuming it to be at the maximum length specified by the management frame and two short inter-frame spaces.

Upon reception of a polling frame the station being polled may:

- Ignore the polling request as it has no data to send.
- Respond with a data frame not greater in size than specified by the polling request.
- Issue a negative acknowledgment frame to indicate that no data was available to send.

Upon receiving a negative acknowledgment the polling initiator expires its timer and can continue with other transmissions without waiting the full duration of the worst-case polling transaction time, thus enabling it to make efficient use of the limited transmission time.

3.6 Membership Protocol

The critical requirement imposed by the slotted TDMA structure is the need to ensure that all stations have a consistent view of the slot allocations. This is critical to guarantee contention-free transmissions in each slot.

In order to achieve this, Cristian's synchronous atomic broadcast protocol [Cristian, 1990] has been implemented, to provide a slot management function in a totally-ordered fashion. While this work adopts a similar approach to that used by TBMAC [Cunningham, 2003, Cunningham and Cahill, 2002], the implementation differs significantly.

Unlike TBMAC, as described previously in section 2.4.2, HD-TDMA does not require the dedicated TBMAC Contention Period (CP) instead unallocated slots are used to transmitt slot allocation requests by stations with no currently allocated slots. Through the elimination of the contention period, two performance improvements may be obtained. Firstly an increase in throughput by increasing the number of slots available for application data and secondly reduced worst-case time-bounds for membership-management operations.

The remainder of the discussion of the membership protocol is structured as follows: first a review of how membership data is propagation between stations in section 3.6.1, secondly how membership data is propagation between stations is discussed in section 3.6.2, while the membership operations provided are listed in section 3.6.3. Several membership modes are supported, these are discussed in section 3.6.5; in conclusion time-bounds both worst and best-case are derived in section 3.6.6.

3.6.1 Exchanging Membership Data

Provision to exchange membership data is made for stations with and without existing slot allocations. As this provision is made at the MAC level, the process is generic and is not tied to any particular agreement protocol, as it merely provides a means to exchange membership data in such a way as to be independent from real-time data traffic.

Two specific cases exist for the exchange of membership data, for stations with no existing slot allocation (section 3.6.1.1) and for stations with at least one existing slot allocation (section 3.6.1.2). The first method, in turn, relies upon the second to ensure propagation of the data to all cell members.

3.6.1.1 New Stations

Access to the medium is split between two modes, HD-TDMA slots which have been allocated to a specific stations and those which have yet to be allocated. As discussed in chapter 2, the disadvantages of the TBMAC protocol with respect to membership-management were identified, principally the overhead of the contention period (CP) used to handle membership-management requests.

This contention-period is a fixed portion of the TBMAC cycle used exclusively by stations which do not have a slot allocated to them to request a slot. When all slots are allocated any request issued in the CP area will be ignored, as no slot can be allocated, thus the CP period is of no use and reduces the overall bandwidth available for application data. In addition as the CP period is distinct from the application data transmission slots, introducing a gap between adjacent CFP slots resulting in potential jitter in the delivery time of real-time data.

As dedicated resources are required to ensure that membership requests can be issued, a virtual contention period (VCP) is therefore introduced, where a station determines that a HD-TDMA slot has not been allocated, that slot is used as a contention period. The number of contention periods will therefore be equal to the number of unallocated slots.

Within each VCP the requesting station will transmit its membership request several times using a CSMA channel access protocol, this approach increases the probability of successful reception of the membership request by at least one active station and also allows for the case where two or more stations issue membership requests in the same VCP slot.

As each allocated HD-TDMA slot begins with a beacon-frame, which contains information sufficient to identify slots which are vacant, correct reception of a single beacon-frame is sufficient to provide the current slot allocation state of all slots in this cell, together with a list of pending changes contained within the additional information portion of the frame. Thus a station can determine which slots are unallocated and membership operations are kept separate from all application data both real-time and best-effort.

If all HD-TDMA slots are allocated, no virtual CP slots will exist. This is acceptable as with no slots available no allocation will be possible, thus no bandwidth is reserved maximising the bandwidth available for real-time application data.

3.6.1.2 Existing Stations

A station with a pre-existing slot allocation does not rely on the virtual contention period slots to exchange membership-management requests. Membership requests are conveyed inside the beacon-frame structure which contains a dedicated portion to convey a configurable number of membership updates. As the beacon-frame is not subject to contention and is sent at the basic transmission rate it has the highest probability of successful reception by other stations.

Stations with existing slot allocations on reception of a beacon-frame containing membership data, the slot bitmap and proposed slot allocations, the atomic broadcast messages, will merge that information with their local view of membership-management requests and populate their beacon frames accordingly.

3.6.2 Agreement Protocol

The agreement protocol provides a method to ensure that each station within each cell has a consistent view of all slot allocations. Given dynamic membership, mobility and risk of communications failure in the wireless domain [Nicopolitidis et al., 2004,Eckhardt and Steenkiste, 1998], a distributed solution for achieving and maintaining agreement over slot allocations has been chosen.

The agreement protocol must address two distinct challenges: firstly, when a station wishes to change its slot allocation state, either requesting allocation of a slot or to deallocate one or all its currently allocated slots. This proposed membership state change must be propagated reliably to all stations within the cell, including those who do not have slots allocated to them; secondly, is the coordinated application of the proposed membership changes by all stations within the cell, such that each station performs the update at approximately the same time and same manner to ensure consistency.

3.6.2.1 Overview of Atomic Broadcast

The goal of the atomic broadcast approach is to ensure that all members of the cell receive at least one copy of each membership-management request in advance of the delivery deadline specified within that request. To ensure reliability this process is decentralised.

By ensuring that all stations have a consistent view of the cell membership and of membership-management requests each station will make the same decisions. Thus the decision to all ocate a slot will be made in an atomic fashion all stations allocate the same slot or no slot is allocated.

The atomic broadcast approach described by Cristian [Cristian, 1990] has been adapted. The key challenge is to ensure each member of the cell receives at least one copy of each membership-management request before its delivery time. In order to ensure the propagation of the membership-management request to all cell members, a delivery delay is introduced, this delay is termed the agreement time and is referred to as Δ.

The agreement time Δ which is set to 2 * cycle times, this ensures that the requesting station is able to transmit the request in each beacon-frame for each slot that station is currently allocated, and therefore will be transmitted at least twice before the delivery time. In simple terms, a station which wishes to change the replicated data proposes an action to update the replicated data at a certain time in the future, the request is propagated to all stations within the cell. Each station, having a replicated list of membership-management requests, will perform each update at the same time as all other cell members.

Total ordering of membership updates is maintained. Should more than one update be proposed for the same time, the order of updates is considered. Updates are ordered by time then by the requesting stations mac address and then by sequence number of the individual request. The goal is to maintain consistent replicas on all stations while minimising the communication overhead while tolerating a level of message-loss which is overcome by the flooding distribution approach. All pending membership updates will be applied at the start of each slot before the beacon-frame is sent/received, thus avoiding inconsistencies arising.

3.6.3 Membership Management Functions

Membership management operations are performed by a set of five functions:

- Slot allocation request
- Slot de-allocation request
- Deallocate all slots to leave cell
- Intercell
- Suspect Slot Failure

3.6.3.1 Slot Allocation

By default, when a station enters or powers on in a cell, it will seek to have allocated to itself at least one slot. Further slots may be requested by the application layer. The station creates an atomic broadcast message of type `allocate slot` containing its mac address and a sequence number which is queued for transmission.

Figure 3.10 presents the sequence of operations which occur during a slot allocation. First, the membership request is created and placed in the local membership update queue. During each beacon-frame transmission the local membership queue is copied into the membership update portion of the beacon-frame and is transmitted to all other cell members which integrate that list with their own local list and include that list with their own beacon transmissions. When the deadline of the membership request arrives, typically two cycle times from its time of creation, the membership request is applied by all members of the cell.

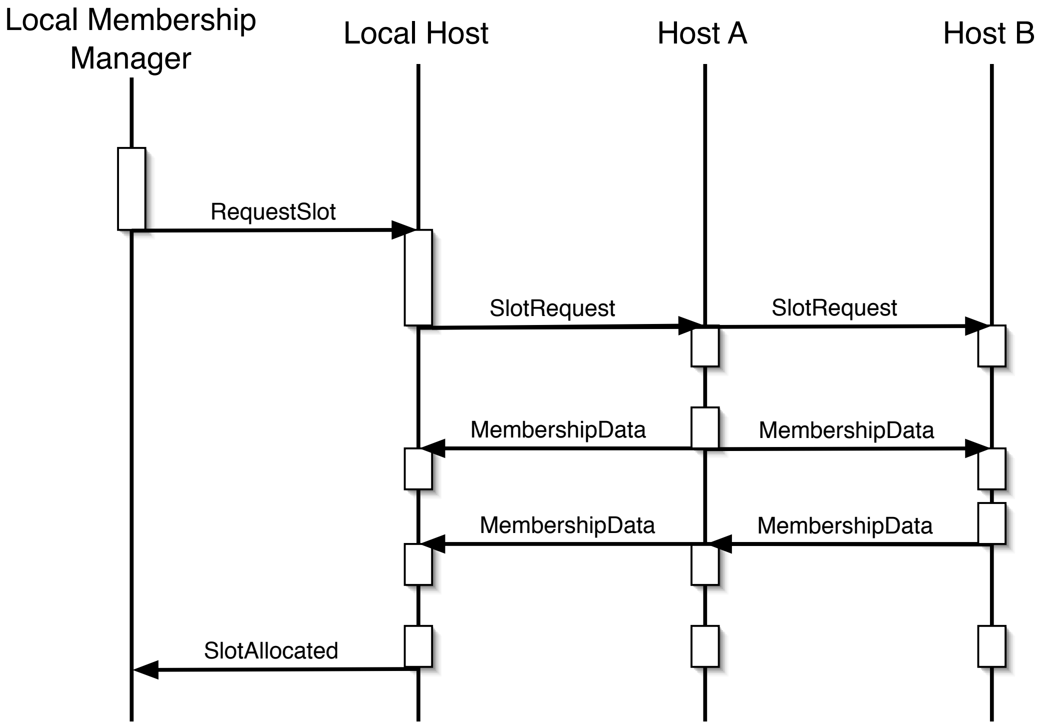

Fig. 3.10: Handling A Slot Allocation

3.6.3.2 Slot De-allocation

Slot de-allocation function is accessed by both the application layer to release unused resources and by the mobility/routing layer which manages the transition between cells, in which case all slots allocated to a station need to be deallocated as a group.

To de-allocate a single slot a station creates an atomic broadcast message specifying its mac address, the slot number to be deallocated and the request type as de-allocate. This message is then queued and time stamped with a delivery time of Δ cycles in the future when it is transmitted for the first time in this station's beacon-frame slot allocation data structure.

Where it is required to de-allocate all slots, an atomic broadcast message is generated specifying the mac address of the station with the request type of de-allocate all. This is then queued as described above for transmission.

3.6.4 Missing Beacon Frames

Loss of a beacon frame potentially indicates that a station has failed, or it may be a result of a transient propagation event. Failure of a station to receive a beacon frame will not impact on the reception of any data frames sent. Failure to receive a beacon will not result in the detection of a slot as being idle as the detection of an idle slot requires the successful reception of a beacon frame with the idle flag set.

If a station counts a number of failures to receive a beacon frame, a query process is initiated through the generation of membership management request to all cell members identifying the effected as station as potentially failed. Through consensus slots allocated to the effected station are either released or retained. Given the variables in propagation it maybe that some stations periodically cannot communicate with other stations, yet those stations remain fully operational, this approach avoids unnecessary and unwarranted slot deallocations.

3.6.5 Membership Modes

Reviewing the environment in which HD-TDMA must operate, three distinct membership scenarios exist. The most complex case occurs when a station enters an empty cell, either on its own or at approximately the same time as other stations. The other two cases, occur when at least one active station exists already in the cell. A station with no current slots will need to request a slot as will a station with a pre-existing slot allocation seeking further transmission resources.

Each case is distinct in how the request to allocate or de-allocate the slot is made, however a significant commonality exists in the process followed to allocate the slot once the request has been made.

3.6.5.1 Bootstrapping

The bootstrapping case occurs when a station powers on in a cell which is has no currently active station, the empty cell case. Motivated by the approach used in TBMAC [Cunningham and Cahill, 2002] which employs an epidemic approach to achieving consensus, a modified approach is employed. The TBMAC implementation relied on aligning bootstrapping process to cycle starts, this is an arbitrary choice as the cycle start has no significance except to form a boundary between the CP and CFP TBMAC periods as well as equipping stations with the knowledge that the bootstrap process would always complete at the cycle start. All stations which have joined the cell in the preceding cycle then enter the bootstrap phase and contend equally. Stations which power up and detect ongoing bootstrapping defer until after the bootstrap phase is complete and use the non-empty cell mechanism to seek a slot in due course.

At worst-case, a wait of two full cycle times exists before beginning the bootstrap process, when in fact a station need only listen for one full cycle time to determine if the empty cell scenario applies. HD-TDMA instead aligns the bootstrapping function at the slot time level, thus a station needs to wait no more than one cycle and one slot time before entering the bootstrapping phase reducing the worst-case delay by almost a full slot time.

If a station in the bootstrapping phase, despite not transmitting in any of the slots it has chosen yet, receives a message from another station also bootstrapping it aborts its bootstrapping process and waits for the other station(s) in bootstrap to achieve consensus before requesting a slot through the non-empty cell mechanism.

The goal is to minimise the number of stations within the bootstrap phase in order to maximise the probability that the epidemic process within the bootstrapping protocol will result in a consistent membership view on all stations. A detailed discussion and evaluation of the epidemic bootstrapping has previously been presented by Cunningham et al [Cunningham, 2003].

Each beacon-frame contains a bitmap which details the stations local view of the slot allocation, four possible states are supported, OWNER, OTHER, COLLISION and AVAILABLE. In order to achieve a consistent view of the slots in use, each station must merge its local view of the slot allocation state with those views of other stations as they are received.

3.6.5.2 Empty Cell Scenario

During the bootstrapping phase, if a station detects the absence of an expected beacon-frame during a bootstrapping phase the most likely cause is that two or more stations have tried to transmit in the same slot, having chosen the same slots when the bootstrap phase commenced, which is possible during the bootstrap phase. If such a case occurs, the slot status is set to collision thus moving the membership view towards consensus faster. If a station fails during the bootstrap its absence will be detected through the lack of transmissions, the collision process will promptly remove the failed station.

In a worst-case, two or more stations may select the same slot in which to transmit in during the bootstrapping phase, which stations receive which transmissions is a complex function of the proximity of the stations to each transmitter and any local propagation conditions, e.g. obstructions in the way. In the majority of cases, no station receives any valid transmission as a result of each transmission interfering with each other.

However there is a possibility that a subset of the stations would receive the transmission of station A, a subset of stations the transmission of station B and so on with the remaining stations receiving nothing meaningful at all. It is therefore necessary to provide an exchange of views of slot allocation. Since, no information of use is retained by slots marked as Available, Collision and Owner as the mac address of the owner is already in the header of each packet transmitted. The additional information portion of the beacon-frame contains the list of mac addresses and the slots allocated as viewed by each station.

The key concern during the bootstrap phase is to ensure it possible for all stations to achieve consensus, logically the fewer stations which take part in the process maximises the probability of each station achieving a consistent view.

The architecture of HD-TDMA provides no dedicated space for slot allocation functions unlike TBMAC, this introduces a small time penalty in the best-case slot allocation times. The clear benefit is that the average and worst-case times from the request to the first transmission in the allocated slot in HD-TDMA are better than in TBMAC. Also, when all slots have been allocated there are no idle resources which represents a key performance benefit.

If the received bitmap state for slot x is owner and the local bitmap for slot x is available or other where the mac address of the owner is unknown the local station updates its data structures to incorporate the information received. If a conflict occurs between the received data and the current local view the slot is marked as collision.

3.6.5.3 Non Empty Cell Scenario

Two discrete cases are be supported, a station may have no slot currently allocated to it and therefore cannot use a beacon-frame to convey its request or the station has at least one slot allocated to it and therefore can use the beacon-frame to convey its request.

HD-TDMA does not reserve a portion of the TDMA cycle for slot management functions, instead it relies on unallocated slots. This may appear counter intuitive however, in the non-empty cell case, a station with no slots can only seek a slot if there is in fact at least one unallocated slot available, thus it is logical to issue the request for a slot within a slot which is unallocated.

For a station which has at least one slot allocated to it it may utilise the slot management portion of the beacon-frame to broadcast its slot management request a number of times. Thus with all slots allocated there is no portion of the cycle idle satisfying the need to maximise performance.

Slot allocations under non-empty cell conditions occur in band with the requesting station utilising any slot which it determines to be unallocated to broadcast its slot request. In order to ensure slot allocations and data transmissions are segregated as far is as is practical, HD-TDMA includes a provision which allows a station to indicate in the beacon transmission that the remainder of the slot is idle and that it has no data to transmit.

A station, upon receiving the request, packages it into an atomic broadcast and sends it in the additional information portion of its beacon frames until its delivery time arrives.

3.6.6 Time Bounds

Within HD-TDMA all membership operations are time-bounded, ensuring that a station will know the outcome of each membership request before a known real-time deadline. Three distinct cases exist: the non-empty cell case which occurs when a station is in a cell which contains one or more active stations; The empty cell case, which represents the most challenging scenario, where no cell members currently are allocated slots; The routine management operations of allocate slot and de-allocate slot(s).

Following the discussion of the time bounds for the three scenarios, a overview is provided to show the best and worst-case times for membership-management operations, including a comparison with the membership protocol adopted by TBMAC upon which the HD-TDMA membership protocol is adapted.

3.6.6.1 Non Empty Cell

The non-empty cell case is the routine case where a station either enters a new cell or powers up within a cell in which at least one other station exists and has allocated to it at least one TDMA slot.

The first delay is fixed at one full cycle and one slot time to allow for one complete cycle to be observed, the addition of one slot time accounts for the worst-case delay if the new station begins to listen just after the first bit of a beacon-frame transmission would have been received.

The station must listen for a full HD-TDMA cycle time before it becomes aware of the presence of all other stations and considers the cell to be a non-empty cell. After one full HD-TDMA cycle plus one slot time has elapsed the joining station has full knowledge of the cells current slot allocations together with the list of pending updates. Under worst-case conditions the soonest unallocated HD-TDMA slot will be one cycle less one slot time away. Thus delaying the slot allocation request by that duration.

By default a membership-management request requires Δ time before it is applied, Δ being 2 full cycle times. Total worst-case time to first slot allocation in the non-empty cell case is therefore given as four HD-TDMA cycles less one HD-TDMA slot time. If the request is successful, the worst-case delay to the first transmission occurs if the slot allocated is the slot before the slot in which the membership request was issued, resulting in a delay of one cycle less one slot time.

The best-case occurs when the listening phase begins the instant before the first bit of a beacon-frame is to be received and that the completion of the listening phase ends with a empty slot, allowing the transmission of the membership request. If the same slot is subsequently allocated to this station the overall delay would be three HD-TDMA cycles.

3.6.6.2 Empty Cell

The bootstrapping phase is entered if after the initial one cycle plus one slot time slot learning phase no HD-TDMA beacon frames are successfully received. The cell is thus viewed as empty.

The worst-case delay case occurs when a station, the second station, enters an empty cell one slot time after the first station entered the same cell. The first station to arrive in the cell will initiate the epidemic protocol which will be detected by the second station which will wait until the epidemic protocol completes, three rounds and then commence the non-empty cell case, being able to skip the detection phase as the cell is known to be non-empty and full knowledge has already been acquired by monitoring the epidemic phase.

Total worst-case under the empty cell case is therefore given as one cycle to the beginning of the epidemic phase which spans three cycles. A full cycle plus the membership delay plus a delay of up to one cycles, resulting in a delay of 8 cycles less one slot time.

The best-case occurs when a station is the first to enter the empty cell, or enters a cell within the same slot time as several other stations enter, in which case after four cycles plus a single slot time all will have converged.

3.6.6.3 Slot Allocation & De-allocation

If a station has an existing slot allocation it will transit membership-management requests within the beacon-frame of its allocated slots, no learning phase is required. As the station is already a member of the cell one cycle time plus one slot time represents the worst-case delay before the request can be sent.

The standard delay of Δ applies until the membership-management request is applied. Worst-case occurs when a station imitates a membership-management request for a further slot the instant after that station sends its beacon-frame, where the station has only a single allocated slot, forcing a delay of one cycle plus one slot time, thus the overall delay of three cycles and one slot time.

Following allocation, a worst-case delay of one cycle less one slot time applies for the first transmission in the slot, if the allocated slot is the slot immediately preceding the slot in which the membership update was applied.

Best-case occurs when a station issues a request at the beginning of a slot allocation to that station, ensuring a maximum delay of Δ before the request will be applied. For slot allocation requests, if the slot immediately following the slot in which the request is made is allocated to this station the delay becomes Δ + one slot time.

3.6.6.4 Review of Time Bounds

Scenario	Best Case	Worst Case	TBMAC Worst
Non Empty Cell	3C	4C	5C
Empty Cell	4C + S	8C - S	6C
Slot (De)Allocation	2C	3C + S	3C + S

C = Cycle Time, S = Slot Time

Table 3.1: Summary of Time Bounds

Table 3.1 presents a review of key time bounds, comparing the membership approach used by HD-TDMA and that of TBMAC. Under worst-case for slot allocation operations HD-TDMA offers reduced time bounds as a result of the elimination of the start of TDMA cycle alignment of membership operations within TBMAC.

3.7 Summary

This chapter presented the HD-TDMA medium access control protocol. HD-TDMA is a flexible TDMA-based medium access control protocol in which a station may make several transmissions within each TDMA slot. HD-TDMA facilities the provision of a real-time wireless communication system by providing a balance between fixed structures in the form of TDMA slots needed to ensure contention and collisions free data transmission and the flexibility to allow for retransmissions and ordering of the transmission schedule at run time. The communication modes supported by HD-TDMA where also presented.

In order to support mobility between cells and also to respond to failures of stations, HD-TDMA incorporates a time-bounded and fault-tolerant membership-management service [Cunningham, 2003] which aims to ensure each cell member maintains a consistent view of the slot allocations. Slot management occurs in band, but in a manner as to be independent of real-time traffic. Time bounds for the three modes of operation of the membership service where presented and compared with those offered by the membership service within the TBMAC protocol.

Chapter 4

Probabilistic Admissions Control & Retransmission Protocol

This chapter describes a probabilistic admissions control and retransmission process which takes advantage of the scheduling flexibility offered by the multiple-transmissions-per-slot feature of the Hierarchal Distributed Time Division Multiple Access protocol introduced in Chapter 3.

The localised decision process consists of a number of functions which manage transmissions, implement the scheduling policy and provide a probabilistic admissions control policy. These functions work together to ensure that frames are transmitted at a time which satisfies their real-time requirements and maximises the probability of successful reception, based on known channel conditions. The admissions control policy, based on a clustered binomial distribution, ensures sufficient time remains unallocated within each slot to accommodate retransmissions.

The remainder of this chapter is laid out as follows: Section 4.1 introduces the admissions control process; Section 4.1.1 outlines the relationship between admissions control and the current and future slot allocations; Section 4.2 describes the mathematical model which underlies the admissions control process, introducing the clustered binomial distribution approach together with the beta binomial alternative; Section 4.3 details the queuing strategy employed; Section 4.4 reviews the nature of errors on the wireless medium and presents approaches to react effectively.

4.1 Admissions Control

In order to maintain real-time guarantees, it is essential to adopt an admissions control policy. The admissions control system should only admit a packet if the reliability requirement specified by that frame's sender can be satisfied, having sufficient space in a transmission slot to accommodate the frame and related transmission overheads such as acknowledgments to ensure transmission before the packet's real-time deadline.

As transmissions may not succeed due to transient effects on the medium, the admissions control process must also account for the transmission resources required to make a number of retransmissions to ensure that the reliability required by the frame is delivered.

Based on the frame type, its payload, transmission data rate, and the requirement for acknowledgment, the station calculates a time budget. This budget is the total time - including all overheads required to send the frame in question. Given the knowledge of the time budgets of all admitted frames, a further frame will only be admitted, if there is sufficient time remaining to accommodate its required time budget.

4.1.1 Relationship With Slot Allocation

An important relationship exists between the admissions control process and the slot allocation and de-allocate element of the membership management service. To implement a real-time data communication service it is necessary to ensure that frames are allocated to transmission slots which will still be allocated to this station at the time of transmission.

A performance trade-off exists between the time delay to perform membership management operations and the ability of the admissions control process to schedule frames in future transmission slots. Two specific cases occur: when a station issues a request to de-allocate a slot and when a slot is allocated to a station. Newly-allocated slots are incorporated into the transmission queue as soon as required.

4.1.2 Slot De-allocations

Slot de-allocations directly impact on the ability to schedule frames-once a frame is accepted that acceptance cannot be revoked. Frames can only be allocated to a transmission slot and related transmission queue which will exist at the time of transmission. No frame can be allocated to a transmission queue for a transmission slot at a time greater than the membership-request delay. Beyond the membership-request delay, the scheduler can have no confidence in the slot allocation state and therefore cannot allocate data to those slots.

Before a new transmission queue is associated with a transmission slot, a search must be performed of the membership-request queue to determine if a de-allocation request for this slot has been issued. As the delivery time of this request is known, it can be determined if the slot will remain allocated during the transmission slot and if so a queue can be created. It is assumed that all de-allocation requests will be successful, as there are no resource conflicts.

4.2 Transmission Resource Allocation

Where the expected delivery reliability ($p_{tx\_reliability}$) is equal to or greater than that demanded by the real-time requirements of the frame ($p_{rt\_requirement}$), then zero retransmissions are expected. This satisfies equation 4.7.

$$p_{rt\_requirement} \leq p_{tx\_reliability} \tag{4.1}$$

This section will describe a number of approaches to determine the transmission time resources required to meet the real-time timeliness requirements of frames. Each approach presented will progressively build on the previously described approach, at each stage reducing the overall transmission resource requirement. An alternative approach in the form of the beta-binomial distribution is also presented.

4.2.1 Individual Frames

As transmissions are not independent, if a transmission is made by station A to station B and that transmission fails if the retransmission happens within a short time frame of the first there is a higher than average probability that that transmission will also lost due to dynamic behavour of the channel, bursty error or channel fading due to propagation effects. For a transmission from station A to station B, the probability of transmission is given by $p_{tx_{ab}}$. Equation 4.2 presents the relationship were transmissions following a failure will not have the same probability of success as the initial attempt, $p_{tx_{ab2}}$ will be less than $p_{tx_{ab1}}$.

$$if\ Tx_{AB_1}\ fails\ p_{tx_{ab2}} \neq p_{tx_{ab1}} \tag{4.2}$$

In the case specific to HD-TDMA, each data frame comprises two portions, the management portion and the data portion, thus the success of a transmission, p_{tx} is the product of the probability of success of these portions, p_{data} represents the probability of failure of the data portion of a frame, while p_{mgmt} provides the probability of failure of the management portion of a frame.

Demarch et al [Demarch and Becker, 2007] provide an insightful analysis of reliability were transmissions are independent. Employing a similar analysis with a view to expanding and optimising its outcome, equations 4.3 and 4.4 define the probability of the successful receipt and acknowledgment of a frame. The acknowledgement case, unique to HD-TDMA requires that only the management frame be received by the destination, as sufficient information is contained within the management frame to generate a negative acknowledgment. Therefore the probability of reception of an acknowledgement of any type, p_{ack} is a function of the probability of management frame reception, unrelated to the data, which sets only the type of acknowledgement.

$$p_{tx} = (1 - p_{data}) \cdot (1 - p_{mgmt}) \tag{4.3}$$

$$p_{ack} = (1 - p_{mgmt}) \cdot (1 - p_{mgmt}) \tag{4.4}$$

The reception of the management portion of a frame requires the correct reception of fewer bits than the full frame, the probability of the reception of either class of acknowledgment frame, positive or negative is greater than that of a successful transmission, thus:

$$p_{ack} > p_{tx} \tag{4.5}$$

Negative acknowledgment enables an informed judgement of the channel conditions. Reception of a negative acknowledgment indicates poor channel conditions, failure to receive a negative acknowledgment indicates the channel cannot support communication even at the lowest system data rate, and may suggest the destination station has failed.

After two consecutive failures to receive at least a negative acknowledgment transmissions to the affected stations are cancelled for the remaining duration of the slot, as it is deemed that the destination station is either out of range, failed or suffering from severe propagation effects, making communication impossible. The probability of a frame failing (p_{fail}) on the nth attempt is given by equation 4.6, where p_{err} is the probability of a transmission failure.

$$p_{fail} = p_{err}^{n_{retries}+1} \tag{4.6}$$

Substituting $1 - p_{success}$ for p_{fail} Isolating $n_{retries}$ results in equation 4.7.

$$\left\lceil \frac{log(1 - p_{rt\_requirement})}{log(1 - p_{success})} - 1 \right\rceil = n_{retires} \tag{4.7}$$

Performing this evaluation for each frame individually results in a large number of required transmissions as the result is rounded up to the nearest integer. This large number of calculated transmissions, heavily constrains the number of admitted frames which can be accommodated as the duration of each transmission slot is fixed.

4.2.2 Transmissions to Same Destination

Considering each frame separately results in excessive and inefficient use of the available transmission resources. A method to group transmissions and pool the transmissions resources between frames would result in reduced resource usage while maintaining real-time guarantees.

As there may be more than one frame to transmit to each destination, a binomial distribution may be used similar to Demarch et al [Demarch and Becker, 2007]. The binomial distribution requires each attempt be a Bernoulli trial, such that each transmission be independent from previous attempts and that the probability of each attempt is constant. However, wireless communication is subject to significant temporal variations in transmission reliability due to issues such as channel fading, channel propagation and mobility.

In order to produce a formalised approach, a number of restrictions are introduced:

- All frames are the same length
- All frames are acknowledged
- Frames to the same destination have the same probability of success

As positive feedback exists, through the correct reception of an acknowledgment that a transmission was successful, the goal is to obtain y successes from n possible attempts, where y is the number of frames to be transmitted. The maximum value of n is bounded by the slot size. For each destination this problem can be modelled using a binomial distribution shown in equation 4.8.

$$p_{success} = \binom{n}{y} p^y (1-p)^{n-y} \tag{4.8}$$

Equation 4.8 provides the probability of exactly y successes, $p_{success}$, what is required is at least y success. Forming the cumulative density function (cdf), by summing the probabilities of all valid results were the delivery requirements are satisfied, e.g at least k successes.

$$p_{success} = \sum_{y=k}^{n} \binom{n}{y} p^y (1-p)^{n-y} \tag{4.9}$$

The goal is to minimise the transmission resources allocated in advance of transmission. By an iterative approach the target value of p_{target} can be achieved by increasing the value of n, until $p_{success} \geq p_{target}$. For equal $p_{success}$, n will be less than or equal to the sum of transmissions as determined by equation 4.7, or for equal n, $p_{success}$ will be greater than or equal in the binomial approach with respect to the individual approach.

This calculation is repeated for each destination and the sum of n determined, which is the total number of transmissions required to meet the reliability requirement and thus the time that must be reserved by the admissions control process can be determined.

4.2.3 Beta Binomial Approach

While the binomial approach results in a significant reduction in transmission resource requirements, it is only of benefit where a number of frames are being transmitted to the same destination. A more desirable approach is one were the result is optimal. Such an approach must consider all destinations by modelling the distribution of probabilities of successful transmission to each destination, weighted by a function related to the number of frames to be transmitted to those destinations.

The fundamental restriction of the binomial distribution is that the probability of attempt must remain constant throughout the process in order to satisfy the requirements of the Bernoulli trial. Situations where the probability varies such that each attempt may have a different probability of success cannot be handled by the binomial distribution. Therefore a distribution must be provided to describe the transmission probabilities. By expanding equation 4.8 to include a further function of p to represent the distribution of attempt probabilities, equation 4.10 is arrived at.

$$f(x|p) = \binom{n}{y} p^y (1-p)^{n-y} f(p) \tag{4.10}$$

In order to model the distribution of transmission success probabilities a beta distribution is introduced in equation 4.11, were $B(\alpha, \beta)$ is the beta function.

$$f(x|\alpha,\beta) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)} \tag{4.11}$$

Substituting equation 4.11 for $f(p)$ in equation 4.10, it can be shown that the beta binomial distribution may to the form of equation 4.12.

$$Pr(Y = y|n) = \binom{n}{y} \frac{B(\alpha + y, n + \beta - y)}{B(\alpha,\beta)} \tag{4.12}$$

Where the α and β parameters describe the distribution of the probabilities and may be estimated by the method of moments shown in equations 4.13 and 4.14, $\overline{x}$ being the mean and υ the variance.

$$\alpha = \overline{x}\left(\frac{\overline{x}(1-\overline{x})}{\upsilon} - 1\right) \tag{4.13}$$

$$\beta = (1-\overline{x})\left(\frac{\overline{x}(1-\overline{x})}{\upsilon} - 1\right) \tag{4.14}$$

As shown by [Garren, 2004, Prentice, 1986] the beta binomial distribution may be represented by equation 4.15 which allows for easier computation.

$$Pr(Y = y|n) = \binom{n}{y} \frac{\prod_{k=0}^{y-1}(\mu + k\theta) \prod_{k=0}^{n-y-1}(1 - \mu + k\theta)}{\prod_{k=0}^{n-1}(1 + k\theta)} \tag{4.15}$$

With θ and μ related to α and β, by equation 4.16. It can therefore be seen that the binomial distribution is in fact a special case of the beta binomial distribution were θ is zero.

$$\theta = \frac{1}{\alpha + \beta} \text{ , } \mu = \frac{\alpha}{\alpha + \beta} \tag{4.16}$$

It is assumed that each transmission has the same duration and target real-time reliability. For a given reliability, to determine the total number of expected transmissions an iterative process is followed by evaluating the cumulative density function of the beta-binomial distribution starting with a value of n equal to the number of messages to be transmitted plus one, as at least one retransmission is expected.

$$P_{cdf} = \sum_{y=z}^{n} \binom{n}{y} \frac{\prod_{k=0}^{y-1}(\mu + k\theta) \prod_{k=0}^{n-y-1}(1 - \mu + k\theta)}{\prod_{k=0}^{n-1}(1 + k\theta)} \tag{4.17}$$

While the beta binomial distribution appears attractive, the environment in which it would be applied does not lend itself to successful application. The number of frames sent within each slot is small-as a result the number of destinations considered by the distribution is also small, insufficient to form a beta distribution which accurately models the distribution of probabilities.

As stated by Cheong [Cheong, 2007], the number of events must be very large and that the distance between probability of different events must be small, neither criteria can be consistently assured as the number of transmissions per HD-TDMA slot is small in comparison and the number of individual destinations variable. If the beta Binomial distribution is applied to a scenario were these conditions are not met, the outcome will be invalid and may underestimation the number of transmissions required, resulting in a real-time reliability failure.

4.2.4 Binomial Clustering Algorithm

The probability of success ($p_{success}$) delivered by adopting the iterative approach in the previous section will be greater than or equal to the probability required, p_{target}. This excess is defined to be the residual:

$$residual = p_{target} - p_{success} \tag{4.18}$$

Equation 4.9 represents the optimal solution for a single destination, as transmissions to each destination will have difference success probabilities. For multiple destinations the sum of residuals may become significant. An initial setup cost also exists within the binomial distribution. As the number of frames to be transmitted increases, the ratio between transmissions and retransmissions reduces. Having identified these two concerns, a process to minimise the transmission resources required is sought.

For two destinations, a and b, where the probability of transmission success is equal such that, $p = p_a = p_b$, were k_a and k_b frames are to be transmitted, the application of equation 4.9 will result in n_a and n_b transmissions respectively were k_a and k_b successes where required.

As transmissions to different destinations are statistically independent [Bhagwat et al., 1997] and as both destinations share the same probability of success, they may be combined into a single binomial distribution. If combined into a single binomial distribution, the overall number of transmissions is given by equation 4.19 as n_{ab}, where $k = k_a + k_b$.

$$p_{success} = \sum_{y_a+y_b=k}^{n_{ab}} \binom{n_{ab}}{y_a + y_b} p^{(y_a+y_b)}(1-p)^{(n-y_a-y_b)} \tag{4.19}$$

As a result of the residual factor and the set up cost of the binomial, it can be shown empirically that equation 4.20 holds for $p = p_a = p_b$ independent of the number of transmissions (n) made by stations a or b.

$$n_{ab} \leq n_a + n_b \tag{4.20}$$

Given equation 4.20, there may exist cases were the probability of p_a and p_b need not be equal, just similar. It may be stated that for an unknown value of $\epsilon \geq 0$, the relationship presented in equation 4.20 continues to hold, such that the probability of success, p_b is slightly greater than p_a.

$$p_b = p_a + \epsilon \tag{4.21}$$

To take advantage of equations 4.19, 4.20 and 4.21, groups of destinations with similar probabilities must be identified. A search algorithm is therefore envisaged which scans the proposed transmission queue to identify destinations sharing similar probabilities of transmission success, with the aim of clustering these destinations to pool the transmission resources. Three separate clustering approaches are considered:

- **By Probability**: By clustering, based on the destination pair in each round which had the least difference in probability.
- **By Variance**: Extending the first approach by seeking the cluster pair with the least variance. This is weighted by the number of transmissions.
- **Exhaustive Search**: Ignoring probabilities entirely, instead seeking to cluster destinations based on the reduction in transmissions required if that pair was clustered.

The computational overhead of options one and two are similar, while the exhaustive search is comparatively expensive. The clustering algorithm is based on a progressive clustering of destinations which share similar probabilities of transmission success. The algorithms begin with each destination considered to be a cluster head, replicating the initial binomial approach from section 4.2.2. Each round of the algorithm identifies the neighbouring clusters which share the least difference in probability, variance or transmission requirements, and merges them. Within each cluster, the destination with the lowest transmission probability becomes the cluster head.

After each round, equation 4.9 is applied to determine the number of transmissions, with the probability of the cluster head as the p parameter and were y is the sum of frames to be transmitted by all cluster members. As the lowest probability within the cluster is used, the reliability requirements of all cluster members are ensured. After each cluster is formed, the overall sum of transmissions across all clusters is determined and recorded.

The clustering algorithm continues until there exists only one cluster, such that all destinations are considered together using the lowest transmission probability of all destinations. The lowest calculated sum of transmissions determined during the clustering process is returned, allowing the admissions control policy to decide if sufficient time exists in the slot to meet the real-time requirements.

While the clustered binomial approach does not guarantee an optimal outcome, it does ensure that the outcome, the transmission time required will be less than or equal to that calculated through the binomial by destination approach discussed in section 4.2.2 and that the result will be mathematically valid and that real-time requirements are met, in fact exceeded by a small margin.

4.2.5 Admission Control Example

The proposed binomial clustering algorithm is best explained through a simple example which highlights the significant resource savings which may be achieved. Each HD-TDMA slot is fixed in duration, thus the goal of the admissions control process is to determine if that slot time is sufficient in duration to transmit a set of frames while accounting for the reliability of each frame and the need for retransmission to meet real-time requirements as a result.

Consider a scenario consisting of a total of 3 destinations and 5 frames to be transmitted, a single frame to destination b, two frames each to destinations a and c. A transmission reliability, p_{target} of 0.95 is sought. Transmission success probabilities were assigned to each destination at random as follows:

- $p_a = 0.578813$
- $p_b = 0.755945$
- $p_c = 0.599707$

Considering each frame separately and applying equation 4.7 results in a sum total of 19 transmissions. Considering each destination using binomial distribution through the application of equation 4.9, results in 7, 3 and 6 transmissions for destinations a, b and c, a total of 16.

Applying the clustering algorithm identifies that destinations a and c have the least difference in probability. Thus destinations a and c are merged with the cluster reliability set to the lowest probability of success of the cluster members, $p_a = 0.578813$. Application of equation 4.9 determines that the cluster *ac* requires 11 transmissions, giving a total of 14 transmissions.

The second round of the clustering algorithm finds that only one possible cluster combination is possible, ac with b. A new cluster is formed *acb*, with the reliability set to the lowest reliability of a cluster member, $p_a = 0.578813$. A total of 13 transmissions are calculated by equation 4.9. Table 4.1 presents a summary of the results.

Approach	Round 0	Round 1	Round 2
Considering individually	19	19	19
Grouped by destination	16	16	16
Clustered approach	16	14	13

Table 4.1: Transmissions Required to Satisfy Transmission Reliabilty Requirement

The result of the algorithm is the lowest calculated number of transmissions recorded during the clustering process. For this example case, the lowest values was recorded in the second round, with 13 transmissions. With the number of transmissions known, the transmission time is determined, including all overheads, e.g. acknowledgments. If this time is less than the slot time, the frames may be transmitted in this slot.

4.3 Queuing Strategy

As it is possible to transmit more than one data frame within each HD-TDMA slot, it is necessary to implement a queuing strategy to ensure that frames are ordered and allocated to transmission slots in such a manner as to satisfy real-time requirements. Each transmission slot allocated to each station is associated with an independent queue on that station. These queues are accessed by the admissions control process as it seeks to determine the slot in which to place incoming real-time frames from the application layer.

A further queue for best-effort traffic is also maintained in which non real-time traffic is queued on a FIFO basis, for transmission when a beacon frame indicates a slot to be in the idle state.

While the modelling of transmission reliability described in section 4.2 seeks to ensure sufficient transmission time is reserved to ensure all admitted frames can be successfully transmitted. On rare occasions, it may occur, due to unforeseen circumstances, that the transmission time allocated is insufficient and that not all frames can be transmitted. In this case priority should be given to frames in order of real-time priority. As this thesis proposes, a firm real-time[1] communications sub-system occasional failures are acceptable provided the reliability requirements of the application data are satisfied.

As frames are ordered initially by priority, should a circumstance arise where the transmission time allocation is insufficient, priority is given to high priority real-time traffic. Should a transmission failure occur a, dynamic reordering of the queue will take place, such that all frames destined to the affected destination will be placed at the

---

[1] A firm real-time system is one where after the deadline the value, or message of interest has no utility [Kopetz, 1997], but unlike a hard real-time deadline failure will not result in a catastrophic event.

rear of the queue in priority order by destination while the remainder of the queue will remain in priority order.

The number of transmissions queues required is a function of the number of slots allocated to this station and the latency in the application of membership requests. Each queue must be associated with a transmission slot which will exist at the time of transmission. Frames cannot be queued for a transmission slot which at transmission time may no longer be allocated to this station.

4.3.1 Admission Control

A probabilistic admissions control protocol is employed to ensure sufficient transmission time is reserved, to accommodate an estimated number of retransmissions. The transmission time reserved for retransmissions will vary depending on factors such as frame size, transmission speed, reliability requirements and number of frames. An evaluation is made at the point of admission based on the characteristics of currently admitted frames and the frame seeking admission.

The scheduler relies on feedback from the MAC layer to adjust its response to conditions, however the failure of a frame transmission may result from a large number of reasons, therefore the scheduler is forced to assume that the primary cause for a failure, the result of bursty error conditions.

To ensure the correct operation of the deferred retransmission scheme to combat burst errors, two scheduling requirements are enforced: first, all frames must be admitted no later that the slot start time; secondly, that the real-time deadline of all frames admitted to a particular slot must be greater than the end of that slot's slot time. This is less restrictive than the Balasubramanian et al [Balasubramanian et al., 2006] approach which requires that all frames within a transmission cycle to be available at the start of the TDMA frame and have a deadline that expires after the end of the frame to eliminate scheduling concerns.

HD-TDMA requires that frames to be transmitted in each slot be queued for transmission before the slot's beacon frame is sent. If no frames are queued, the beacon frame sent will indicate an idle data slot and all stations in the cell, regardless of slot allocation status will use the slot to transmit best effort and sporadic messages.

4.4 Transmission Control

The transmission control sub-system is responsible for selecting the correct sub-queue from which to transmit frames. When the MAC module enters the data transmission portion of the HD-TDMA slot, the transmission control module is informed of the state of the data transmission slot, being either idle, unallocated or owned by this station. If the slot is idle, the transmission queue will select the best effort transmission sub-queue and begin transmitting the frames on this sub-queue through a CSMA channel access method since in idle slots all stations may compete for the medium.

The most common mode of operation is OWNER, indicating the slot is allocated to this station. The transmission control sub-system selects based on the current slot number and time the appropriate transmission sub-queue and begins to transmit the frames in priority order. Should a transmission failure occur, the transmission schedule will be adapted to minimise its impact.

4.4.1 Generating Transmission Independence

Transmissions are not independent, if a transmission is made by station A to station B and that transmission fails, a retransmission is required. If the retransmission happens within a short time frame of the first there is a higher than average probability that that transmission will also lost due to dynamic behavour of the channel, bursty error or channel fading due to propagation effects.

Channel bit error rate (BER) is calculated based on fixed criteria such as the signal to noise ratio, thermal noise, path loss, modulation technique, bandwidth and data rate. This, however, fails to account for dynamic effects such as the bursty error condition.

It is considered that the bursty error condition exists at most once per-destination per slot.

In order to improve frame reliability and to simplify the analysis the concept of a per-destination blackout period is introduced together with a per-destination queue similar to [Bhagwat et al., 1997]. As such when a failed transmission is encountered, the failed communications link is marked as dead. Practical evaluation by Willig et al [Willig et al., 2002] showed a 0.71 probability of a further failure following a failure for the evaluated scenario; highlighting the lack of independence of transmissions. Unlike Balasubramanian et al [Balasubramanian et al., 2006] as there is the ability to make multiple transmissions per slot, slots are not exchanged between stations to facilitate the delayed retransmission process.

Transmissions to other stations are dynamically reordered within a slot to take advantage of the idle medium and the statistical independence of the communication links to other stations. Once the transmission of frames to unaffected destinations are completed, transmissions are again permitted to the affected station. This approach ensures the independence of each transmission and allows the application of an analysis similar to Demarch et al [Demarch and Becker, 2007] in a realistic environment.

4.4.2 Localised Decision Making Process

Having described the MAC approach in chapter 3, the localised decision making process which supports the admissions control process is now discussed. This process builds on the features of the HD-TDMA protocol primarily the slot structure which makes it possible to implement a localised decision process while retaining distributed control of slot allocation.

The localised decision process consists of two components, a scheduling policy and an admission control policy. These work together to ensure that frames are transmitted at a time which satisfies the real-time requirements and maximises the probability of successful reception based on known channel conditions. The admission control policy ensures sufficient time remains unallocated to accommodate retransmissions and sporadic messages.

4.4.2.1 Scheduling Policy

A station has total control over its local transmission schedule and requires no agreement from other stations. There exists, however, a precondition that all transmissions are ordered in a way that satisfies real-time deadlines. This is a two stage process: first, the station secures sufficient slots, choosing slots which best suit the timeliness requirements of its transmissions; secondly, to maximise the possibility of recovery from a blackout, a portion of each slot is reserved to accommodate retransmissions, the size of this portion varies and is decided by the probabilistic admissions control process based on frame reliability requirements and transmission reliability estimates.

The communications link between each station is considered to be statistically independent [Bhagwat et al., 1997]. If a blackout condition occurs, e.g. a communications channel goes through a bursty error phase, the immediate retransmission of a frame or further immediate transmissions to that station are considered a poor choice [Willig et al., 2002]. However communication links to other stations may be unaffected. To maximise the number of successfull transmissions, transmissions to other stations are dynamically reordered within the slot to take advantage of the idle medium and the statistical independence of the communication links to other stations in a similar fashion as [Bhagwat et al., 1997, Balasubramanian et al., 2006]. Once frames to unaffected stations are complete, transmissions are again permitted to the affected station(s).

4.4.3 Transmission Queue Example

In order to demonstrate the operation of the protocols reaction to a transmission failure an example is presented in figure 4.1. A single queue representing the queue associated with a slot is shown and this queue is populated with a number of frames with varying real-time priorities to be sent to several destinations. The queue state at the commencement of data transmission from the queue is shown in figure 4.1.a, frames to all destinations are ordered by priority.

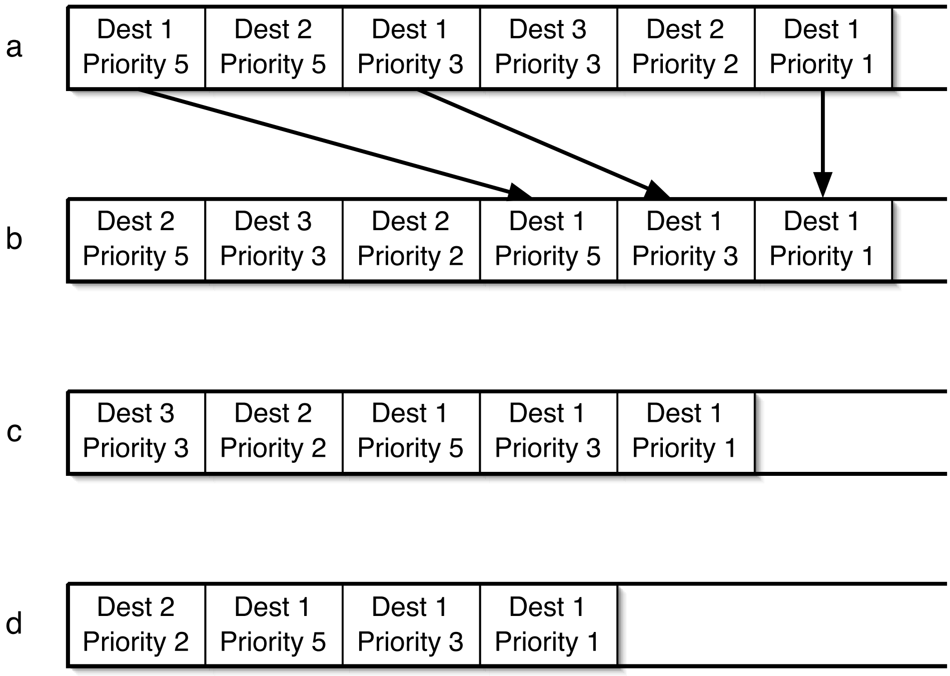

Fig. 4.1: Slot Transmission Queue Following Transmission Failure

An attempt is made to transmit the frame at the top of the queue, which is destined for destination 1, with an assigned real-time priority of 5. The transmission attempt fails. The failure is notified to the admissions control and retransmission protocol following the failure to receive a valid acknowledgment frame. The retransmission control system reacts by placing the failed frame back into this slots transmission queue.

The ordering of queue is changed in order to defer transmissions to the destination of the failed transmission attempt in order to reduce the probability of a second failure. All queued frames to destination 1, including the frame requiring retransmission are placed to the rear of the queue, while frames to other destinations are promoted, in real-time priority order to the front. Frames to destination 1 remain sorted by priority at the rear of the queue, as shown in figure 4.1.b.

Transmissions from the queue continue, with the queue order maintained as shown in figures 4.1.c and 4.1.d. As a result of the retransmission protocol and the queue ordering all high priority frames unaffected by transmission failures will have been transmitted before the slot end. If at the end of the slot time, frames remain in the queue, they are destroyed.

4.5 Summary

This chapter described the admissions control and scheduling approach utilised to support firm real-time communication in the wireless domain. Two specific problems were addressed in this chapter, first, that of scheduling and rescheduling of transmissions to maximise probability of success and, secondly, an efficient means of admissions control.

In order to reduce the number of frames lost due to interference and propagation effects, a per-destination deferral process is introduced. Application of this deferral process restores independence between transmission attempts permitting the use of a Bernoulli process to model transmissions.

A number of approaches to determining the transmission resources required were presented. A clustering algorithm based on binomial distribution was developed to group transmissions on the basis of destinations, and destinations with similar probabilities, in order to minimise the transmission time allocation, while maintaining real-time guarantees.

Chapter 5

Implementation of the HD-TDMA Protocol

The previous two chapters presented a detailed description of the design of the HD-TDMA protocol together with the probabilistic admissions and retransmission protocol. This chapter outlines the implementation of those two designs in the OPNET modeller network simulator [OPNET Technologies, 2008]. An overview of the simulator environment is presented in section 5.1.

A high level review of the implementation is provided first, with the implementation structure described in section 5.2. The MAC, Queue and Membership models are described in section 5.3.

Section 5.4 details the implementation which permits the modelling of errors. The principal functions and data structures are outlined, together with the execution flow of key operations of the HD-TDMA MAC. Section 5.5 focuses on the process of sending and receiving frames. Section 5.6 reviews the admissions and queuing process implementation, while membership operations are described in section5.7. The chapter then concludes with a summary.

5.1 Implementation Environment

The OPNET modeller simulator is a discrete event-based simulator. The software code which executes in response to events takes zero execution time. As a result the simulation result is idealised by not considering execution time.

The simulator operates on a queue of scheduled events, either explicitly set by the software author or set by the simulation software as a result of an action, e.g. sending a frame requires the calculation of the propagation delay which will schedule a frame arrival event at each destination after the appropriate interval has elapsed.

The basic functional element is that of a process model, which consists of a number of states related by conditional transitions with related functions to execute, on transition, on entry and on exit from a state. There must exist exactly one valid transition for each event.

A node is made up of several process models joined, typically by FIFO queues to exchange information, together with statistical links to determine the state of a node. For example, a wireless receiver node provides a statistic output as to the wireless channel being active or idle. A number of pre-built models of common wired and wireless networking hardware is provided.

OPNET relies heavily on graphical tools to allow rapid development of state machines as well as the relationships between processes. Development consists of developing each component by drawing a state machine diagram and adding transitions with related conditions and operations. Two programming languages are supported by OPNET, C and C++. The work presented in this thesis was implemented entirely in C.

The OPNET simulator provides an ideal real-time environment, as all scheduled events occur at the scheduled time, if more than one event is scheduled for the same time, the priority of the event is considered. Where an event is scheduled for several stations at the same time, the event occurs sequentially, starting with the node with the lowest id number as assigned by OPNET, thus execution is deterministic.

Communication between processes is primarily based on FIFO queues, as each process adds a packet to a FIFO queue an interrupt is generated for the process, connected to the FIFO allowing that process to retrieve the packet. The software api provided is not dissimilar to real world real-time Linux implementations such as RTAI [Racciu and Mantegazza, 2006] which share a similar real-time FIFO based inter-process communication paradigm.

5.2 Design

A modular design and implementation approach was adopted. The HD-TDMA MAC layer relies on both a membership and queuing service to operate. HD-TDMA does not mandate a specific membership or queuing protocol. To ensure as far as practical that each module was independent, the implementation is broken into four distinct modules, each implementing a separate element of the overall communications architecture.

5.2.1 MAC Module

All medium access control elements reside in this module. The timing and generation of events related to the HD-TDMA cycle originate within this module. The module interfaces to the physical layer devices which provide the ability to send and receive frames to/from the wireless medium.

5.2.2 Queue Module

The queue module is the most complex module, implementing not only the probabilistic admissions control process but also the retransmission protocol. Internally the module is modelled as a pipeline: Frames are accepted (or rejected), stored in queues and then forwarded to the MAC module.

5.2.3 Membership

The membership module implements the necessary queues and procedures which are required to establish, maintain and update the allocation of the HD-TDMA slots. The slot allocation information is shared with the queue module to ensure the admissions control process is aware of which slots are allocated and those which are scheduled to be de-allocated.

5.2.4 Application Interface

The application interface acts as a bridge or middleware between the HD-TDMA communications sub-system and its real-time FIFO based communication model. This bridge enables the functionality within the HD-TDMA communications sub-system to be accessed using the traditional function calling model as to allow applications to execute function calls to send frames, to request allocation and de-allocation of slots and to register callbacks for events such as frame received and frame could not be accepted for transmission.

This interface has been successfully employed to enable SEAR [Hughes et al., 2006] to be ported to the OPNET simulator and communicate with the HD-TDMA subsystem [Zhang, 2008].

5.2.5 Inter-Module Communication

Communication between modules is through a series of FIFO queues which, by default, generate an interrupt at the destination when a packet is inserted, the interrupt will be processed in order of priority once the current thread of execution has completed. Figure 5.1 shows the relationships between the modules, FIFO queues in solid blue lines and statistic links in dashed red. The normal direction of information flow is shown by an arrow.

At times it is necessary for a module to invoke a process on another module and be

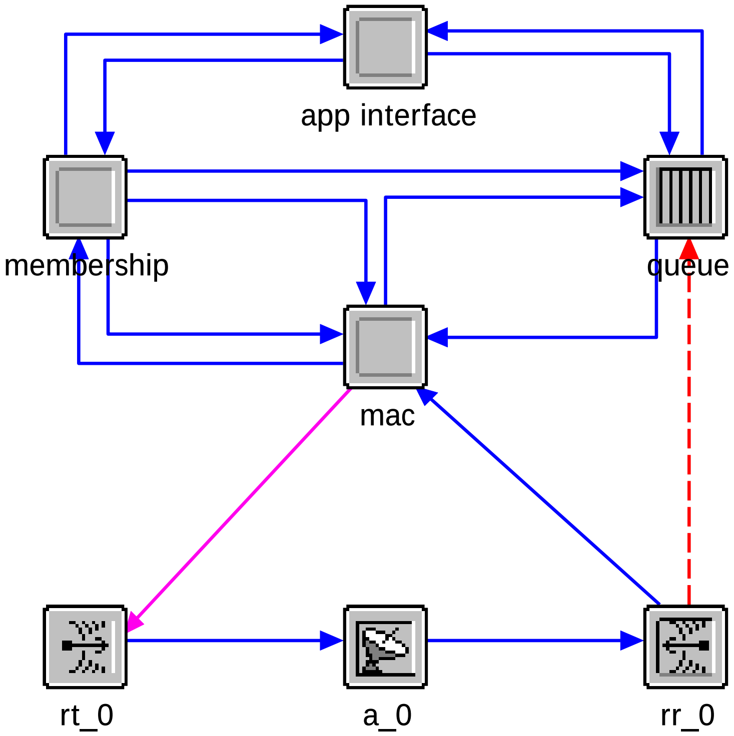

Fig. 5.1: Module Relationships

returned a value before continuing. Use is made of the OPNET access interrupt feature which allows the destination of a FIFO to raise an interrupt on the module connected to the source of the FIFO. This provides an instant switch of context to the remote module to process the interrupt, once the interrupt has been processed, control returns to the invoking process, thus allowing the currently executing function to collect the result from the FIFO and continue. The use of this feature avoids the need for complex synchronisation and additional states, thus simplifying the development process.

5.2.6 Shared Memory

An implementation difficulty was the lack of a flexible shared memory system within the simulation software. Three modules (MAC, Membership, Queue) require access to the slot allocation structure. While it is possible to use a series of interrupts to produce the same result this results in significant overhead and implementation complexity. While not effecting the outcome of simulations, the simulation run-time would be significantly increased and make debugging the code significantly more difficult as a result of the need for several extra states and interrupt handlers.

A shared memory structure is equivalent to an expected hardware implementation where dual port memory may be used, in a fashion similar to controller networking interface (CNI) within TTP/A [Kopetz, 1997]. At the initialisation stage the membership module allocates the necessary data structures and transmits the memory address to the other modules through the FIFO queues provided, using an internal management frame. Consistency of the shared data is ensured as shared memory access is based on a single writer (membership module) multiple readers model (MAC and Queue modules). No locking protocol is required, as there can be only a single thread of execution within each node[1].

5.3 Modules

5.3.1 MAC Module

The MAC module implements the HD-TDMA channel access method and provides the overall control and synchronisation, triggering events in other modules, through the generation of interrupts, by sending frames through real-time FIFO queues provided by the OPNET environment.

The state machine shown in figure 5.2 implements the necessary logic to execute the HD-TDMA slot structure, with the state machine broken into four distinct elements: Two active elements in which transmissions are made or received, and two passive elements which implement the required inter-slot and inter-frame spacing.

---

[1] Multiple threads are possible, but restrict the api available for simulations. Single thread execution was used in the described implementation

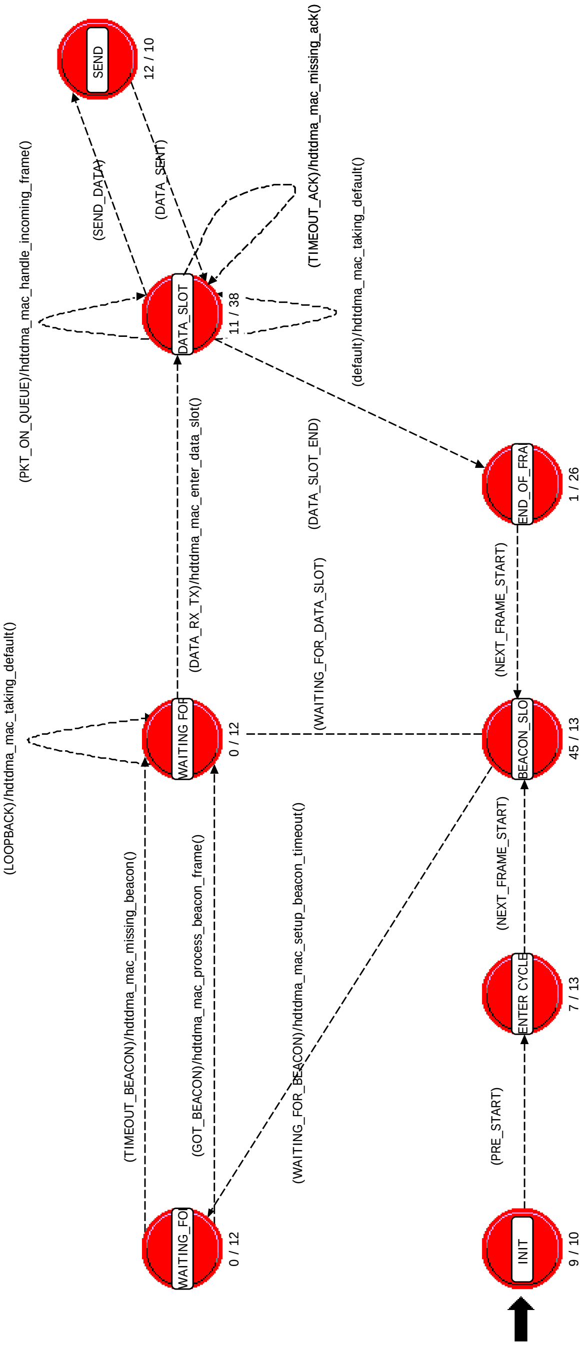

Fig. 5.2: HD-TDMA MAC State Machine

If a slot is allocated to this station, the OPNET `op_strm_access()` function is employed which notifies the Membership module that a beacon-frame is required. The membership module generates the required data for a beacon-frame and passes it to the MAC module for transmission. Following the beacon-frame transmission the system leaves `BEACON_SLOT` state and enters the `WAITING_DATA_SLOT` state awaiting an interrupt to signal the start of the data transmission phase. Otherwise, a beacon-frame may be expected from another station and a transition to the `WAITING_BEACON_FRAME` state occurs, if the slot is allocated. Upon reception of a beacon-frame it is forwarded to the membership module for processing, otherwise, if a beacon-frame is not received a management packet is created and forwarded to the membership module, indicating that an expected beacon-frame was not received.

Following the beacon-frame stage, the MAC enters a `WAITING_DATA_SLOT` state for the main data transmission phase. When leaving the `WAITING_DATA_SLOT` state the MAC enters the data transmit and reception state, `DATA_SLOT` . Based on the slot state frames are requested from the Queue module using the OPNET `op_strm_access()` function applied to the FIFO queue between the Queue and the MAC, invocation of the access call hands execution to the Queue module which places a frame in the FIFO if one is available. If no frame is available the system will enter a wait state until a frame is received from the queue module.

Immediately before a frame is sent, a calculation is performed as to the duration the frame will take to transmit, based on its transmission speed, plus a short inter-frame space and an allowance for an acknowledgment frame, if required. A task is scheduled to raise an interrupt when the transmission transaction is completed.

If no acknowledgment is received within the expected window of the previously scheduled interrupt, a management frame is generated and sent to the Queue module. The MAC will then request the next frame from the queue module.

When a frame is received an interrupt is generated by the wireless receiver module, allowing the MAC to read the frame, determine its error state and then pass the frame to the Queue module and/or Membership module, while issuing an acknowledgment or negative acknowledgment, as required.

An interrupt signals the end of the slot, forcing a change of state to the second waiting state for the commencement of the next slot so as to ensure that at least a full inter-frame space is provided between adjacent slots transmissions.

5.3.2 Queue Module

The principal function of the queue module is to take data from an application and to ensure its successful transmission in line with the data's specified real-time reliability and timeliness requirements. The Queue module implements the admissions control and retransmission protocol by ensuring that only frames which the system can meet the real-time requirements of, are accepted for transmission.

The OPNET simulator provides a special process model type which includes advanced packet queue management. A single queue exists within which numerous independent sub-queues are provided, which may be addressed by index or by a statistical criteria such as the queue with the most frames, for example. In this implementation a unique sub-queue is used for each instance of each assigned HD-TDMA slot.

Two operations are supported by the queue module: admission of frames and management of transmissions where the initial attempt to transmit to a destination has failed.

5.3.3 Membership Module

The membership module implements the real-time membership service required to support not only the MAC module, but also the Queue module. At the start of each slot an interrupt is received from the MAC module which causes the Membership module to inspect the queue of membership requests and to apply all membership-management requests due for this slot time. When an interrupt from the MAC module is received, a beacon-frame is generated containing the current snapshot of the slot structure, incorporating all changes made by the application of membership-management requests at the beginning of this slot time. The beacon-frame also contains a number of future membership-management requests.

5.4 Error Modelling

The default setting within OPNET is for any frame received with one or more bits in error is dropped without notifying the layers above. Such frame drops occur due to a station being out of reception range or due to collisions on the medium. While this represents a perfectly reasonable approach where a frame has a single CRC check, in the case of HD-TDMA data frames are made up of two distinct portions each with a dedicated 16 bit CRC for error detection purposes.

To provide error modelling functionality an error sub-system is implemented within the MAC module which processes all incoming frames. Information determining if a frame is valid is obtained from the error model applied to simulate both single bit and burst error conditions. Figure 5.3 presents the interface to the error model. The incoming frame is passed to the error model which responds with the CRC result.

```
int crc_state = hdtdma_dbu_error(incoming_frame);
```

Fig. 5.3: Determining Error State Of A Frame

The return value, `crc_state` describes the error state of the incoming frame. As each HD-TDMA data and beacon-frame is composed of two distinct portions, four error permutations exist. In the case where both portions fail, `FAIL_BOTH_CRC` no useful information can be obtained from the frames reception. If the management portion is received correctly, but the data portion is not, the result `FAIL_DATA_CRC` is returned. As the frame-type together with the originator and destination address are known, so a negative acknowledgment frame may be sent.

Failure of the CRC of the management portion, `FAIL_MGNT_CRC` or failure of both CRC's `FAIL_BOTH_CRC` are considered equivalent. As the management portion specifies the length and type of the data sub-frame, failure of the management CRC results in an inability to decode the data portion correctly. The optimum outcome is both CRC's being valid, `VALID_CRC`.

5.4.1 Burst Errors

In order to consistently model burst error conditions during simulation, a set of trace files may be read in for each transmitter-receiver pair. A trace file may be either generated through a random process or from real-world data. The trace contains the start and end times of burst error conditions on the medium, together with the source associated with that burst event. Each transmitter receiver pair has a unique trace.

By taking the arrival time of the packet (the time the last bit of the frame is received), the time of the arrival of the first bit can be determined from the packet length and transmission speed. A search is performed on the trace to determine if the frame was impacted by a burst error condition, and if so, what portion(s) of the frame were effected. Returning an appropriate error code as discussed in section 5.4.

An optimisation is applied, if the burst model returns `FAIL_MGNT_CRC` the single bit model case is not calculated, as the entire frame has already been been determined as being irreparably lost.

5.4.2 Single Bit Errors

While OPNET provides a function to perform a single bit error calculation determined by the bit error rate (BER), in the case of the HD-TDMA model a specific need exists in this implementation as to determine the location within a frame of an error if an error occurs, so the correct CRC error code can be returned. Upon receipt of a frame, the single bit error model is invoked.

Two single bit error operations are performed, first, the management sub-frame and secondly over the data sub-frame, if it is required to specify the length in bits to apply the error calculation function over. Based on the frame-type, appropriate calls to the function in figure 5.4 are made. For optimisation purposes, the data sub-frame is only processed if the management sub-frame bit error calculation returns a valid CRC result.

```
int error_model_single_bit_error_calc(double pe, double random,
                                      int seg_size)
```

Fig. 5.4: Calculating Single Bit Error Of HD-TDMA Sub-Frames

The single bit error function is parametrised by three parameters, `pe` the probability of a single bit error, `random` a randomly distributed variable between 0 and 1 and `seq_size` the number of bits to operate over, thus allowing the frame to be processed in two chunks, the management sub-frame and the data sub-frame, to meet HD-TDMA requirements.

5.5 Communication

Once at least one station within a cell has been allocated a slot, real-time data communication is possible. Three distinct types of frame exist, a non real-time i.e. best effort packet, a scheduled real-time packet and a sporadic real-time packet. The communications sub-system comprises functionality in the MAC, Membership and Queue modules.

Such distributed functionality is necessary to ensure a tight level of integration to ensure that each component has an awareness of changes which impact on the services they provide. For example the Queue module must be aware of which slots are allocated to the local station, which slots are scheduled to be de-allocated and how far into the future confidence exists in a slot remaining allocated to this station.

This section will describe key elements of the communication system through the presentation of examples of how frames are sent and received, together with the steps followed when a transmission fails.

Frame Type	ID
SCHEDULED_ACKED_DATA_TX	0x0
SCHEDULED_NOACK_DATA_TX	0x1
RETX_ACKED_DATA_TX	0x2
RETX_NOACK_DATA_TX	0x3
SPORADIC_ACKED_DATA_TX	0x4
SPOARDIC_NOACK_DATA_TX	0x5
OPTIMISTIC_ACKED_DATA_TX	0x6
OPTIMISTIC_NOACK_DATA_TX	0x7
BEACON_TX	0x8
BEACON_TX_IDLE	0x9
JOIN_REQUEST	0xA
MAC_MANAGEMENT_FRAME	0xB
RESEVERED	0xC-OxD
POS_ACK	0xE
NEG_ACK	0xF

Table 5.1: HD-TDMA Frame Types

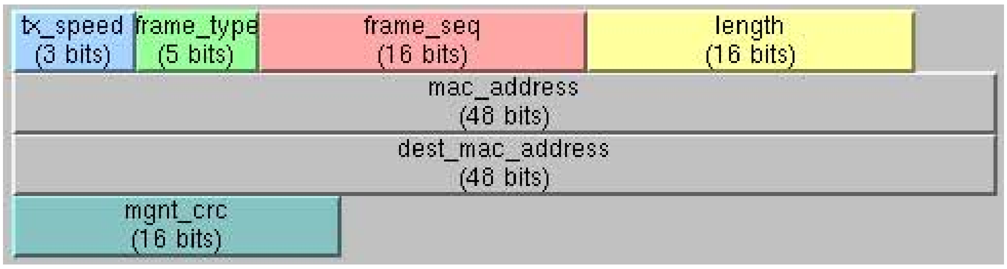

Fig. 5.5: Management Frame

Table 5.1 details the various frame-types used by HD-TDMA. The physical layer frame-types used by HD-TDMA are presented as implemented in the OPNET packet format tool, in figures 5.5, 5.6 and 5.7. All frames share the same basic structure of

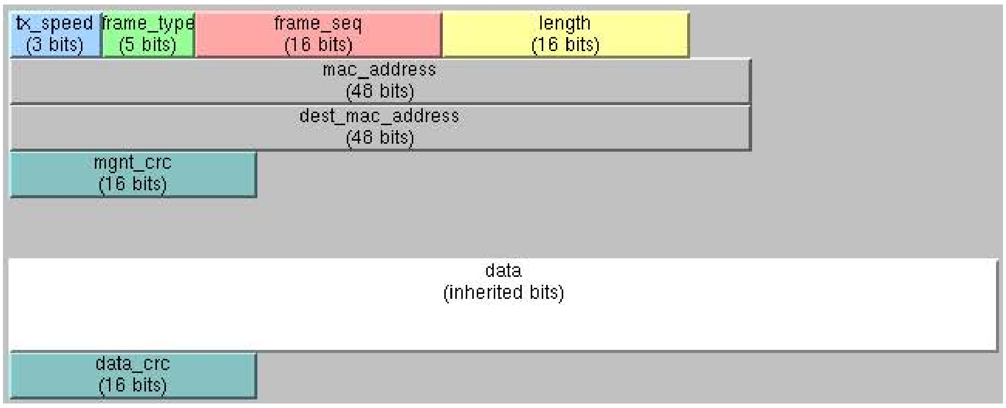

Fig. 5.6: Data Frame

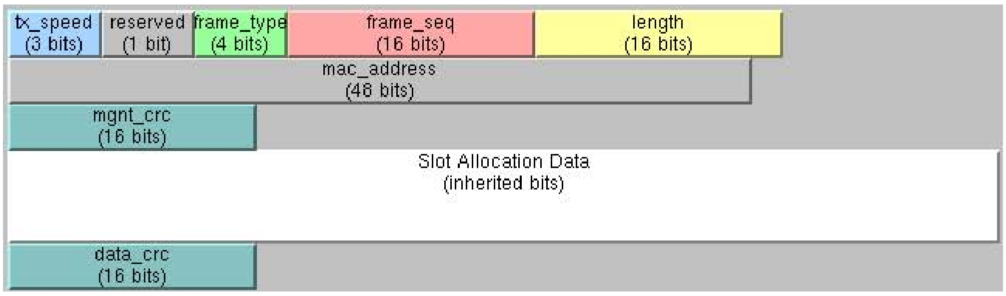

Fig. 5.7: Beacon Frame

control and addressing information followed by an optional data sub-frame.

5.5.1 Transmission of a Real-Time Data Packet

Transmission of a real-time data packet consists of four distinct stages:

- Receipt of Packet from Application
- Queue Selection and Admissions Control Process
- Removal of frame from queue for transmission
- Transmission

First, upon receipt of the frame from the application, the type of frame must be determined. If the frame is a best effort frame, it is simply queued in the best effort FIFO queue for transmission in slots marked as idle by received beacon frames, or in slots where time exists after all real-time frames have been successfully transmitted.

Otherwise the frame needs to be inserted into a transmission queue which has sufficient resources to meet the reliability requirements of the frame, and that ensures the frame will be transmitted before its real-time deadline. The admissions control process must search through the transmission queues seeking a transmission queue which meets the needs of the frames real-time reliability and deadline.

The determination of which queue to use is made by the binomial clustering algorithm described in section 4.2.4, which seeks to determine the minimum number of transmissions required to meet the real-time requirements of a group of frames.

The second stage of a data frame transmission is the removal of a frame from a transmission queue and transmission over the wireless medium. If data transmissions are possible in the current slot, the MAC module will request a frame from the Queue module, which will reply with the real-time frame of the highest priority.

```
void hdtdma_scheduler_pkt_from_application();
```

Fig. 5.8: Interrupt Handler For Packet From Application

The application is joined to the Queue module through a real-time FIFO, when an application inserts a frame for transmission an interrupt is generated on the Queue module, which will be processed using the function shown in figure 5.8.

A real-time data frame has both a real-time deadline and a real-time priority. A search is triggered by the admissions control protocol to identify slots in which this frame may be sent. The search is constrained by the membership-management latency. The search begins with the slot with the soonest transmission time.

```
result = hdtdma_scheduler_insert_into_tx_queue(pkt_to_send,pkt_deadline);
```

Fig. 5.9: Inserting A Frame Into A Transmission Queue

Where `insertion_result` returns the outcome of the admissions control process, three possible outcomes exist:

- **SUCCESS**: If the packet was successfully added to a queue within the packets real-time deadline, which also meets the real-time reliability requirement.
- **CANNOT_SATISFY_REAL_TIME_REQ**: If it is not possible to meet the real-time requirements i.e. no space exists in any queue until beyond this frames deadline.
- **FAILED_TO_ALLOCATE_SUBQ**: None of the existing queues could accommodate the requirements of the frame, but the real-time deadline has not yet been exceeded, so the admissions control system attempts to allocate a further queue to a future transmission slot. Queues can only be associated with slots with slot end times of less that Δ HD-TDMA cycles in the future, due to the delivery delay of membership requests.

5.5.2 Assigning Real-Time Packet To A Queue

To assign a packet to a transmission queue, a search is performed of the transmission queues associated with each transmission slot between the next slot time and the next slot time plus the membership agreement time, Δ which is typically two full HD-TDMA cycles.

A progressive search through the list is made, starting with the soonest transmission slot allocated to this station, until a slot is found which can meet the real-time requirements of the packet to be transmitted. If no suitable slot is found an error code is returned.

```
int hdtdma_scheduler_insert_into_tx_queue_2(QueueList_First_Ptr,
    pkt_to_insert, pkt_deadline, dest_mac_address);
```

Fig. 5.10: Assigning A Frame To A Transmission Queue

While a restriction ensures packets are not allocated to transmission slots too far in the future, the membership process may allocate slots. Thus, if the search algorithm finds a slot which has not associated with a transmission queue, a transmission queue will be created and inserted at the correct location in the transmission queue list as determined by the slot time.

Figure 5.11 shows the function which creates and adds a new transmission queue and places its after the referenced queue. A pointer to the newly created queue is returned.

```
NewQueuePtr = hdtdma_scheduler_insert_into_queue_list(QueueList_Last_Ptr,
                                                      AFTER);
```

Fig. 5.11: Adding A Queue To The Transmission Queue

When a suitable queue has been determined by the admission control process, the packet is placed at the rear of the sub-queue using the OPNET provided insertion function shown in figure 5.12. OPNET ensures the frame will be placed into the queue at the appropriate location for the frames priority.

```
op_subq_pk_insert(CurrentPtr->sub_q_number, pkt_to_insert, OPC_QPOS_TAIL);
```

Fig. 5.12: Insert A Frame Into A Queue

5.5.3 Sending A Frame

In the previous sections, the process followed to insert a frame into a transmission queue was discussed. The process to remove a frame from a queue and to send the frame is now discussed. Upon entering the data transmission and reception phase of the HD-TDMA slot, the MAC module will request a frame from the Queue module. If the slot is allocated to this station and real-time data is queued for this slot, the Queue module will select that queue and return a frame.

Otherwise, regardless of ownership, if the idle state has been indicated by a beacon-frame, best effort traffic may be transmitted by all stations regardless of their slot allocation state. If the slot is unallocated the slot is reserved for use by membership requests for station which currently have no slots allocated to them.

The MAC module generates an interrupt to the Queue module seeking a frame to transmit by passing an internal management frame. The Queue module will identify the queue from which draw a frame from, based on the current time through a search of the transmission queues to locate the queue allocated to this slot time. When a queue has been identified the transmission control function will determine which frame to give to the MAC module.

The frame is then removed from the queue, passed to the MAC module, which executes the function shown in figure 5.13, which completes the sending of the frame by passing the frame to the wireless transmission module.

```
void hdtdma_mac_send_data_pkt();
```

Fig. 5.13: Sending A Packet

Based on the frame-type, the MAC module will schedule an interrupt at a time in the future when the transmission is complete, accounting for the frame being acknowledged or not. In order to schedule the transmission ended interrupt a time budget must be calculated, the function provided (figure 5.14) accounts for inter-frame spaces and the possibility of the data portion of the frame being transmitted at a different rate to the base rate of the management sub-frame in determining the total transmission time. If an acknowledgment was required and none is received before the interrupt executes a management frame is sent to the Queue module indicating a failed transmission.

```
double hdtdma_mac_calc_timebudget(Packet * pkt);
```

Fig. 5.14: Determining Timebudget For A Packet

5.5.4 Receiving A Frame

Reception of a data frame occurs the instant following the simulated arrival of the last bit at each station. The simulator calculating the appropriate arrival time, based on distance from transmitter and transmission speed. Upon arrival an interrupt is generated, which triggers a handler to run within the MAC module shown in figure 5.15. This function removes the frame from the real-time FIFO connected to the receiver module and commences processing.

```
void hdtdma_mac_handle_incoming_frame();
```

Fig. 5.15: Reception Of A Frame

After receiving a frame it is first processed by the error model described in section 5.4. If the CRC response is `CRC_VALID` or `FAIL_DATA_CRC` the frame is processed. Arrival of a frame with a `FAIL_MGMT_CRC` outcome from the error model, results in the frame being destroyed.

If the frame is a data broadcast frame, or is addressed to this station, it is passed to the Queue module which will, in turn, pass the frame to the application. Reception of an positive acknowledgment by the Queue module confirms that the real-time requirements of both timely and correct delivery were achieved for the frame in question. Otherwise the frame is discarded. Frames of type `JOIN_REQUEST` are passed to the membership module for further processing.

The MAC module will generate and transmit an acknowledgment of type `POS_ACK` in the `CRC_VALID` case or `NEG_ACK` in the case of `FAIL_DATA_CRC`. This acknowledgment will have the same frame sequence number of the frame to which the acknowledgment is in response.

5.5.5 Transmission Queue

Each station maintains a list of queues associated with each transmission slot up to Δ HD-TDMA cycle times into the future. As each frame is received from the application layer, it is assigned to the soonest transmission slot which satisfies its real-time requirements. If a slot has been scheduled for de-allocation, no frames will be allocated to transmission queues associated with that slot beyond the delivery time of that de-allocation request.

Figure 5.16 presents the data structure employed to store the details of each transmission slot. The key data are the end time of the, slot `end_of_slot_time` and the unallocated time remaining in the slot, `available_time`.

The OPNET simulator provides a comprehensive queuing package which allows frames within queues to be accessed in a direct fashion, e.g. head or tail of queue, by index in the queue or by a indirect form based on statistical quantity such as frame size, priority, length of time in the queue.

In this implementation, only direct queue addressing mechanisms are employed and all queues are ordered automatically by frame priority. Each sub-queue provided by OPNET has a unique index number, which is stored by the `sub_q_number` variable. A linked list is maintained of all current transmission queues. The process of assignment of frames to transmission queues is discussed in detail later in section 5.6.2.

```
struct QueueList{
     struct QueueList * NextPtr;
     struct QueueList * PrevPtr;
     double end_of_slot_time;
     double available_time;
     double recovery_time;
     signed int sub_q_number;
     unsigned int slot_number;
     unsigned int no_of_pkts_admitted;
     struct DestinationProbability * HeadPtr;
     struct DestinationProbability * TailPtr;
};
```

Fig. 5.16: Data Structure Transmission Queues

5.5.6 Beacon Frame

At the beginning of each allocated slot, the station allocated this slot will send a beacon-frame containing membership information to all other members of the cell. The beacon-frame is critical to the management of slot allocations and the detection of failed stations.

Upon entering a slot, if the MAC address of the slot owner is that of the local station, the function shown in figure 5.17 is called, which, in turn, will invoke the function shown in figure 5.18 to hand the thread of execution to the membership service thus allowing it to populate the beacon-frame with the latest membership information.

```
static void hdtdma_mac_send_beacon_frame();
```

Fig. 5.17: Sending A Beacon Frame

Control returns to the MAC module with the populated beacon-frame being received

```
op_strm_access (BEACON_TO_MAC_FIFO);
```

Fig. 5.18: Requesting A Beacon Frame

by the MAC module through a FIFO. The MAC module then sends the beacon-frame. Stations aware of the allocation of this slot will schedule an interrupt through the function call shown in figure 5.19. The absence of a beacon-frame may indicate failure of a station.

```
static void hdtdma_mac_setup_beacon_timeout();
```

Fig. 5.19: Setup Beacon Timeout

If no beacon-frame has been received, an interrupt will be generated as a result of the previously set-up time out (figure 5.19), invoking the function shown in figure 5.20. This function generates an internal management frame directed to the membership module.

```
static void hdtdma_mac_missing_beacon();
```

Fig. 5.20: Detection of a Missing Beacon Frame

Failure to receive two beacons from the same station will trigger a suspicion phase within the membership protocol, which will attempt to coordinate with the other active stations in the cell to confirm the failure of the station and thus de-allocate the slots allocated to that station. The code required to generate the management frame is shown in figure 5.21.

Upon commencement of the optimisation algorithm, each destination is considered independently, each entry in the list is deemed to be a cluster, progressively clusters are merged in a manner as to minimise the number of overall transmissions required. Within each cluster, a single destination is assigned to be the cluster head: this is always the destination with lowest transmission success probability, in order to ensure real-time

```
mgmt_frame = op_pk_create_fmt ("hdtdma_management_frame_format");
op_pk_nfd_set(mgmt_frame,"frame_type",MAC_MANAGEMENT_FRAME);
op_pk_nfd_set(mgmt_frame,"sub type",FAILED_TO_RECIEVE_BEACON);
op_pk_send(mgmt_frame,MAC_TO_MEMBERSHIP_FIFO);
```

Fig. 5.21: Generation Of Management Frame Following Missing Beacon

```
struct DestinationProbability{
     unsigned long long dest_mac_address;
     double probability;
     unsigned int count;
     unsigned int total;
     double target;
     unsigned int cluster_head;
     unsigned int cluster_count;
     unsigned int cluster_total;
     struct DestinationProbability * Next;
     struct DestinationProbability * Prev;
};
```

Fig. 5.22: Data Structure Describing Destination Probabilities

reliability requirements are met.

The variable `cluster_count` represents the number of frames to be transmitted by this cluster, e.g the sum of `count` across all cluster members. The total number of transmissions required to meet the reliability requirement `target` for a cluster is given by `cluster_total` which has been determined by the clustered binomial approach (section 4.2.4).

To insert a further frame into a slots transmission queue, the function presented in figure 5.23 is invoked. This function first determines if the queue already contains an entry for this destination. If so, the `count` variable is incremented and the `total` variable is recalculated to reflect the additional frame.

```
int hdtdma_scheduler_add_frame(unsigned long long dest_mac_address,
            double target, struct QueueList * ActiveQueue);
```

Fig. 5.23: Adding A Frame

If this is the only frame for this destination, a new entry is placed in the queue. The location determined by probability of transmission success. The function ensures that the queue remains ordered by ascending transmission success probability. Once a frame has been added to the queue the optimisation function is called by passing a reference to the queue to the function presented in figure 5.24.

```
int transmissions = hdtdma_scheduler_optimise_count_by_probability(struct
                              QueueList * ActiveQueue);
```

Fig. 5.24: Determining Optimum Transmission Count

The outcome, `transmissions`, is converted into a duration which represents the total transmission time required, including inter-frame spaces and acknowledgments if required. If the calculated duration is less than the slot time, sufficient resources exist to meet the real-time requirements and the frame is accepted for transmission in this transmission slot.

If the slot is unable to meet the real-time requirements of the frame seeking admission, a function (figure 5.25) is executed to rollback the insertion of the frame by reducing the `count` variable by one, if the `count` becomes zero as a result, the destination is removed from the queue.

```
static void hdtdma_scheduler_remove_frame(unsigned long long
     dest_mac_address, struct QueueList * ActiveQueue);
```

Fig. 5.25: Adding A Frame

5.5.7 Failed Transmission

When a failed transmission is detected, either through the omission of an acknowledgment or the receipt of an negative acknowledgment, the retransmission sub-system must react in order to minimise the potential for further frame loss while also ensuring the lost frame is re-transmitted to meet the real-time requirements. A failed transmission notification will generally be triggered by a time out leading to an interrupt handled by the function shown in figure 5.26.

```
static void hdtdma_mac_missing_ack();
```

Fig. 5.26: Detection Of A Failed Acknowledgment

5.6 Admissions Control

The admissions control process requires the ability to assign packets for transmission in a certain slot at a time in the future: to enable this, dedicated queuing services are provided which incorporate knowledge of both the membership and transmission control systems.

As the transmission resources available to the admission control system are limited to the membership service latency, frames can only be queued for transmission where confidence exists that the assigned transmission slot will be available at transmission time.

As an atomic broadcast system is used by the membership service, the atomic broadcast parameter Δ indicates the number of cycles required from issuing a membership request until it is be applied. Typically a value of two HD-TDMA cycles is used. As confidence exists only to a maximum of Δ cycles in the future, the queue module maintains a sub-queue for each transmission slot up to a maximum Δ cycles in the future.

Figure 5.27 presents a block diagram of the queuing process, incorporating the admissions control process. As each packet is received from the application interface, an attempt is made by invoking the admission control process to queue the packet in such a way as to satisfy its real-time requirements, by selecting an appropriate transmission queue and slot. When the MAC layer requests a packet, the sub-queue for the current transmission slot is selected. If the current slot has no queued frames, the idle case exists and the best effort sub-queue will be selected. The admissions control process is informed upon receipt of an acknowledgment, negative acknowledgment or failure to receive an expected acknowledgment.

The key operation within the admissions control process is the repeated application of the binomial distribution to determine the number of transmissions required to meet real-time application requirements. The number of calculations is considerable thus optimising this process is desirable. To minimise the number of iterations of the binomial distribution within each calculation, the inverse probability is calculated. The number of iterations is then set by the number of frames to transmit, not the number of total transmissions required, significantly reducing the computational overhead. Equation 5.1 is therefore implemented, where k is the number of frames requiring transmission.

$$p_{cdf} = 1 - \sum_{y=0}^{k} \binom{n}{y} p^{y} (1-p)^{(n-y)} \tag{5.1}$$

A secondary aspect of this approach is a further reduction computational overhead as for $y = 0$ will result in $p^y = 1$ and for $y = 1$ will result in $p^y = p$. Figure 5.28 presents the optimised algorithm employed to determine the binomial coefficients.

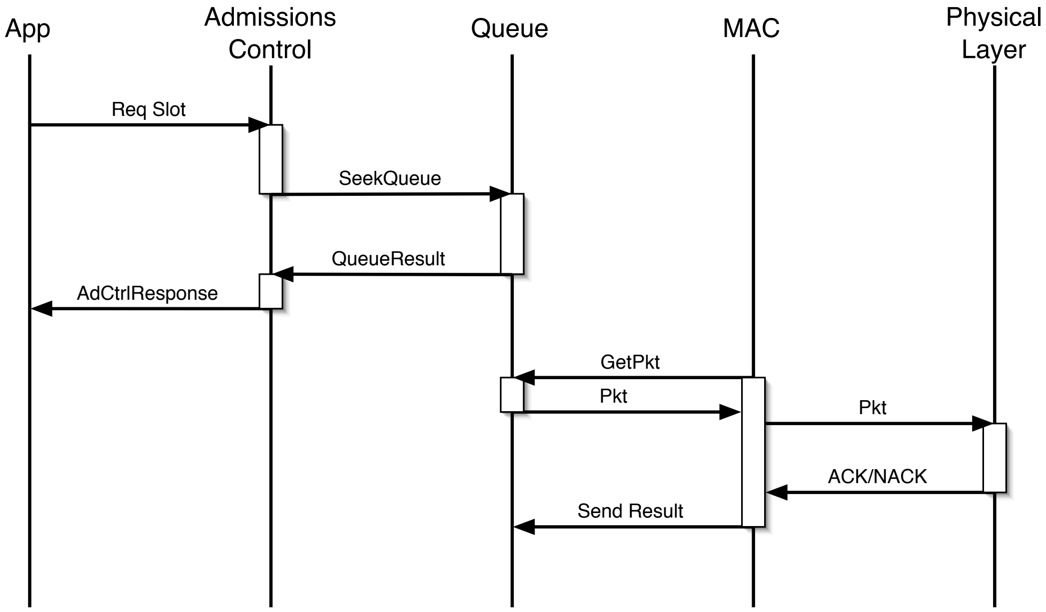

Fig. 5.27: Sequence Of Queue Operations

```
top = bottom = 1;
for(i = 0 ; i < x ; i ++){
        top = top * (n - i);
        bottom = bottom * (x - i);
}
return top / bottom;
```

Fig. 5.28: Binomial Coefficient Calculation

5.6.1 Slot Transmission Queue

Each transmission slot is associated with a dedicated queue for each occurrence of that transmission slot between the current time and *current time* $+ \Delta$ in the future. By default, frames are queued in priority order in order to satisfy the need to prioritize transmissions should a failure occur preventing successful transmission of all frames.

The retransmission protocol as described in section 4.2.4 may override the default ordering, in order to maximise the probability of successful transmissions. The transmission queue is implemented using the sub-queue api provided within the OPNET simulator, which ensures automatic ordering by priority. A new sub-queue is assigned to each slot when the first frame is allocated to that transmission slot.

5.6.2 Cluster Algorithm

In order to implement the clustering admissions control process described in section 4.2, a series of data structures and functions are provided. Within each HD-TDMA slot, there are potentially many frames to different destinations.

Each transmission slot contains a reference to a linked list storing details of the destinations and number of queued packets to each to support the admissions control process. The data structure shown in figure 5.22 is provided to describe the number of frames from each destination, as well as details of the number of transmissions required to meet the real-time needs of those frames.

5.6.3 Assigning A Frame To A Slot

Once accepted, a frame has a commitment that it will be transmitted, thus a high level of confidence must exist that the slot in which the transmission will be made will exist when it is called upon to send the frame.

A further restriction is that of the atomic broadcast function of the membership service, which requires a number of HD-TDMA cycles before a membership-management request can be applied, this delay is given by the Δ parameter, as discussed in section 3.6. It is assumed that all requests to de-allocate a slot will be successful, therefore no frames can be allocated to slots for slot times beyond the delivery time of a request to de-allocate that slot.

Frames can only be assigned to slots for transmissions less than Δ time in the future, as beyond *current time* $+ \Delta$ no knowledge is available as to the slot allocation state of all stations. Where this station has issued a request to de-allocate one or more slots, the admissions control process will no longer assign frames to those slots beyond the delivery time of the membership request. It is assumed that de-allocation requests always succeed, whereas it is assumed any pending slot allocation request will fail. Combining these two conditions ensures frames are only allocated to valid and available slots for transmission. When an application wishes to send a frame it creates a message containing:

- Destination
- Data
- Real-Time Reliability
- Real-Time Deadline
- Transmission Mode[2]

Upon receipt, the admissions control process seeks to locate the soonest transmission slot in which sufficient time remains available to transmit the frame seeking admission, accounting for real-time factors of reliability and timeliness.

The probabilistic admission control process as described in section 4.2 is then applied to determine the overall transmission time required by the frames currently allocated to this slot, together with the frame seeking admission, to consider if sufficient resources exist to accommodate the number of estimated transmissions to ensure that all frames meet their target reliability requirements.

If the slot is unsuitable, the next available slot is checked, a process continuing until a termination condition is met. If the frame cannot be allocated to any slot a matching error code is returned. Three possible failure conditions exist:

- A suitable slot is found

---

[2] Does the frame require an acknowledgment

- The end time of the slot exceeds the deadline of the frame
- All available slots have been exhausted due to the Δ membership time restriction

5.7 Membership Operations

The membership service must learn of the current state of the cell in which is physically located, through the beacon frames it receives from the MAC module. Based on these received frames, or the absence of them, the membership module determines the current cell state and so reacts accordingly through the state machine outlined previously in section 5.7.2.

```
typedef struct{
     unsigned long long requesting_address :48;
     double delivery_time;
     short seq_number :16;
     char number_allocation;
     char request_type :3;
}AdditionalInfo;
```

Fig. 5.29: Structure Of Membership Request

The key data structure is that of the membership request, shown in figure 5.29. This describes a request to change the slot allocation state, by either requesting that a slot be allocated, that a specific slot be de-allocated, that all slots be de-allocated, or that a conflict exists in the ownership of a slot.

5.7.1 Membership Manager

The membership manager receives all membership requests and beacon frames from the MAC module. A key component of the membership-management sub-system is the ordered list of pending membership-management requests. This queue is not sorted, but kept in permanently correct order by carefully inserting all new membership requests in the appropriate location, thus avoiding a complex and time-consuming sorting algorithm.

The queue of pending membership updates is implemented as a doubly linked list with membership requests dynamically added and removed. However as membership requests may be re-transmitted numerous times, by different stations as part of the atomic broadcast function it is also necessary to ensure each entry in the queue is also unique.

When a membership request is received it must be inserted in the local membership request queue or discarded if already in the queue. If the current membership queue is empty, the request becomes the first entry. If the request has a delivery time more distant in the future than any other it becomes the tail element. Otherwise a process is followed to determine the correct insertion location of the membership request based on the delivery time, then the requesters MAC address and finally the request sequence number.

5.7.2 Membership States

The membership service employs a state machine to model the state of the membership service. The bulk of states relate to the process of a station gaining its first slot. The state directly effects the manner in which a station submits membership-management requests, as three distinct approaches exist.

Figure 5.30 presents the membership state machine and the relationships between the various membership states. The membership state machines consists of 8 independent states:

- INIT
- EMPTY_CELL
- NON_EMPTY_CELL

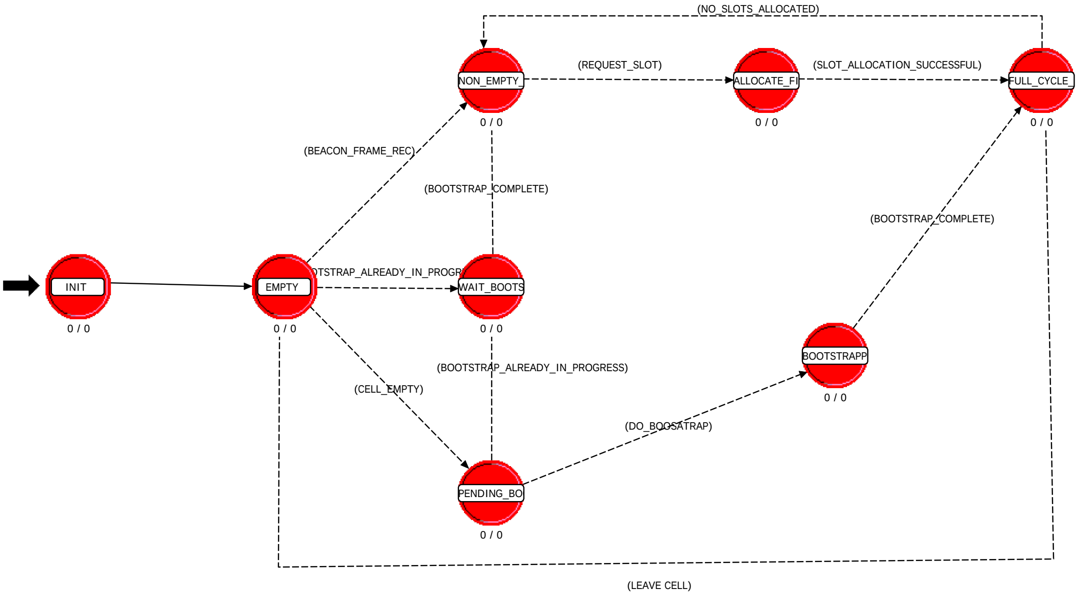

Fig. 5.30: HD-TDMA Membership State Machine

- PENDING_BOOTSTRAP
- WAIT_BOOTSTRAP_END
- BOOTSTRAPING
- ALLOCATE_FIRST_SLOT
- FULL_CYCLE_MEMBER

Each state is described in detail, including the conditions for transition to other states:

- **INIT**: The `INIT` state is the default entry state entered when the membership service is activated. Initialisation tasks take place in this state. Once completed, an automatic transition to the `EMPTY` state occurs.

- **EMPTY**: By default, the membership system enters the empty cell state where it waits while the state of the cycle is determined. This will take at most $cycle\ time + slot\ time$, regardless of the point in the cycle a station powers on/enters the cell.

The addition of one slot time is necessary to account for the fact the station will have to wait for the next slot to begin, before it can begin to listen.

Transitions from the this state occur, if after one complete cycle time elapses, at which point the cycle is confirmed to have no allocated slots resulting in a migration to the `BOOTSTRAP` state.

If a valid beacon-frame is received, the station is in an active cell, in which at least one other station exists. The station transitions, therefore, to the `NON_EMPTY_CELL` state. If the beacon-frame is a normal transmission i.e. not associated with the bootstrapping protocol, and this station requires a slot, the station transitions to the `ALLOCATE_FIRST_SLOT` state. If the beacon is marked as a bootstrap transmission, a transition occurs to the `WAIT_BOOTSTRAP_END` state to allow the bootstrapping protocol to complete.

- **NON_EMPTY_CELL**: If a valid beacon-frame from the normal data transmission state is received within the one cycle listening window, the station is in an active cell and that other stations are transmitting normally. If a station requires to transmit, it may request a slot, to do this it must transition to the `ALLOCATE_FIRST_SLOT` slot state. If the station merely wishes to receive data, HD-TDMA allows receivers to directly acknowledge packets and supports a polling mode without being allocated a slot, so the station remains in this state.
- **ALLOCATE_FIRST_SLOT**: If the station has powered on in a non-empty cell, during the `WAIT_BOOSTRAP_END` or `NON_EMPTY_CELL` case and requires to transmit it must first seek a slot. In this state the station chooses a slot which it knows to be unallocated and makes a number of transmissions of its request at random points within that slot.
- **PENDING_BOOTSTRAP**: If station powers up/enters a cell and after one full cycle time finds that the cell is empty, it will decide to bootstrap. The member-

ship service will generate a list of random slots to allocate to itself. The pending bootstrap is the state between deciding on the need to bootstrap and transmitting in the first slot chosen randomly by the station. If the station detects that another station has sent a bootstrapping beacon-frame before it has sent its first bootstrapping beacon-frame, the station will abort the scheduled transmissions and await the outcome of the bootstrapping process, in doing so minimising the number of stations in the bootstrap phase.

- **FULL_CYCLE_MEMBER**: Through either the bootstrapping or allocate first slot process, the station now has at least one slot allocated to it and thus is a full member of the cell and uses the membership portion of the beacon-frame to modify its slot allocation(s) from now on. If the station de-allocates all its slots it returns to the `NON_EMPTY_CELL` state.

- **BOOTSTRAPPING**: When a station commits itself to the use of the bootstrapping process it enters the `BOOTSTRAPPING` state from the `PENDING_BOOTSTRAP` state. The station will send a HD-TDMA beacon in the random set of slots the station chose while in the `PENDING_BOOTSTRAP` state, these beacons will be marked as `BOOTSTRAP` type thus clearly identifying them to all stations in the cell that the cell is being bootstrapped.

  The station will then, following a number of rounds come to agreement as to the state of the slot allocations, based on the beacon frames it has received from other stations in the cell. The use of the `PENDING_BOOTSTRAP` and `WAIT_BOOTSTRAP_END` states aims to minimise the number of stations in the `BOOTSTRAPING` phase. If two or more stations enter the bootstrapping phase at approximately the same time, whichever station transmits a `BOOTSTRAP` frame first will cause the other stations to move to the `WAIT_BOOTSTRAP_END` state. If two stations or more stations choose the same slot for there first bootstrapping slot transmission the station must resolve the slot allocation state.

- **WAIT_BOOTSTRAP_END**: If a station joins during the bootstrap phase or was in the `PENDING_BOOTSTRAP` state and detected another station beginning the `BOOSTRAPPING` phase the station moves to the `WAIT_BOOTSTRAP_END` phase after which it will enter the `NON_EMPTY_CELL` state.

Each allocated slot begins with a beacon frame transmitted by the slot owner containing membership information. Figure 5.31 presents the function call triggered to populate an empty beacon-frame with membership data in the form of two membership requests which reside in the 'Slot Allocation Data' portion of the beacon.

```
static void hdtdma_membership_populate_beacon_data();
```

Fig. 5.31: Generating The Beacon Frame

When a beacon-frame is received by the MAC module it is passed by a FIFO to the membership module, which results in the raising of an interrupt which is handled by figure 5.32. The membership requests within the beacon-frame are then integrated with the current membership request queue.

```
hdtdma_membership_handle_beacon_frame(received_frame_from_mac);
```

Fig. 5.32: Processing Beacon Frame

The bootstrapping phase represents a challenge to ensure stations can acquire sufficient information concerning the status and slot allocation of all slots, such that consensus can be achieved. It is key that stations exchange their view of slot ownership, not just the allocation state as is the normal case.

The beacon-frame, instead of membership updates, carries a list of slot owners of which it is aware. A special frame-type of `BEACON_TX_BOOTSTRAP` is used to ensure receiving stations can distinguish between the contents of the slot allocation information portion in the data sub-frame. This special frame-type allow stations to also recognise if the cell is currently in bootstrap phase.

By default, each beacon-frame carries only two membership updates, each update being 192 bits in length, thus a space of only 384 bits is available to convey the slot ownership information, only seven 48 bit MAC addresses will fit in the available space as each MAC address must be accompanied by the slot number to which it refers.

The slot bitmap structure provides significant information concerning slot ownership. It is possible to infer from the slot bitmap the ownership of many slots. For entries set as owner, the mac address of the senders is contained in the beacon-frame. For the collision or unallocated state, no mac address is relevant. The other field is the only field of interest. A more efficient way to represent the ownership of slots is shown in figure 5.33. Where the mac address is the owner of the specified slots. Where a station is aware of only one or two slots being allocated to a station, the remaining entries repeat the same slot number. A MAC address of zero represents an unused entry. By this mechanism it is possible to convey ownership information on a maximum of six distinct stations and up to 18 slots in total.

```
(slot x, slot y, slot z, mac_address)
```

Fig. 5.33: Compact Slot List

As each station during the bootstrapping process allocates three slots at random to itself each beacon-frame sent within the bootstrap phase best case contains the owners of up to 21 slots, 3 owned by the beacons generator and up to 18 others owned by 6 other distinct stations.

When a station enters the bootstrapping state, it generates a list of slots at random to allocate to itself, using figure 5.34, typically a station will choose three slots. If the station has not detected another station active in the cell by the time of the transmission time of the first slot chosen, the station commits to the bootstrapping process figure 5.35.

```
static int * hdtdma_membership_generate_slot_list(int end,
                    int number_of_results_req);
```

Fig. 5.34: Generating Slot List

```
void hdtdma_membership_do_bootstrap();
```

Fig. 5.35: Performing Bootstrap

If during the pending bootstrap phase a station detects another station before it has made its first transmission in a bootstrapped slot, the bootstrap process is cancelled by this station and will wait for the outcome of the bootstrapping operation. The execution of the function 5.36 terminates the scheduled bootstrap process.

```
void hdtdma_membership_cancel_bootstrap();
```

Fig. 5.36: Cancel Bootstrap

If a station enters a non-empty cell or cancels a pending bootstrap process, the station will perform a normal membership request after the bootstrapping process has completed, by executing the function show in figure 5.37.

```
static void hdtdma_membership_do_non_empty_cell_join();
```

Fig. 5.37: Joining A Non Empty Cell

5.8 Summary

This chapter described the implementation of HD-TDMA, its membership services as well as the probabilistic admissions control and retransmission protocol in the OPNET network simulator.

The overall structure and design approach adopted was described. The modular nature of the implementation in the form of the MAC, Queue and Membership modules was discussed together with the key elements of the implementation. Key data structures and the flow of control through state machines for key system events were presented. This implementation will be used in the evaluation presented in chapter 6.

Chapter 6

Evaluation

Chapters 3 and 4 described HD-TDMA and the probabilistic admissions control and retransmission process which together address the challenges of wireless communication as discussed in chapter 2, as well as meeting the goals for real-time communication as laid out in chapter 1. These protocols implement a real-time communications sub-system which provides time-bounded TDMA slot allocation, probabilistic admissions control and support for retransmissions to minimise communication failures while optimising the resources required.

This chapter will outline the experiments carried out, their goals, results and conclusions which can be drawn from them. Firstly the equipment used and configuration is outlined. The probabilistic admissions control and retransmission protocol as described in chapter 4 was evaluated using the HD-TDMA protocol as described in chapter 3.

This chapter is laid out as follows, section 6.1 outlines the experimental setup and equipment used together with the protocols used for evaluation purposes. Section 6.2 outlines the evaluation scenarios and assumptions, section 6.3 described the channel model employed during simulations. Section 6.4 presents results from the single bit error case, while section 6.5 presents results for scenarios incorporating both single bit errors and bursty error conditions. Section 6.6 evaluates the performance of the binomial clustered admissions control process. The evaluation concludes with section 6.7 which evaluates the performance of the admissions control process combined with the HD-TDMA MAC. Finally the chapter concludes with a summary.

6.1 Approach

In order to evaluate the performance of HD-TDMA and the probabilistic admissions control and retransmission protocol, a series of experiments were conducted employing the OPNET network simulator [OPNET Technologies, 2008].

As this thesis proposes a mechanism to overcome varying and error-prone medium conditions it is necessary to ensure that an accurate comparison is made between the evaluated protocols, such that the medium conditions experienced by each protocol are identical in time of occurrence, duration and severity. Such requirements rule out practical evaluation, as such repeatability is unachievable in the real world.

It is infeasible to capture the behaviour of the medium in the form of a trace to play back during simulation. To determine the impact of dynamic propagation effects from the environment would require a constant bit stream between each sender and receiver recording the start and stop times of errors, similar to the one way single transmitter receiver experiments performed by Willig et al [Willig et al., 2002]. However, such a process would make it impossible to analyze the return path or indeed the communication between any other surrounding stations, as communication is half duplex and all stations share the same frequency, thus only one station can transmit at the same time.

Based on the observations of Willig and Othani, [Willig et al., 2002,Otani et al., 1981] a mathematical model is used to model the medium using a long tailed distribution in the form of a Pareto distribution to model the duration of multiple bit in length error bursts. Single bit errors are modeled using the BER model. The occurrence of burst errors is modeled as a random process. A continuous trace file is generated for each transmitter-receiver pair for a range of error frequencies and durations achieved by varying the Pareto shape parameter and the frequency of burst error occurrence.

Should a frame be corrupted as a result of a collision between two or more concurrent transmissions, the simulator will automatically drop the frames involved at the effected receiving stations, otherwise the trace file is consulted to determine if the frame contains errors as a result of bursty error conditions and the specific location within the frame those errors are, a single bit error model is applied as the final stage.

6.1.1 Simulation Hardware & Software

The evaluation environment consisted of the implementation of the HD-TDMA protocol within the OPNET simulator. Version 14.5.A of the OPNET Modeler package for Linux was used. A dedicated server with a dual core 64bit Intel Zeon processor with a clock speed of 1.86 GHz and 8GB of memory was utilised. A GNU/Debian 5.0.2 Linux installation was installed running kernel 2.6.26-2. Version 4.3.2 of the GCC complier was utilised by OPNET. OPNET was configured to run all simulations in optimised mode with sequential execution with the default complier parameters. The kernel, OPNET and simulations ran in 64bit mode.

OPNET provides a distributed simulation capability allowing simulations to be distributed to a number of networked computers. For multiprocessor machines this capability extends to allowing separate simulation runs to run on each processor or core. This feature was exploited to allow two simulations to run on the server in parallel.

6.1.2 Evaluated Protocols

Five protocols where analysed as part of the evaluation. Two versions of HD-TDMA together with, CSMA, TDMA and TBMAC.

- **HD-TDMA**: Two distinct versions of HD-TDMA where evaluated. HD-TDMA utilised the per-destination deferral process and probabilistic admissons control as previously described in section 4.4.1, while HD-TDMA-I employed no deferral process or admissions control and thus maintained queue order during transmissions following a transmission failure.

- **TBMAC**: Time Bounded Medium Access Control, as proposed by Cunningham [Cunningham, 2003, Cunningham and Cahill, 2002] implements a TDMA system with dynamic and time-bounded membership services. No support is provided for acknowledgments or intelligent scheduling.

- **CSMA**: Carrier sense multiple access provides a baseline. CSMA is the fundamental channel access method in many technologies, including IEEE 802.11 and IEEE 802.15.4. The use of CSMA provides a baseline with which to compare a large volume of prior work. Acknowledgments are employed to maximise reliability.

- **TDMA**: Time division multiple access with no acknowledgments, featuring a fixed membership. TDMA in this configuration offers the maximum data throughout at the price of flexibility and reliability. Multiple variants of TDMA like approaches where discussed in chapter 2, such as FTDMA, RTMAC, MACA/PR and HCT-MAC. Under steady state conditions where each station has acquired the transmission resources it requires the classic fixed TDMA structure may be used to evaluate the performance of many TDMA like protocols.

6.1.3 General Configuration

Each of the evaluated protocols was configured to share the same inter-transmission gaps and the same number of slots where applicable. The default configuration consisted of 32 TDMA slots. With long inter-frame spaces (used between transmissions of different stations) set at 50μs and short inter-frame spaces of 10μs used between data and acknowledgment frames. Traffic is at constant bit rate (CBR), with experiments are performed under a range of different bit rates. Where the MAC protocol supports acknowledgments, acknowledgments are required.

In all cases worst-case conditions are applied, as all stations generate traffic at the same time and thus seek to access the medium to transmit at the same time resulting in worst-case contention. Thus in a data sink scenario, all the sources are seeking access to medium at the same time, resulting in significant contention, lost packets, transmission delays and dropped packets representing the most challenging conditions for CSMA in particular.

The OPNET simulator was configured to consider the simulation environment to be an 'office' location. The simulator used IEEE 802.11b physical layer characteristics, frequency, channel bandwidth and modulation to model the physical layer of the radio. All data transmissions occur at 2Mbits/sec in order to support the broadcast functionality which is provided by the IEEE 802.11 protocol. IEEE 802.11 channel 1, 2412 MHz was employed.

An overview of the slot and cycles times for protocols with TDMA like rounds is presented in table 6.1.

Protocol	Frame Size	Slot Length	Cycle Time
TDMA	64 bytes	385 μs	12.32 ms
	128 bytes	580 μs	20.16 ms
	256 bytes	1070 μs	35.84 ms
TBMAC	64 bytes	640 μs	20.48 ms
	128 bytes	835 μs	30.09 ms
	256 bytes	1325 μs	46.75 ms
HD-TDMA	64 bytes	5.1 ms	163.2 ms
	128 bytes	5.1 ms	163.2 ms
	256 bytes	5.1 ms	163.2 ms

Table 6.1: TDMA Slot Durations And Cycle Times

6.1.4 Protocol Specific Configuration

Specific configuration details are presented for the four evaluated protocols:

- **HD-TDMA**: The execution of the HD-TDMA MAC protocol relies on several parameters, the number of slots per cycle, the duration of the data transmission slot and the number of membership-management messages supported.

  As HD-TDMA supports variable packet sizes, the cycle duration and configuration was fixed for all experiments. The beacon frame contained provision to carry two membership-management messages. The beacon frame duration was 340 μsec, the data transmission slot as 4.7 msec allowing for a short frame space of 10 μsec between beacon frame and the start of the data transmission slot and a long inter-frame space of 50 μsec between the data transmission slot end and the start of the next slot. Full slot time was 5.1 msec, resulting in a 163.2 msec cycle time for a 32 data slot configuration.

- **TBMAC**: TBMAC was configured to have 32 data transmission slots, described as Contention Free Period (CFP) slots in TBMAC and two Contention Period (CP) slots, all slots have the same duration, which is dictated by the maximum data payload that may be transmitted. Similar to the HD-TDMA configuration there is provision for two membership-management requests per data slot.

- **TDMA**: Similar to both TBMAC and HD-TDMA, for each experiment scenario TDMA was configured to have 32 data transmission slots. These slots being sized to match the size of the data frames being transmitted, with the resultant impact on cycle duration. The slot and cycle durations are listed in table 6.1.

- **CSMA**: The CSMA parameters employed where the default IEEE 802.11b settings and are presented in table 6.2.

6.2 Scenarios & Assumptions

The OPNET simulator does not provide support to generate or to verify CRC values for frames. CRC's are not considered within the vendor supplied models, e.g. Ethernet.

Parameter	Value
EIFS	364 μsec
DIFS	50 μsec
SIFS	10 μsec
Slot Time	20 μsec
Retries	7
CWmin	15
CWmax	1023

Table 6.2: IEEE 802.11b DCF/CSMA Parameters

The default behavour of the simulator is to destroy packets with errors before they are received by each node, making the verification of a CRC value redundant.

While the default simulation configuration is such that any frame with one or more bits in error dropped. This may be reconfigured such that a frame with a number of bits in error up to a threshold may be accepted. However, the design and layout of HD-TDMA frames (section 3.4) is such that separate CRC calculations are required for the management and data sub-frames. The exact location of the bits in error being required to determine which portion of the frame has been damaged and therefore the functionality provided by OPNET was not deemed sufficient.

It is assumed throughout the simulations, that a 16 bit CRC strength will be sufficient to detect any combination of bit errors which occur within the portion of the packet upon which that CRC was calculated. In order to ensure this, the CRC polynomial must be chosen carefully. It may be appropriate to employ a different CRC polynomial for different data sub-frame lengths. Koopman et al [Koopman and Chakravarty, 2004], presents a detailed analysis of CRC performance for various frame lengths for various polynomials.

When a packet is received, the error models are applied. The implementation of these models is such that they are aware of the basic packet structure employed by HD-TDMA and will upon generation of an error(s) appropriately inform the MAC layer of the location of the error within the received frame. There are three possible valid outcomes:

- MGNT_CRC_VALID
- BOTH_VALID
- CRC_INVALID

The knowledge of the error location within the frame allows an appropriate CRC response to be provided. While a separate DATA CRC test exists, it cannot be valid if the `MGNT_CRC` check fails, as the management portion contains the length and transmission speed of the data sub-frame, without which the sub-frame cannot be decoded. In the following sections a number of scenarios are formulated and discussed with respect to the characteristics and behaviour of a variety of MAC techniques.

6.2.1 Scenario 1

Figure 6.1.1 represents a small sensor network consisting of eight sensors and one recording station. The sensors monitor a phenomena and notify the recording station, station A periodically of the value. Scaling the scenario to eight stations dramatically increases the risk of packet collisions and contention for medium access. Worst-case conditions are simulated by ensuring that all sensors will attempt to send a packet to the recording station at approximately the same time, resulting in an event storm.

CSMA is expected to perform poorly in this scenario as all eight sensors will be contending for the medium at approximately the same time. This competition for the medium results in packet loss through collisions. Multiple stations will see the medium busy at the same time and back-off multiple times resulting in a high latency of message delivery and poor utilisation of the medium which will spend a considerable time idle while back-off counters are decremented on each host.

TDMA like protocols should perform well in this scenario as they implement a time-triggered communication approach where data is only sent to the recording station at

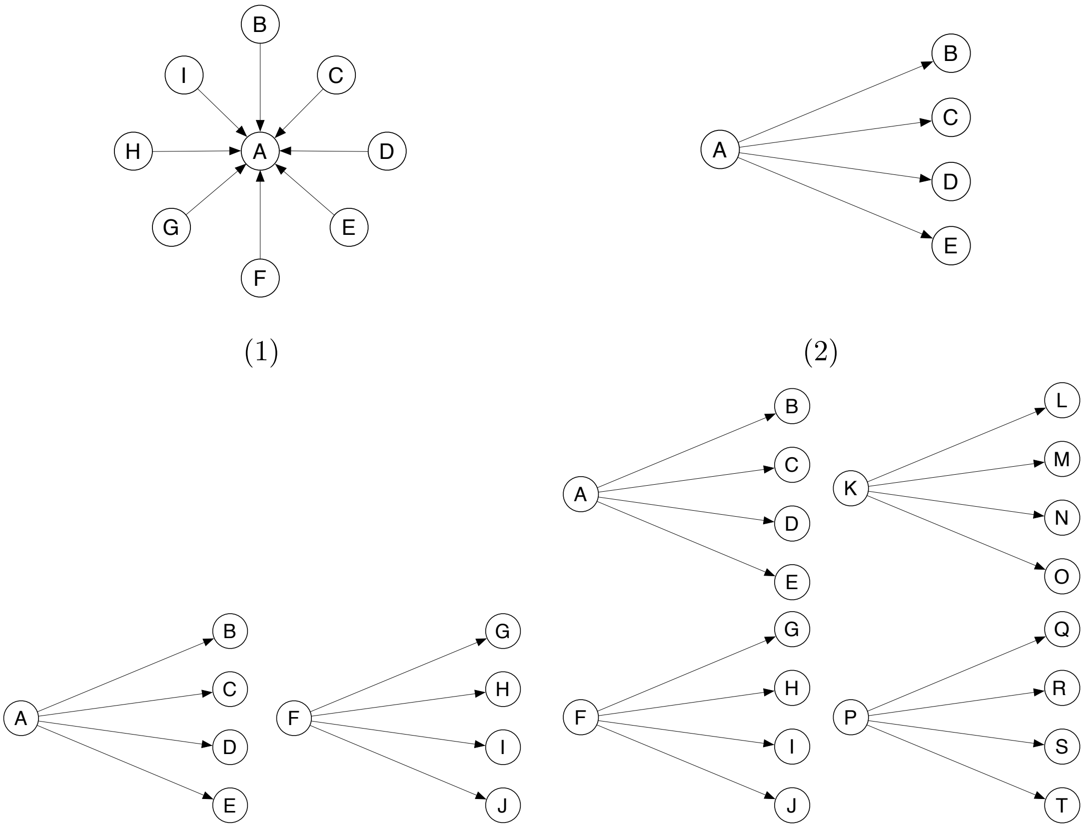

(1) (2)

(3) (4)

Fig. 6.1: Scenarios

specific well known times, i.e. the slots in which those stations have been allocated. This order comes at the price of latency as each station is forced to wait until its next allocated slot time. The TBMAC protocols results in the worst performance owning to the excessive latency it introduces owing to the slot configuration it demands. A key weakness shown by slotted protocols is that under certain emergency conditions it may not be acceptable to wait until the allocated slot to send the update to the recording station, only HD-TDMA supports this functionality through its ability to adapt the local transmission schedule.

6.2.2 Scenario 2

In the second scenario shown in figure 6.1.2, a single station is transmitting a stream of packets to four other stations. One or more of the communications links suffer intermittently from packet loss. This scenario aims to demonstrate the performance of the MAC approach in terms of minimising the impact of failures on communications on unaffected communications links.

When a transmission fails in the CSMA protocol, a retransmission is attempted immediately until either the packet is successfully transmitted or the maximum number of retries is exceeded. Under poor conditions e.g. bursty error channel this leads to poor channel utilisation. In this case the channel is blocked and the transmission queue on all stations begins to grow quickly as multiple attempts to send fail, resulting in head-of-queue blocking where forward progress is impossible until the packet at the top of the queue is removed. Using CSMA the medium may become idle for long periods owing to the randomly calculated back-off period and results in poor channel utilisation.

TDMA like protocols perform quite differently in this scenario, since channel failures are independent, packets to other stations suffer no impact or delay which is a significant improvement compared to the baseline CSMA and its head-of-queue blocking characteristic. In HD-TDMA-I, immediate retransmission is a poor choice under bursty error channel conditions e.g. fading there is a high probability given the medium is bad that it will remain bad but the communication link to other stations are considered statistically independent so if host A cannot communicate with B it is likely to still be able to communication with C, D, E and so on. Thus when unable to transmit to station B, transmissions to other stations could be made with the host transmitting the packets for B at a point in the future where it is deemed the channel will have recovered.

On demand allocation of slots while high latency will result in a high level of reliability at the cost of latency on the effected link. This will in many cases lead to a failure to achieve real-time deadlines. The localised protocol dynamically reorders transmissions to maximise the time before a second attempt to transmit is made and in doing so promotes the transmissions to other destinations avoiding the head-of-queue blocking and thus maximising the throughput of packets.

6.2.3 Scenario 3

The third scenario shown in figure 6.1.3, is two instances of scenario two. The presence of two transmitting stations A and F resulting in a potential for transmission collisions which where not present in scenario 2, these collisions will reduce overall CSMA throughput and increase transmission lattency as the expoential back off process increases the back off level due to a greater number of transmission failures.

The TDMA family of protocols retain the independence of transmissions and there scheduled transmission slots ensure, as in all scenarios that transmissions are contention free.

6.2.4 Scenario 4

The fourth scenario shown in figure 6.1.4 builds on scenario three by containing 4 instances of scenario 2. CSMA performance is likely to be hampered to a greater extent than in scenarios 2 and 3 as there are now four transmitting stations.

6.3 Error Model

The OPNET simulator lacks support for effective channel modelling, supporting only free space path loss with no channel fading model [Takai et al., 2001, Agba et al., 2008], whereas Ns-2 and GloMoSim both offer two ray ground models. While in many cases this would be a negative, for this thesis, it allows a channel model to be added into OPNET without conflicting with in built models. A pipeline stage may be added allowing an external source to provide channel propagation information similar to Agba et al [Agba et al., 2008].

In order to assess the response of various medium access control protocols to various error burst lengths and frequencies, it is necessary to generate a number of trace files which give the start and stop times of bursty conditions on the medium.

Previous experimental work [Otani et al., 1981, Willig et al., 2002] suggests that the burst length of an error follows a long-tailed distribution, therefore a Pareto distribution is introduced; to model the length of error bursts. A normally distributed random process models the arrival rate of bursts. The burst model is then combined with the single bit error random process, in order to model both single and burst errors. This approach may be thought of as a constrained Fritchman model [Fritchman, 1967].

Owing to the independence of transmissions, errors are modelled independently by each receiving station. All frames received by a station pass through the error model. The model incorporates memory in order to correctly simulate the duration of bursty errors which may effect multiple frames.

$$f(x) = \frac{kx_m^k}{x^{k+1}} \tag{6.1}$$

The burst duration is modeled using a Pareto distribution, a long-tailed distribution in line with observations by Willig et al [Willig et al., 2002]. The probability density function is given by equation 6.1 and shown in figure 6.2. The parameter x_m represents the minimum x axis value which is set to two bits: the single bit error case modeled by the BER model. The shape parameter, k is varied between 1 and 0.45, reducing this value resulting in increased error duration. The occurrence of burst errors was modeled as uniform distribution varying between 1.5 and 0.25 seconds. A fixed BER rate of 10^{-5} was used throughout.

In order to assess the impact of burst errors on the HD-TDMA deferral process, HD-TDMA is evaluated in two forms: HD-TDMA-I which in case of a transmission failure immediately retransmits a frame and HD-TDMA, which defers a retransmission in an attempt to avoid a burst error condition.

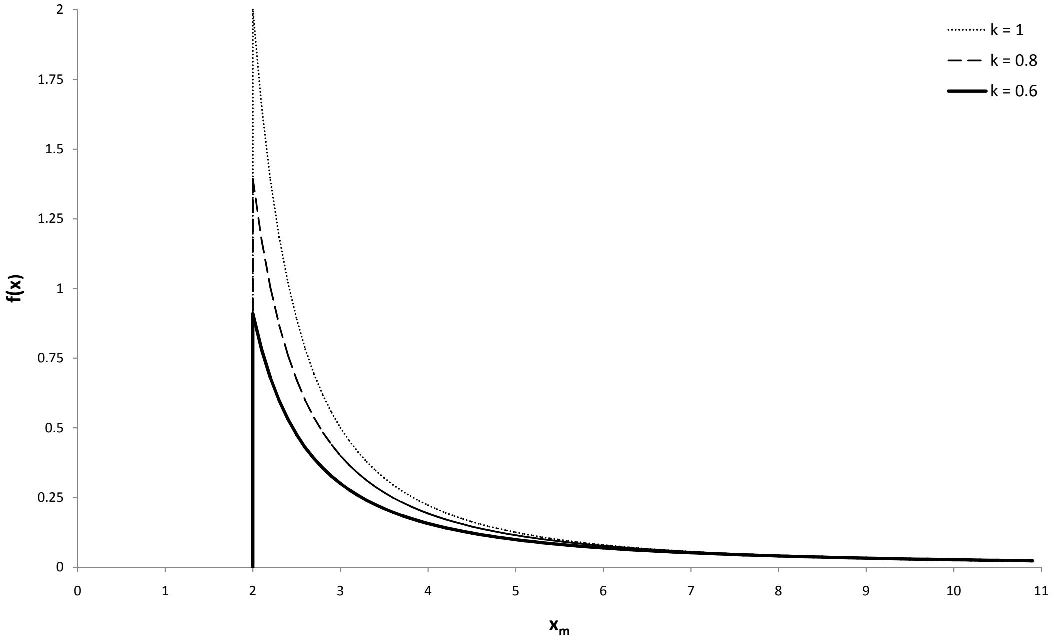

Fig. 6.2: Pareto Distribution

6.4 Under Single Bit Error Conditions

All TDMA-style protocols were configured to use 32 data transmission slots and to transmit 64 byte data payloads. In all cases, slots are allocated equally to transmitting stations. The exact slot durations and cycle times for HD-TDMA, TBMAC and TDMA are as shown in table 6.1. Table 6.3 presents the per station constant bit rate loads (CBR) loads. Load is defined to be the maximum theoretical capacity of the channel, which in this simulation is 2Mbps, achievable application data throughput will be lower at approximately 1Mbps owing to transmission overheads such as framing and inter transmission gaps. A real-time deadline of 50 msec is assigned to all frames upon generation.

Three performance cases are considered, firstly the plain transmission success ratio ignoring real-time deadlines, secondly the transmission success ratio considering real-time deadlines and finally the transmission success ratio following admissions control of

HD-TDMA. HD-TDMA admissions control will reduce the admissions rate in order to maintain transmission reliability.

In order to evaluate the frame delivery reliability of the individual protocols, the bit error rate (BER) is varied from 0 to 0.005 and the fraction of correctly completed transmissions is recorded. In the case of an acknowledged protocol such as CSMA, a successful transmission is one where the sender receives a valid acknowledgment frame.

All data frames sent with HD-TDMA, CSMA and TDMA are acknowledged. Two frames must be sent and received correctly, the data frame and the corresponding acknowledgment, doubling the number of overall transmissions. As two frames are required to complete each transmission, the probability of a failure, i.e failure to receive a valid acknowledgment is greater than in the unacknowledged case represented by TBMAC.

Scenario	Load	CBR Rate
1	25%	8 kb/sec
	37.5%	12 kb/sec
	50%	16 kb/sec
	75%	24kb/sec
2	25%	64 kb/sec
	37.5%	96 kb/sec
	50%	128 kb/sec
	75%	192 kb/sec
3	25%	32 kb/sec
	37.5%	48 kb/sec
	50%	64 kb/sec
	75%	96 kb/sec
4	25%	16 kb/sec
	37.5%	24 kb/sec
	50%	32 kb/sec
	75%	48 kb/sec

Table 6.3: Data Rates Per Transmitting Station For Presented Scenarios

6.4.1 Results

Figures 6.3 and 6.6 show that HD-TDMA can provide 100% successful real-time delivery in the presence of error rates under which other protocols fail to meet the delivery deadlines. Performance of all slotted protocols is virtually identical across all 4 scenarios owing to contention free transmission.

HD-TDMA incorporates an admission control and retransmission process, ensuring that time to retransmit frames is reserved in proportion to the estimated medium reliability, this combined with a low management overhead compared to TBMAC results in good performance. Figures 6.5, 6.8, 6.11 and 6.14 demonstrate how the admissions control process ensures that packets are admitted only when a high probabilty exists of successful transmission. As admissions control is applied at the time the packet is submitted for transmission, instant feedback is provided to the requesting application if HD-TDMA considers that it cannot meet that packets requirements.

TBMAC piggybacks membership-management data on all data frames, as a result these frames are larger than those used for data transmissions in TDMA and HD-TDMA and therefore are at greater risk of single bit errors. This is shown in each graph where TBMAC offers the worst performance of the slot protocols. This is most noticeable in heavy load conditions shown in figures 6.9 and 6.13 where for moderate error rates throughput is constrained as a result of TBMAC's significant protocol overhead not the state of the medium.

TDMA performance lacking the ability to make retransmissions results in degradation of performance faster than HD-TDMA, as shown in figures 6.3 and 6.6. Under heavy load, TDMA outperforms HD-TDMA as show in figures 6.9 and 6.12. This is princpaly a function HD-TDMA reducing its admission rate to ensure reliabilty of the admitted frames, figures 6.11 and 6.14 clearly show that the reduction in admissions allows HD-TDMA to provide reliable delivery.

CSMA in general, performs well due to a combination of retransmissions and the

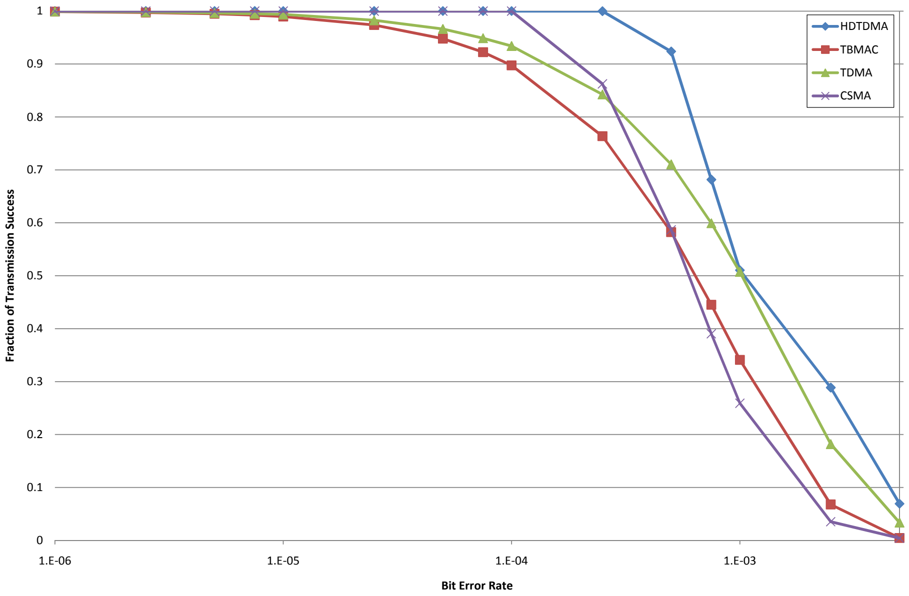

Fig. 6.3: Scenario 2 - 25% load without real-time constraints

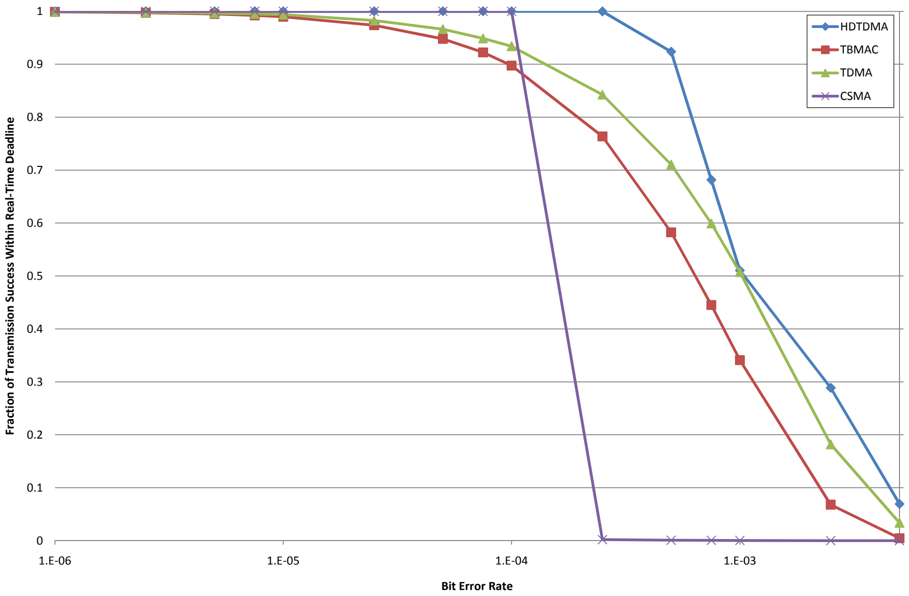

Fig. 6.4: Scenario 2 - 25% load with real-time constraints

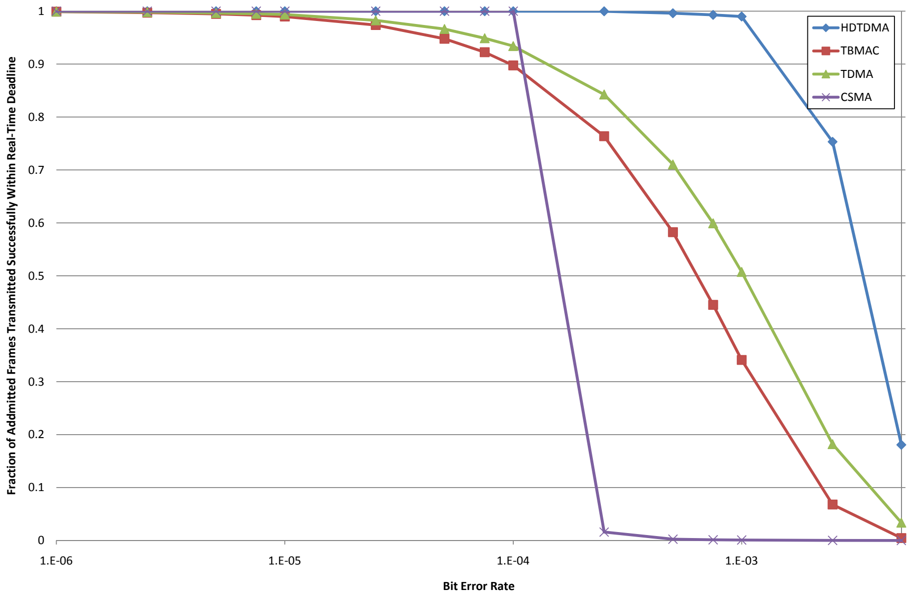

Fig. 6.5: Scenario 2 - 25% load with RT constraints + HD-TDMA admissions control

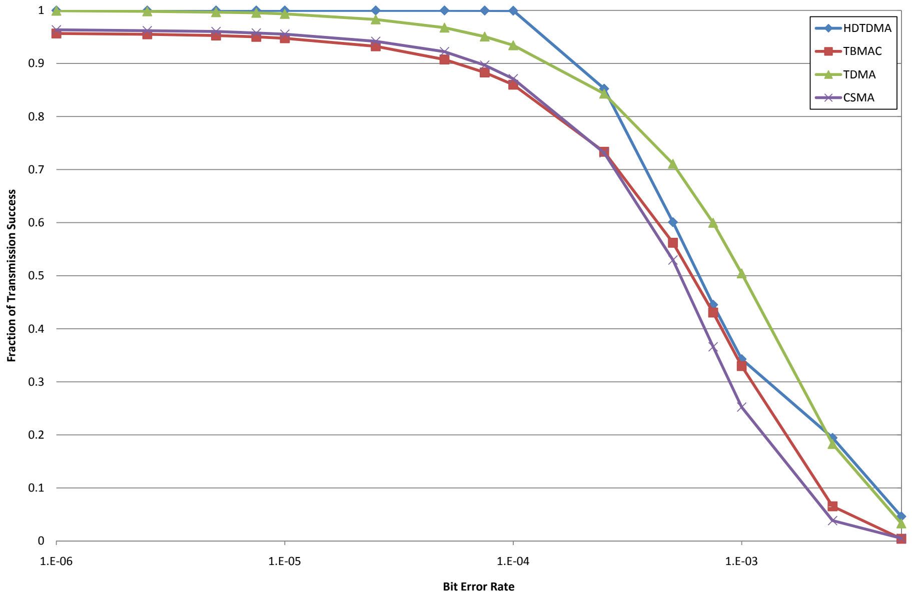

Fig. 6.6: Scenario 3 - 37.5% load without real-time constraints

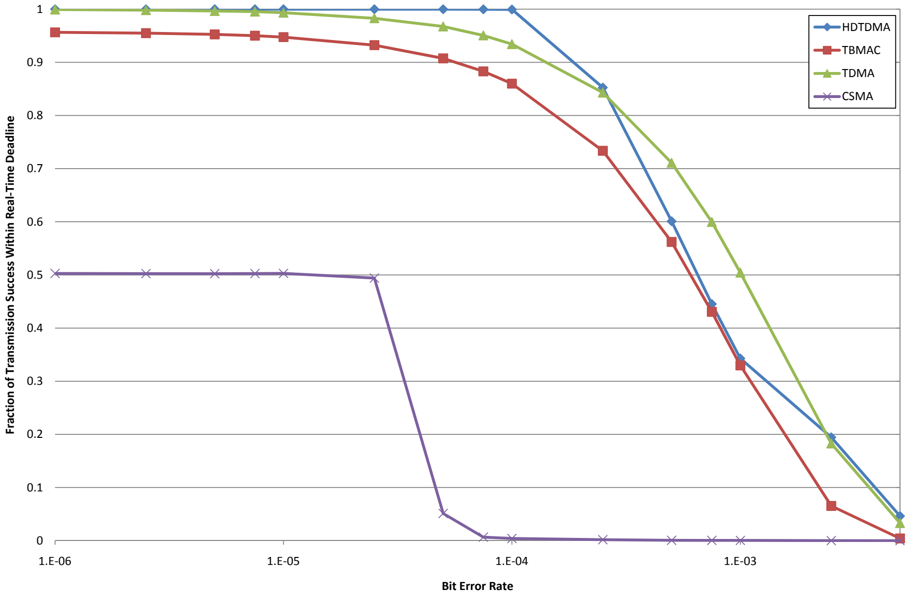

Fig. 6.7: Scenario 3 - 37.5% load with real-time constraints

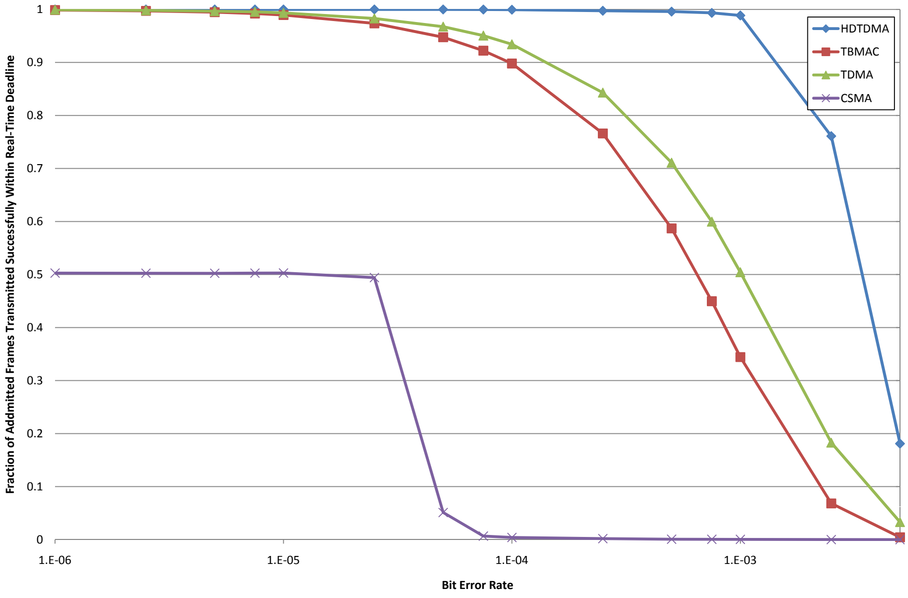

Fig. 6.8: Scenario 3 - 37.5% load with RT constraints + HD-TDMA admissions control

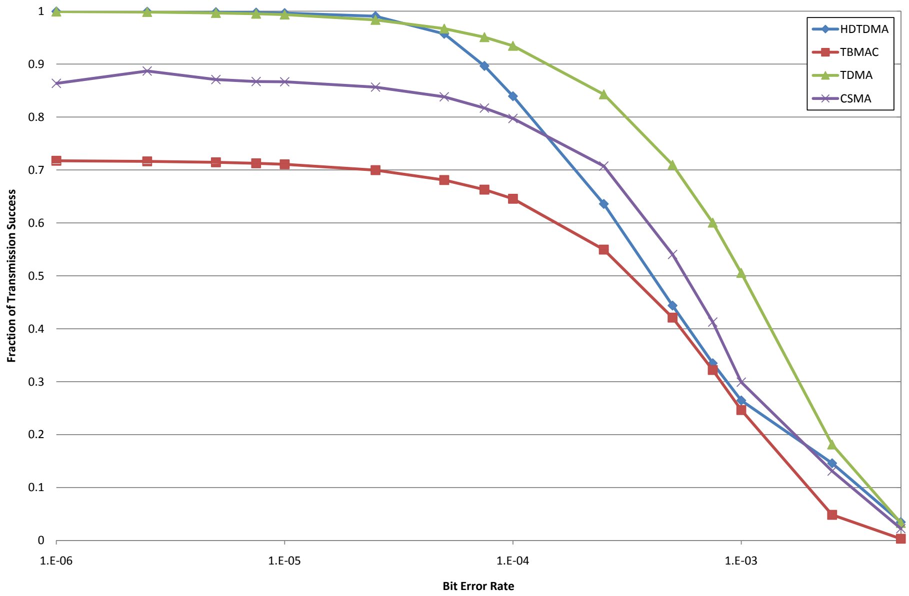

Fig. 6.9: Scenario 4 - 50% load without real-time constraints

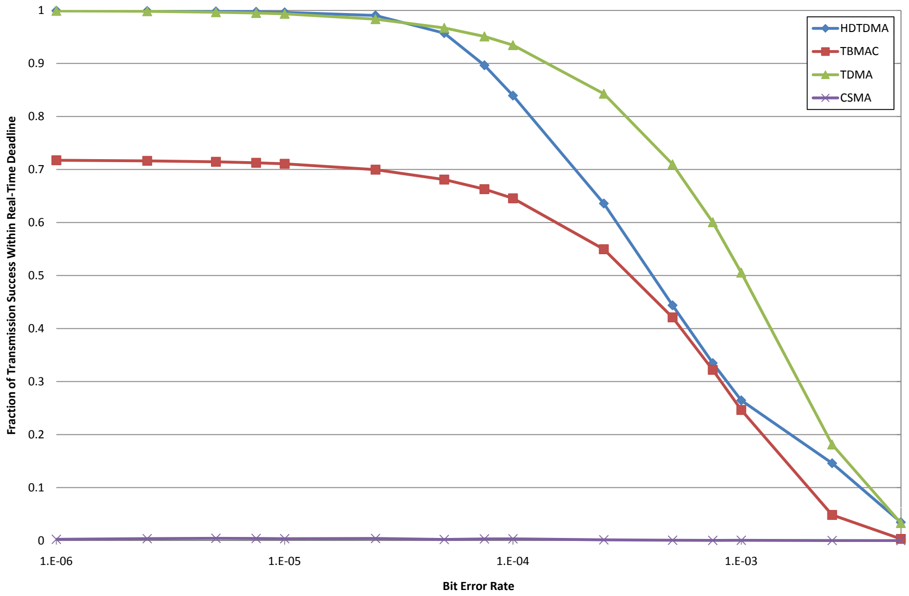

Fig. 6.10: Scenario 4 - 50% load with real-time constraints

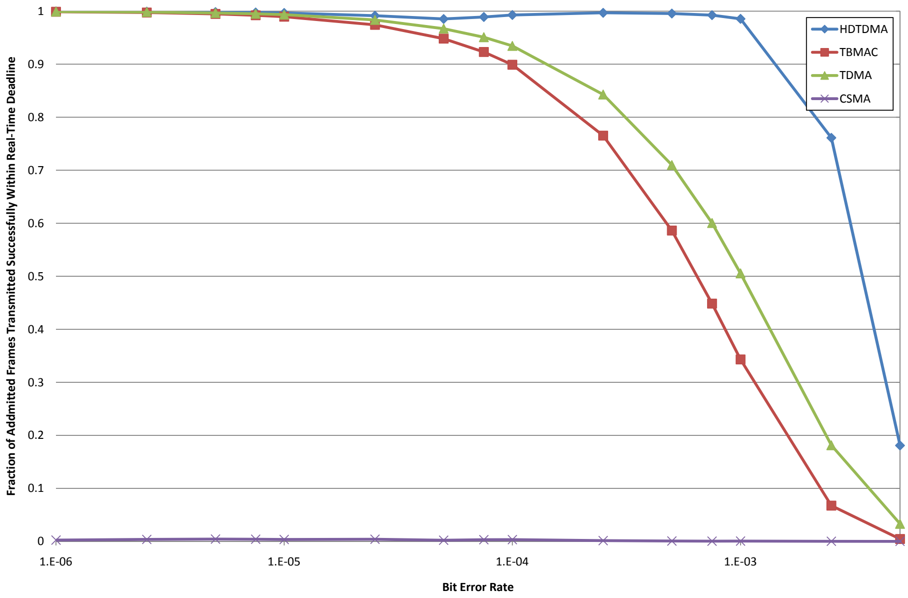

Fig. 6.11: Scenario 4 - 50% load with RT constraints + HD-TDMA admissions control

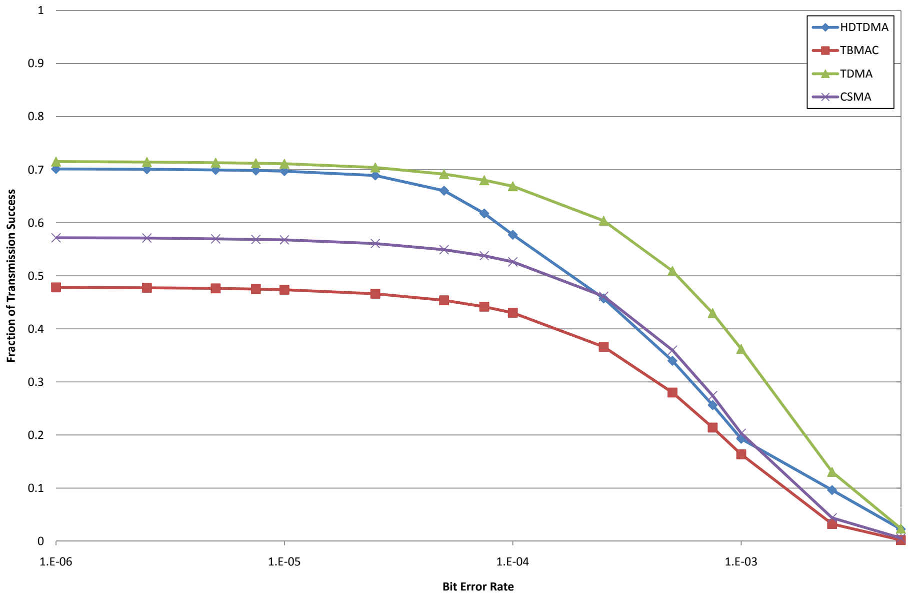

Fig. 6.12: Scenario 1 - 75% load without real-time constraints

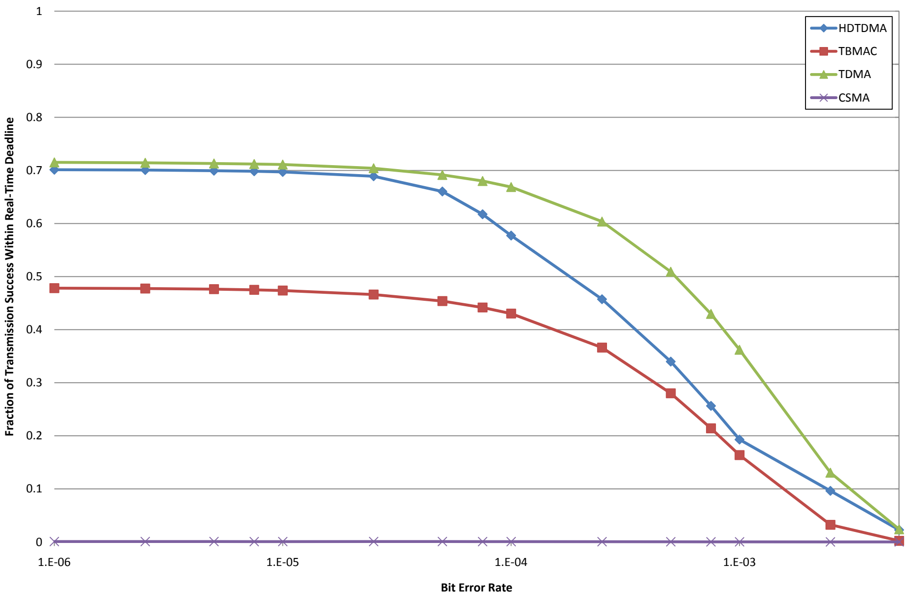

Fig. 6.13: Scenario 1 - 75% load with real-time constraints

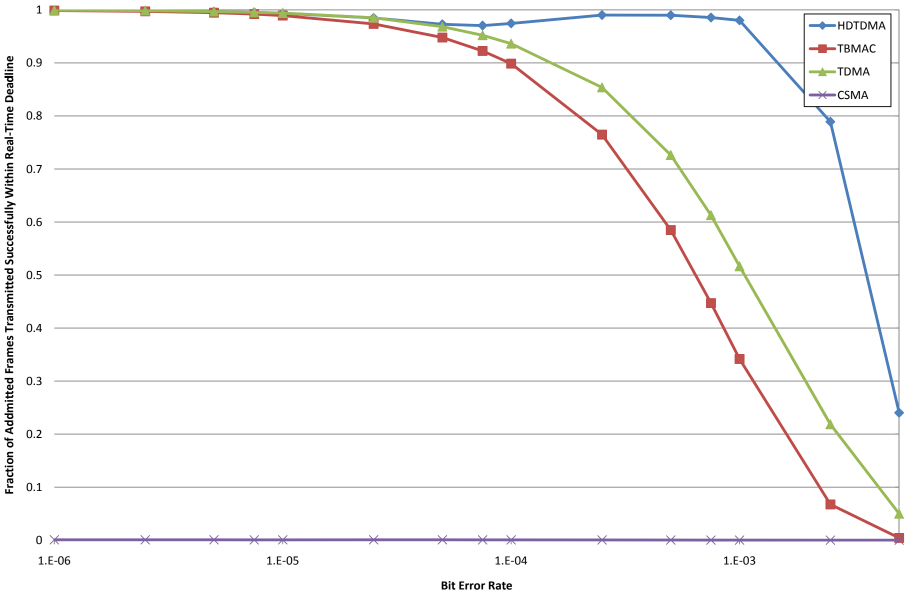

Fig. 6.14: Scenario 1 - 75% load with RT constraints + HD-TDMA admissions control

flexibility offered by not being constrained by slots. However, in all scenarios CSMA performance when real-time considerations are applied degraded at a faster rate than that of HD-TDMA, as figures 6.4 and 6.7 show. When a higher rate rate is present, such as in figures 6.10 and 6.13, real-time delivery under CSMA is virtually nil.

Evidence of capture is presented in figure 6.7 where CSMA successfully recieved approximately 50% of transmitted frames within deadline, in contrast to the near 100% delivery before real-time concerns are applied, as shown in figure 6.6. Analysis of the results obtained for figure 6.7 show that some stations had a high level of real-time delivery and others little or none suggesting channel capture effects due to high back-off counter levels.

This poor performance of CSMA may be attributed to the nature of the IEEE 802.11b DCF CSMA/CA protocol. If a packet is received from any station which fails the CRC test, the receiving station must wait for the medium to be idle for an EIFS interval before (re)commencing the back-off function, the EIFS interval remains in force until a frame with a valid CRC is received, combined with the need to wait for the medium to be idle for the duration of the back off counter results in head of queue blocking. With seven retires permitted and the data load close to or exceeding the maximum capacity of the channel due to addressing and control overhead, packets become delayed resulting in an inability to offer real-time performance.

6.5 Under Bursty Channel Conditions

As the propagation path between each pair of stations is independent, communication between stations A and B is considered different to communication between stations A and C but also consider that the channel is not symmetrical such that communication between A to B is distinct from communication B to A. Thus for n static nodes there are $n^2 - n$ independent transmission paths, each with a distinct error characteristic.

In order to assess the impact of burst errors on the HD-TDMA deferral process, HD-TDMA is evaluated in two forms: HD-TDMA-I which in case of a transmission failure immediately retransmits the effected frame and HD-TDMA, which implements probabilistic admissions control and defers the retransmission in an attempt to avoid a burst error condition and transmits frames to other destinations before reattempting transmission to the effected destination. A simplified admissions control process is used, in the case of HD-TDMA time was reserved to permit at least one retransmission attempt in each slot. The data rate and packet size parameters presented in table 6.3 where again used.

6.5.1 Results

Consistent with the results obtained under the single bit error case CSMA offered poor real-time delivery performance. CSMA performance is erratic especially when the results with and without real-time considerations are compared, such as figures 6.18 & 6.19 and again in figures 6.24 & 6.25.

In all scenarios the post admissions control results, shows that HD-TDMA by reducing the admissions rate provides the highest level of delivery, as shown in figures 6.17, 6.20, 6.23and 6.26.

For all scenarios, as shown in figures 6.16, 6.19 & 6.22 the use of the deferred retransmissions approach in HD-TDMA shows a significant benefit compared to instant retransmissions (HD-TDMA-I) thus validating the deferral approach in the presence of bursty errors. Performance in the case of scenario 1, showed little difference as the deferral process provides no advantage, as each station is transmitting to only a single destination. Thus the results for HD-TDMA-I and HD-TDMA are similar, but HD-TDMA offers greater throughput as it dynamically adapts the admission rate to match channel conditions, whereas HD-TDMA-I uses a fixed allocation permitting at least time for one retransmission per slot.

TDMA and TBMAC perform strongly for the presented scenarios owing the slot allocations, as slots are allocated on a round robin basis to stations when the system is

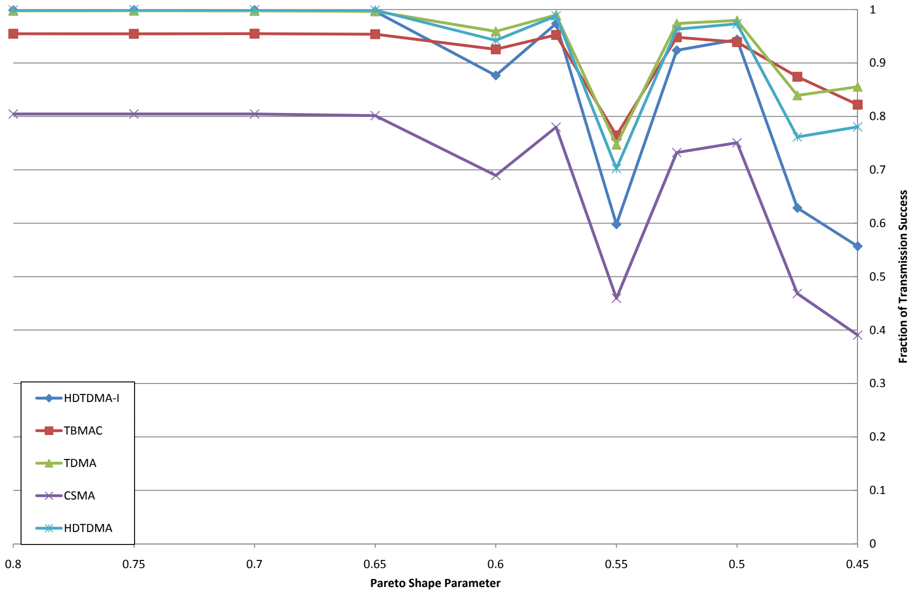

Fig. 6.15: Scenario 2 - 37.5% load, 0.625 sec burst interval no real-time constraints

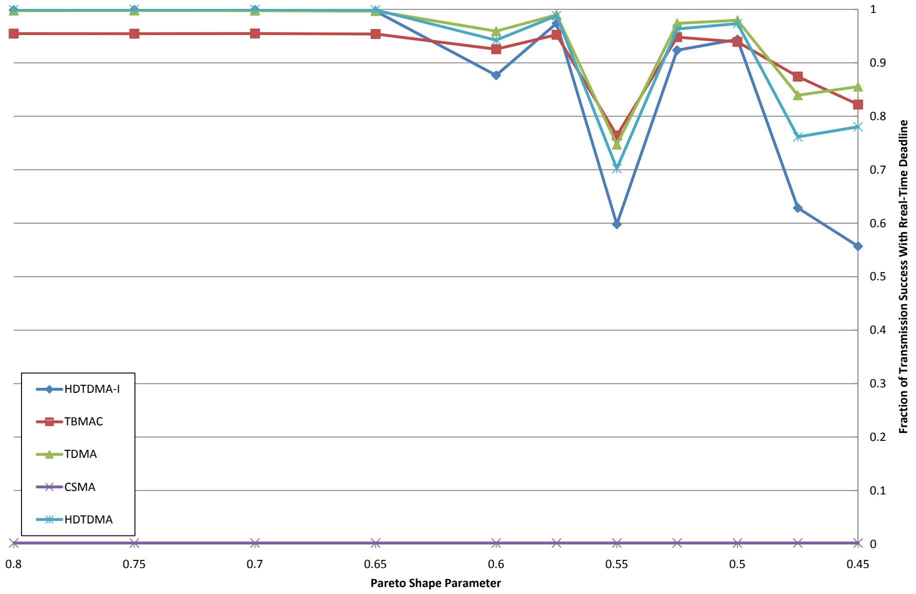

Fig. 6.16: Scenario 2 - 37.5% load, 0.625 sec burst interval with real-time constraints

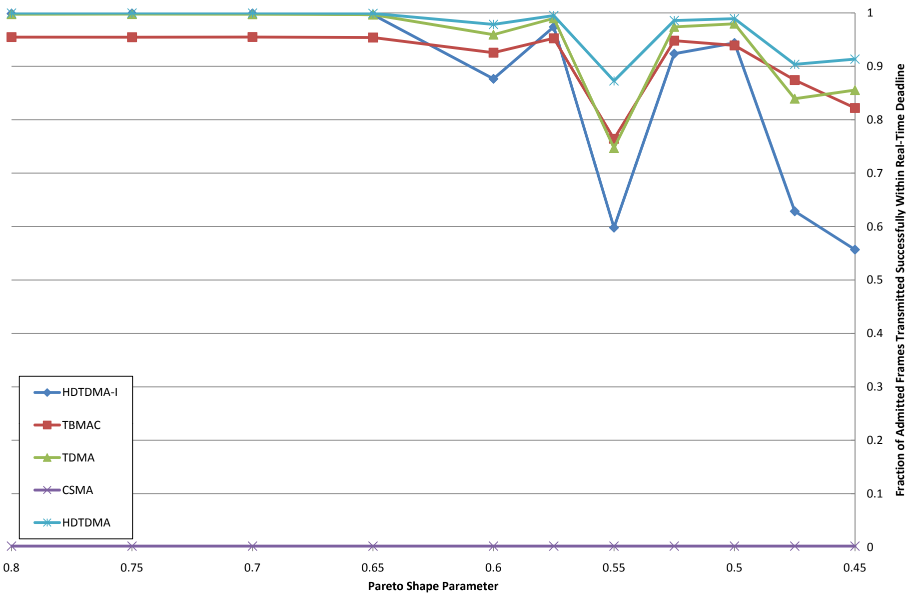

Fig. 6.17: Scenario 2 - 37.5% load, 0.625 sec burst + HD-TDMA admissions control

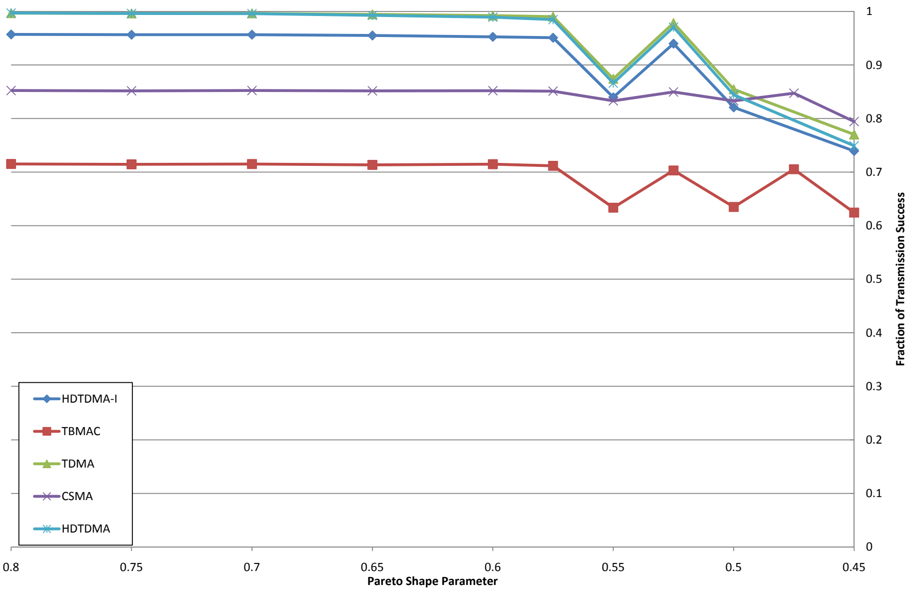

Fig. 6.18: Scenario 1 - 50% load, 0.375 sec burst interval without real-time constraints

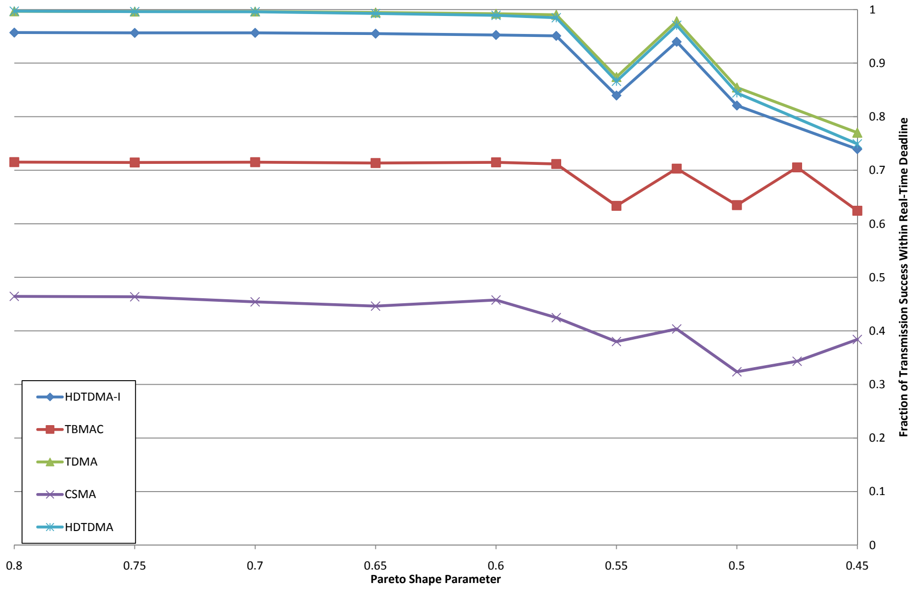

Fig. 6.19: Scenario 2 - 50% load, 0.375 sec burst interval with real-time constraints

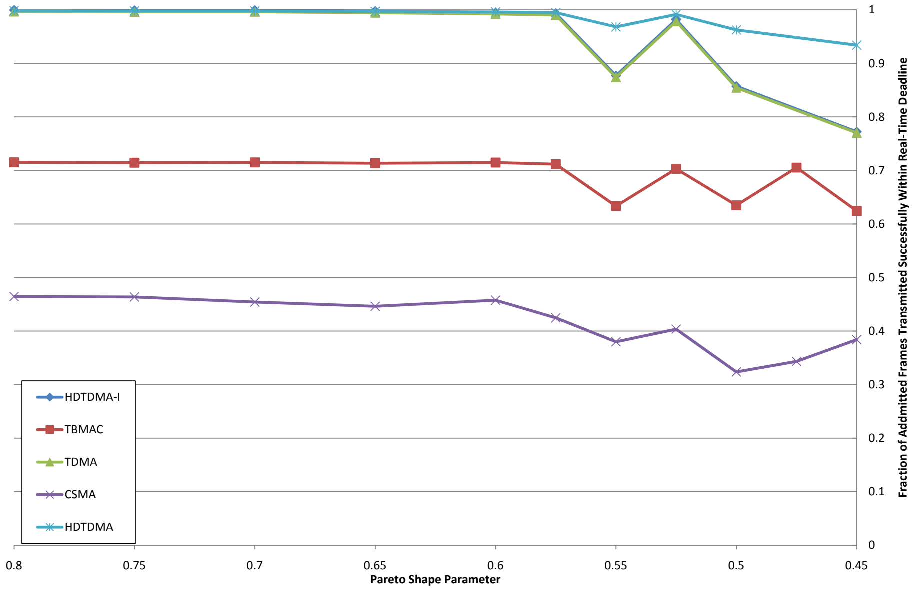

Fig. 6.20: Scenario 2 - 50% load, 0.375 sec burst + HD-TDMA admissions control

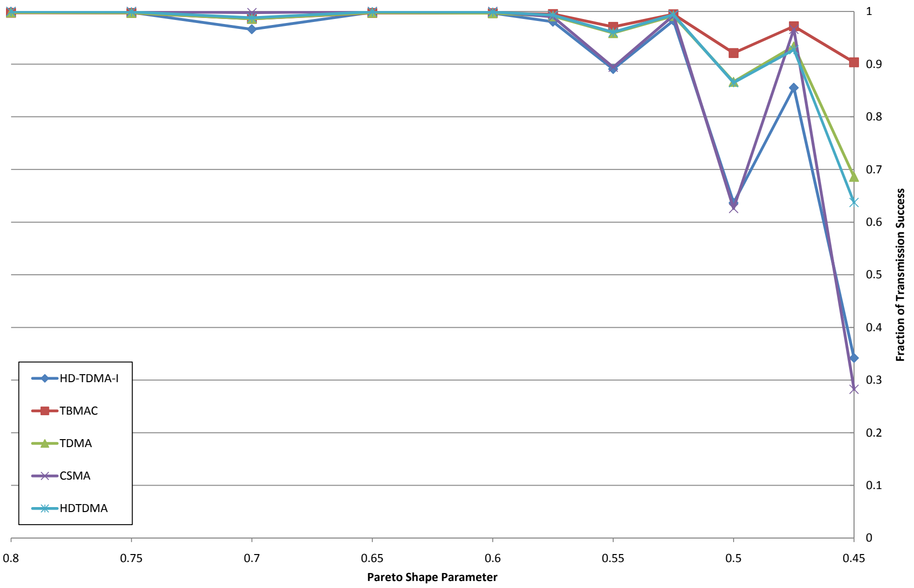

Fig. 6.21: Scenario 4 - 25% load, 0.75 sec burst without real-time constraints

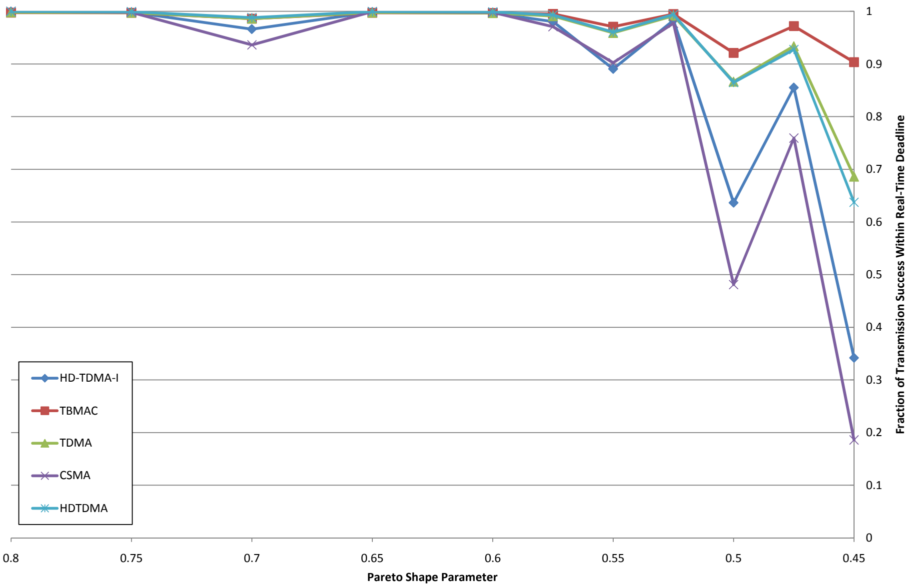

Fig. 6.22: Scenario 4 - 25% load, 0.75 sec burst interval with real-time constraints

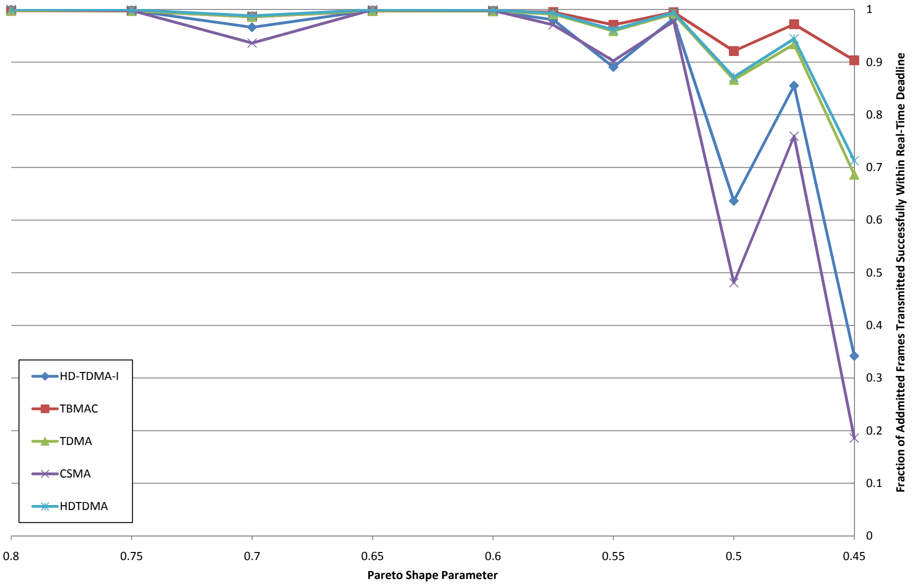

Fig. 6.23: Scenario 4 - 25% load, 0.75 sec burst + HD-TDMA admissions control

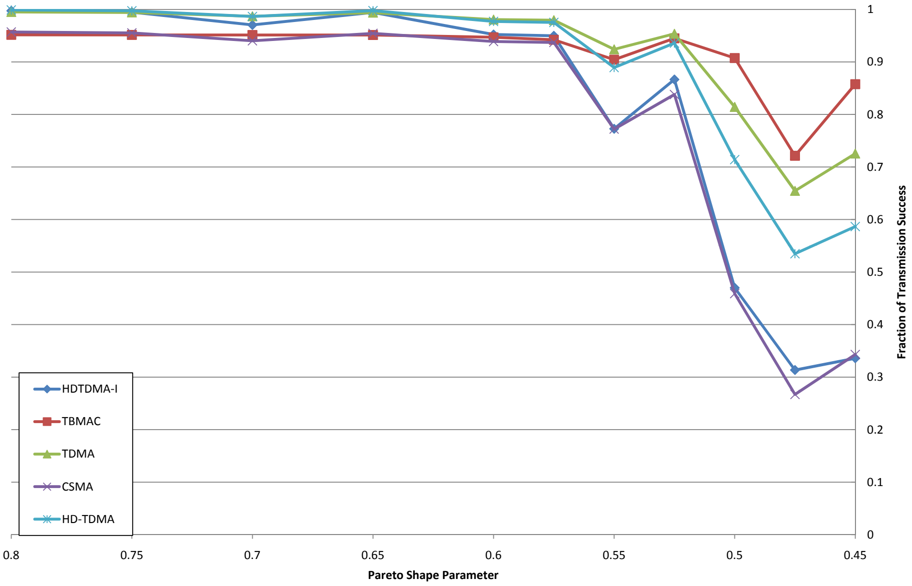

Fig. 6.24: Scenario 3 - 25% load, 0.25 sec burst interval without real-time constraints

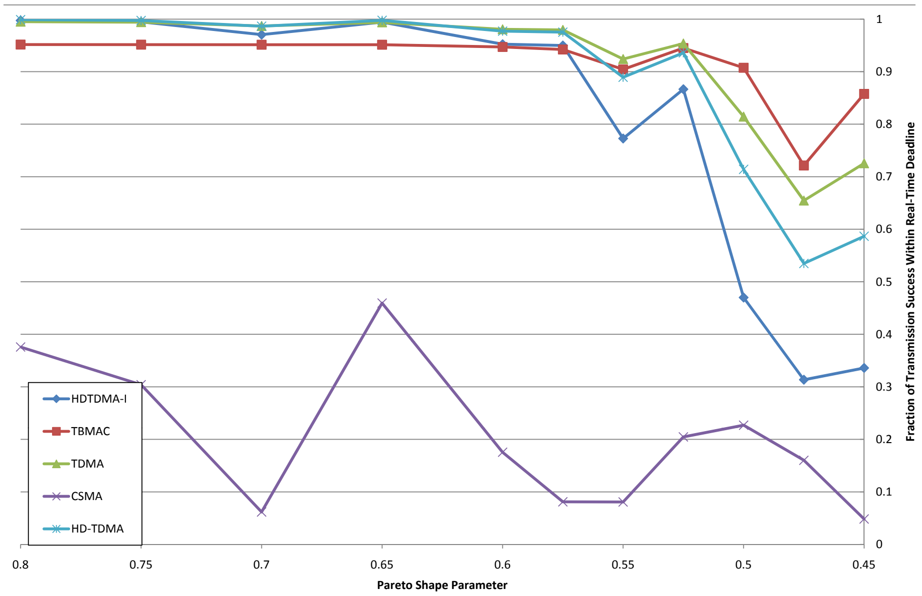

Fig. 6.25: Scenario 3 - 25% load, 0.25 sec burst interval with real-time constraints

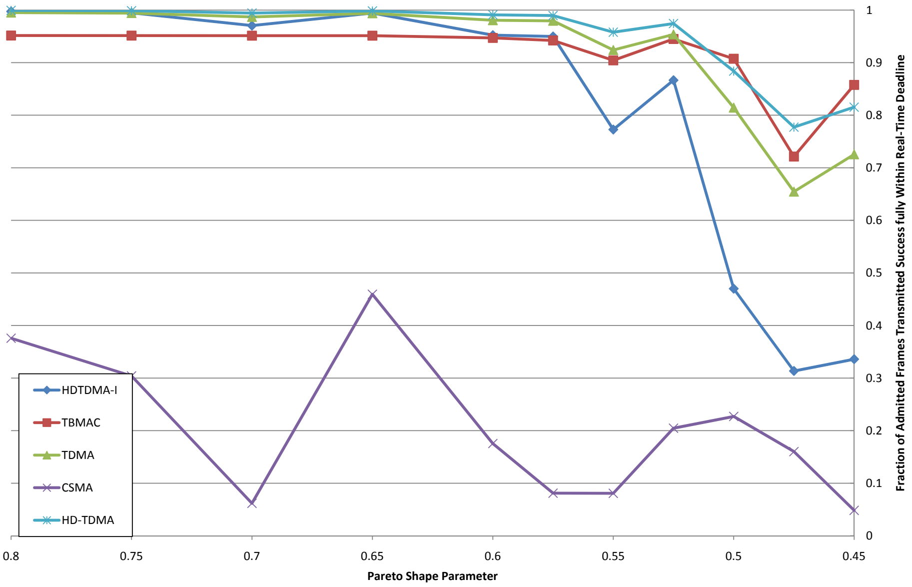

Fig. 6.26: Scenario 3 - 25% load, 0.25 sec burst + HD-TDMA admissions control

bootstrapped, no two adjacent slots share the same owner and thus are protected from back to back transmission failures. Neither protocol offers the admissions control and acknowledged data service of HD-TDMA.

TBMAC performance suffers principally under heavy load as the protocol overhead impacts on application data throughput as is shown in figures 6.21 & 6.24, in such cases HD-TDMA provides a greater throughput under the majority of error conditions.

In many of the graphs spikes exist where performance drops considerably only to recover, despite the fact that the error characteristic in terms of burst duration increases from left to right. Two factors explain this, firstly the trace files are generated off line using the error parameters as inputs to a statistical process, and therefore may result in more aggressive or more favourable error traces depending on the input seeds. A factor also is the alignment of errors with the slotted structure. Figures 6.15 & 6.21 particularly demonstrate this where the performance of HD-TDMA-I and CSMA suffer to a much greater extent than HD-TDMA.

6.6 Probabilistic Admissions Control Protocol Performance

Chapter 4 introduced the probabilistic admissions control protocol. Prior to the evaluation of the integrated HD-TDMA and probabilistic admissions control protocol it is appropriate to assess the performance and behaviour of the probabilistic admissions control protocol separately to determine the benefit it offers.

The evaluation compares the three options presented in section 4.2, where each transmission is considered separately, where transmissions to each destination are considered as a group and where a binomial clustering approach are employed. The scenarios consider between three and five destinations seeking to transmit up to three frames each. Table 6.4 lists the number of frames for each destination for eight scenarios.

Each scenario consists of 500 trials in which at the start of each experiment run, each destination is assigned a random probability of success between 0.45 and 0.95. The total number of transmissions required to meet the target reliability of 0.95 is then calculated.

Scenario	Dest 1	Dest 2	Dest 3	Dest 4	Dest 5
a	1	2	1	0	0
b	1	2	1	1	1
c	1	2	2	0	0
d	1	1	1	1	0
e	2	1	2	1	0
f	1	1	1	1	1
g	3	2	1	1	0
h	3	3	1	0	0

Table 6.4: Admission Control Scenario Parameters

A number of different clustering approaches are evaluated:

- **By Probability**: Clustering based on the destination pair in each round which had the least difference in probability.
- **By Variance**: Extending the first approach by seeking the cluster pair with the least variance in probability, this is weighted by the number of transmissions.
- **Exhaustive Search**: Ignoring probabilities entirely instead seeking to cluster destinations based on the reduction in transmissions achieved if that pair was clustered in that round, all destination combinations are considered.

6.6.1 Results

Table 6.5 presents an overview of the results from the eight evaluated scenarios, showing the total number of transmissions calculated from 500 trials. The evaluation shows that employing a clustering approach results in an reduction of between 13.24% and 24.93% in the evaluated scenarios, compared to considering transmissions independently and an reduction of between 6.02% and 13.24% compared to considering transmissions by destination alone.

Scenario	Individual	Destination	Clustered	Variance	Exhaustive	Reduction Individual	Best Reduction
a	6429	5910	5493	5449	5454	15.24%	7.80%
b	9594	9075	8094	7962	7973	17.01%	12.26%
c	8039	6986	6573	6487	6485	19.33%	7.17%
d	6429	6429	5691	5720	5715	11.47%	8.60%
e	9620	8578	7775	7767	7910	19.26%	7.47%
f	7975	7975	6919	6951	6984	13.24%	13.24%
g	11210	9572	8740	8716	8915	22.25%	8.94%
h	11210	8954	8429	8415	8572	24.93%	6.02%

Table 6.5: Admission Control Scenario Transmission Counts

Scenario		Individual	Destination	Clustered	Variance	Exhaustive
a	ave	12.83	11.79	10.78	10.89	10.96
	var	7.13	5.42	4.16	4.13	4.25
b	ave	19.15	18.11	15.89	15.91	16.15
	var	9.31	7.67	5.02	4.68	5.16
c	ave	16.05	13.94	12.95	12.94	13.12
	var	18.86	7.48	5.67	5.61	5.69
d	ave	12.83	21.32	19.52	19.48	19.93
	var	4.47	12.17	8.67	8.49	8.83
e	ave	19.20	17.12	15.52	15.50	15.78
	var	11.18	8.12	5.42	5.32	5.51
f	ave	15.92	15.92	13.81	13.87	13.94
	var	5.69	5.69	3.72	3.40	3.69
g	ave	22.36	19.10	17.44	17.39	17.79
	var	16.48	10.09	7.10	6.82	7.44
h	ave	22.38	17.87	16.82	16.82	17.11
	var	21.23	11.11	8.89	8.85	9.27

Table 6.6: Admission Control Scenario Average and Variance

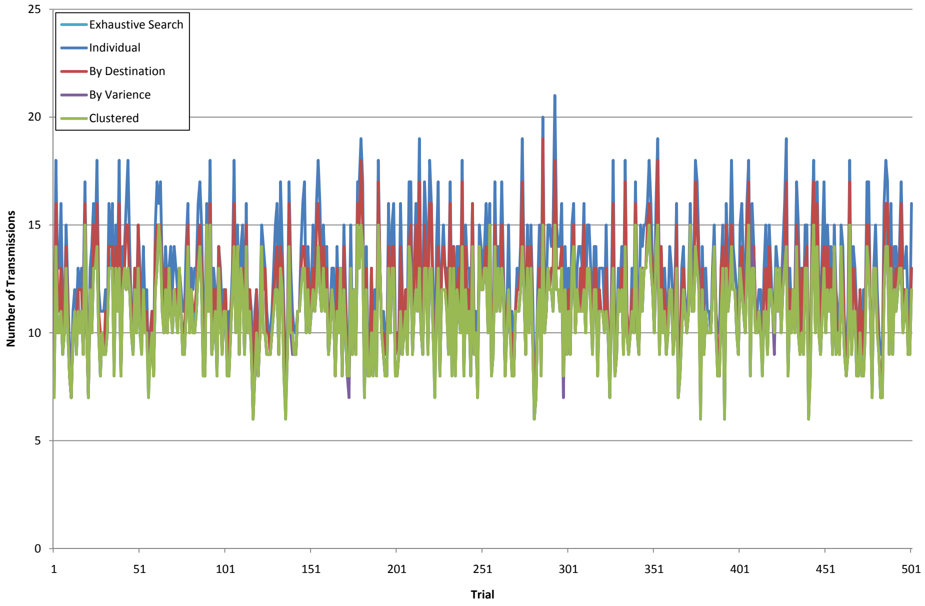

Fig. 6.27: Admissions Control - Scenario a

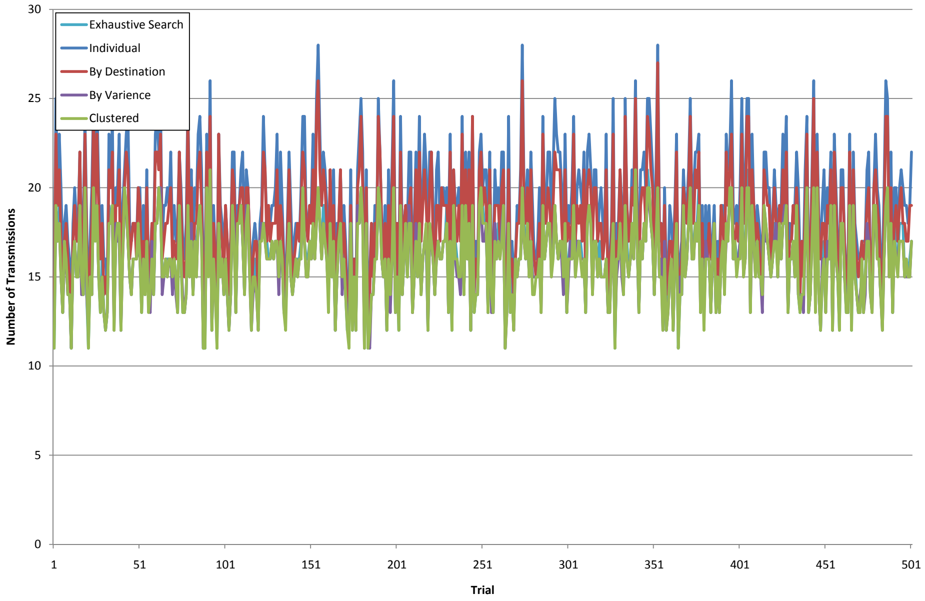

Fig. 6.28: Admissions Control - Scenario b

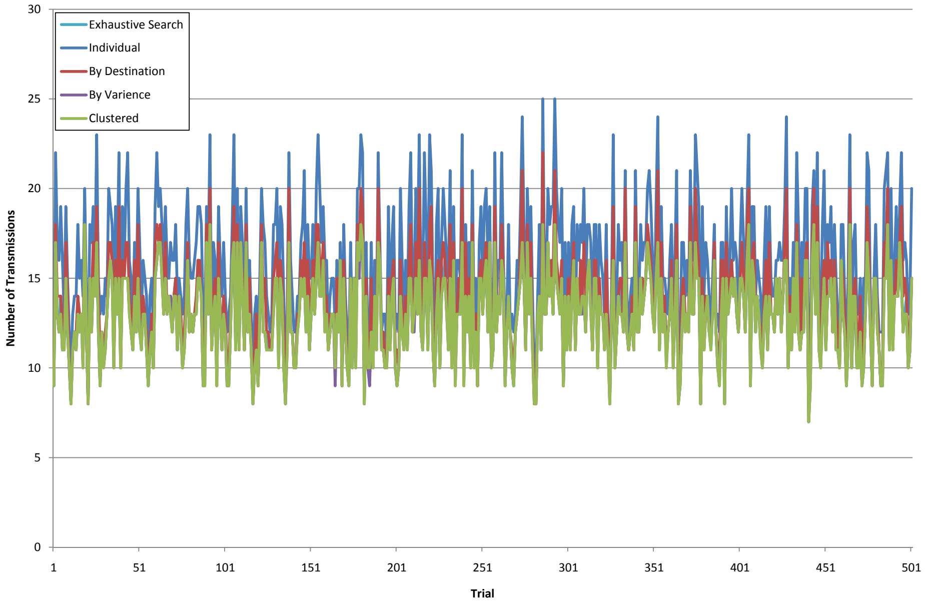

Fig. 6.29: Admissions Control - Scenario c

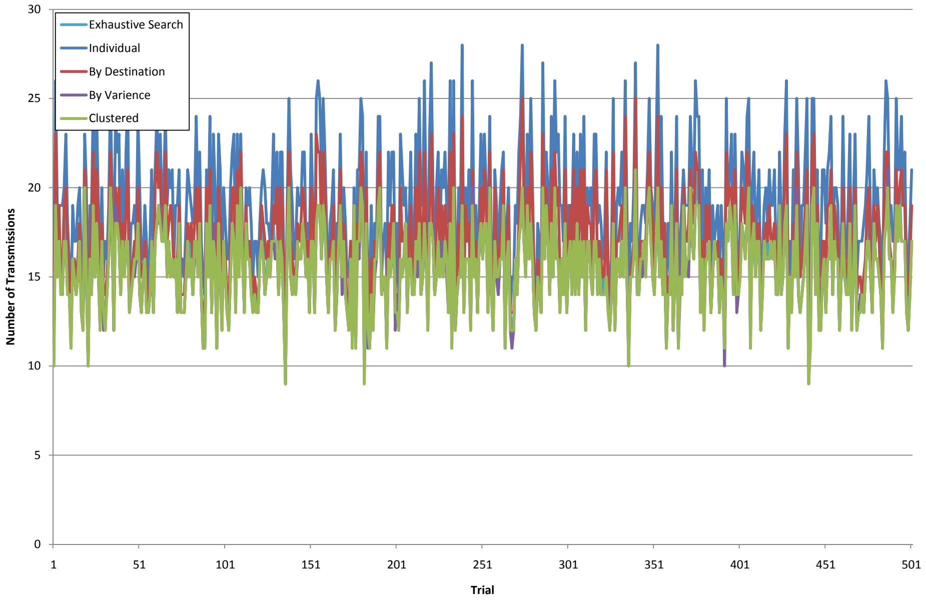

Fig. 6.30: Admissions Control - Scenario e

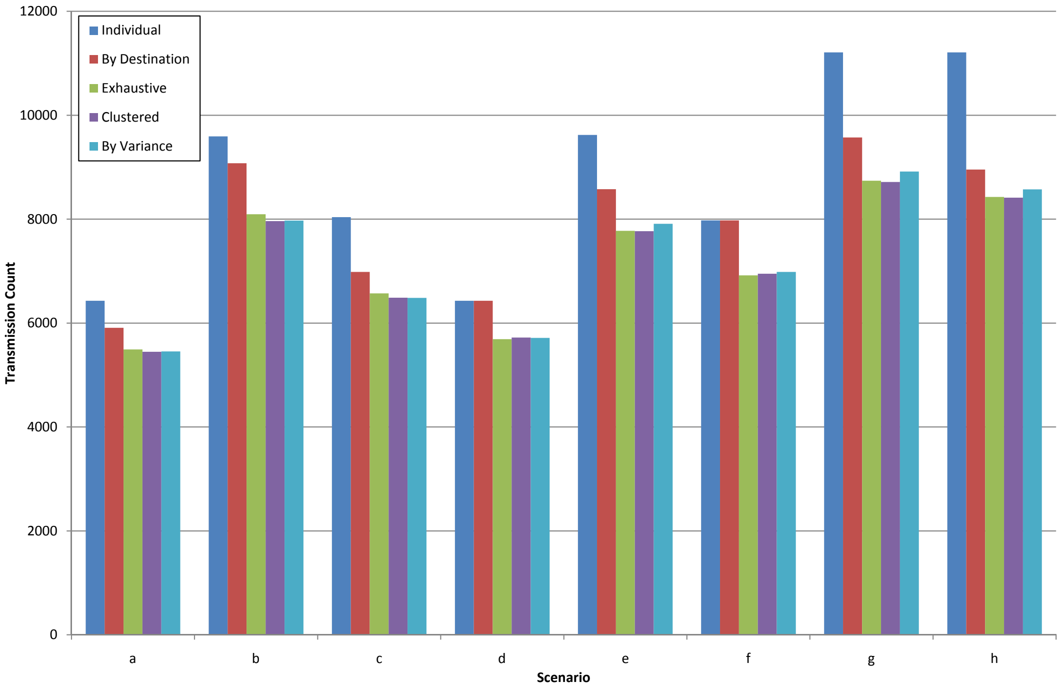

Fig. 6.31: Admissions Control - Overall Results

The inefficiency of considering each frame individually is shown by table 6.5, with the benefit offered by the binomial approaches clear, reducing the resource requirement by up to 25%. Figures 6.27-6.30 show that employing a clustered approach consistently provides an improvement over considering transmissions solely by destination.

Of the three clustering approaches, the exhaustive search approach consistently offered the least benefit, despite being the most complex algorithm. However, in the majority of cases it performance matched that of the other two clustering approaches at the price of significantly greater number of computations. In five of the eight scenarios the by variance clustering approach resulted in the best outcome, while in the remaining three the clustered by probability approach led to the best result.

The best outcome is achieved by performing both the by variance and clustered by probability approaches and selecting the algorithm which results in the least transmission requirements at run-time. Both approaches have low implementation complexity, requiring only a single pass of the queue, whereas the exhaustive approach implementation is similar to the classic bubble sort, requiring nested loops.

6.7 Integrated Protocol

The integrated protocol combines the functionality evaluated in sections 6.4, 6.5 and 6.6 in order to evaluate the complete real-time communications sub-system. The HD-TDMA MAC is evaluated with three distinct admissions control processes. Again a deadline of 50 msec applies to all frames.

The baseline, simple admissions control approach as used in sections 6.4 and 6.5, where sufficient time is reserved in each slot to guarantee the ability to retransmit at least one frame. This protocol is fixed and does not adapt to medium conditions. The second approach considers each frame independently, applying a Binomial distribution to determine the number of transmissions which may be required in order to achieve the target reliability. Frames are then allocated to slots which have sufficient time to accommodate the transmission of the frame and the estimated number of retransmissions.

The final approach, is the clustered admissions process in which transmission resources are pooled between destinations which share similar transmission success probabilities, thus reducing the overall time which must be reserved to carry out a group of transmissions. This allows a greater data throughput while maintaining real-time reliability and timeliness requirements. In light of the results of the previous section 6.6 both the clustered binomial approach and the by variance binomial clustered are both employed, with the protocol offering the least transmission time requirement chosen at run time.

To support the admissions control processes, an estimate of channel reliability is provided through a weighted average of reliability of communication to each destination over the last HD-TDMA cycle, with 50% assigned for the most recent third of a cycle,

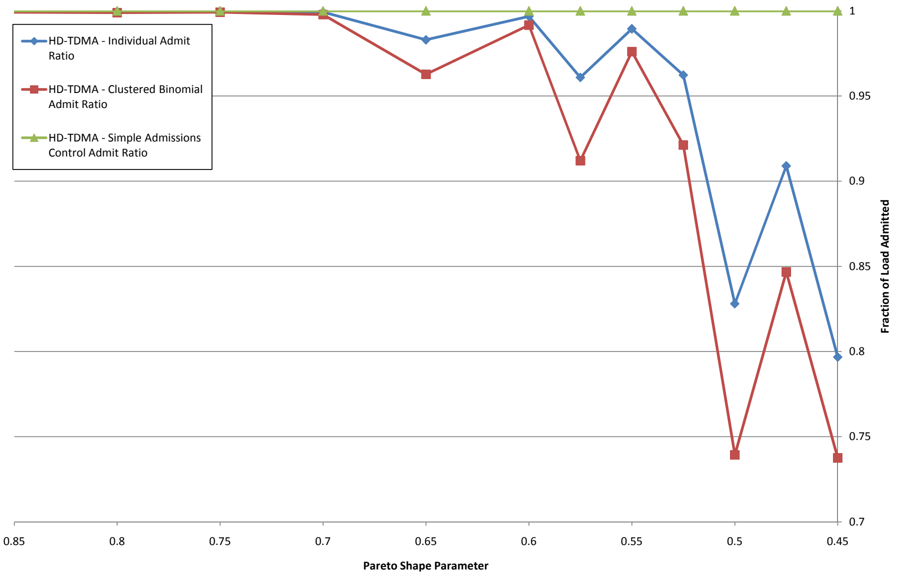

Fig. 6.32: Integrated Protocol - Scenario 2 - Admissions Ratio

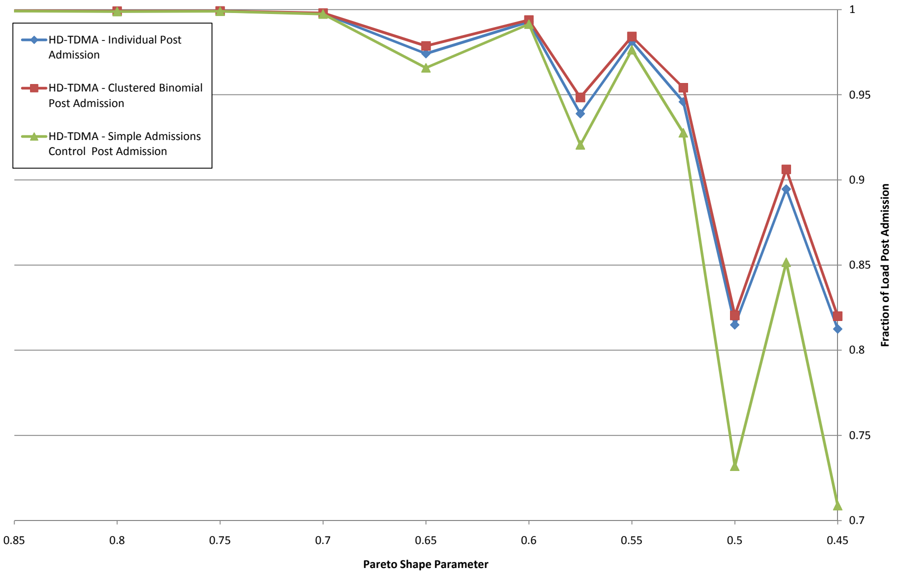

Fig. 6.33: Integrated Protocol - Scenario 2 - Transmission Success Rate

33.3% and 16.6% for the second and final thirds respectively.

In the previous evaluations, each communications link shared the same statistical properties defined by the rate of error burst occurrence and the distribution of the length of error bursts upon occurrence. This represented the worst case scenario where the clustered binomial approach would offer minimal reduction in as the clustering approach relies on the difference between probabilities of different destinations to reduce the overall result. A more realistic scenario is envisaged where each receiver/destination pair is assigned a statistically different error characteristic. For scenario 2 the channel error parameters are presented in table 6.7. As the value of the Pareto shape parameter is reduced, longer error bursts are present for that transmitter/reciever pair.

Error Set	Destination	Burst Frequency	Pareto shape
A	1	1.00	0.80
	2	0.75	0.65
	3	0.50	0.60
	4	0.25	0.45
B	1	1.00	0.80
	2	0.75	0.75
	3	0.50	0.65
	4	0.25	0.45
C	1	1.50	1.00
	2	1.00	1.00
	3	0.50	0.80
	4	0.25	0.45

Table 6.7: Channel Error Parameters

6.7.1 Results

Overall analysis of results suggests that there is marginal difference between the binomial based admissions control approaches. The clustered approach minimises the resources required based on an assumption of transmission independence between destinations and a deferral approach to prevent error bursts from effecting more than one frame to each destination. The binomial by destination approach provides a lower admissions rate owing to allocating a greater level of transmission resources, with a resultant capacity for a greater number of retransmissions.

There was little performance difference between the individual and clustered binomial approaches for the standard experimental configuration where each link shared identical statistical configuration. Where the clustering approach was employed, the number of admissions was reduced to provide a slightly higher delivery ratio of those admitted as may be seen in figures 6.32 and 6.33.

The individual approach admits all frames as it reserves only time for a single retransmission per slot, which leads to poorer performance in post admissions delivery when the channel error rate increases. The admissions rate for both the individual admissions ratio (blue) and the clustered binomial admission ratio (red), varies as a function of the observed channel conditions and the mathematical approach employed in determining the optimum number of transmissions.

Following the application of the admissions control process, the post admission real-time delivery rate is considered, as shown in figure 6.33. The performance of the dynamic, grouped approaches of binomial by destination compared to the simple admissions process shows a significant increase in real-time delivery as a result of the provision of transmission time to enable restransmissions.

The second experiment considers a more realistic case where each link is assigned a different statistical error characteristic. Performance of the clustered approach shows greater benefit as each link as a different transmission reliability which allows the clustering algorithm to provide a reduced transmission count.

In these scenarios, the error characteristic is fixed in line with the parameters set in table 6.7. As shown in figure 6.35, as the load increases the admissions ratio decreases as the various admissions control approaches restrict admissions in order not to exceed

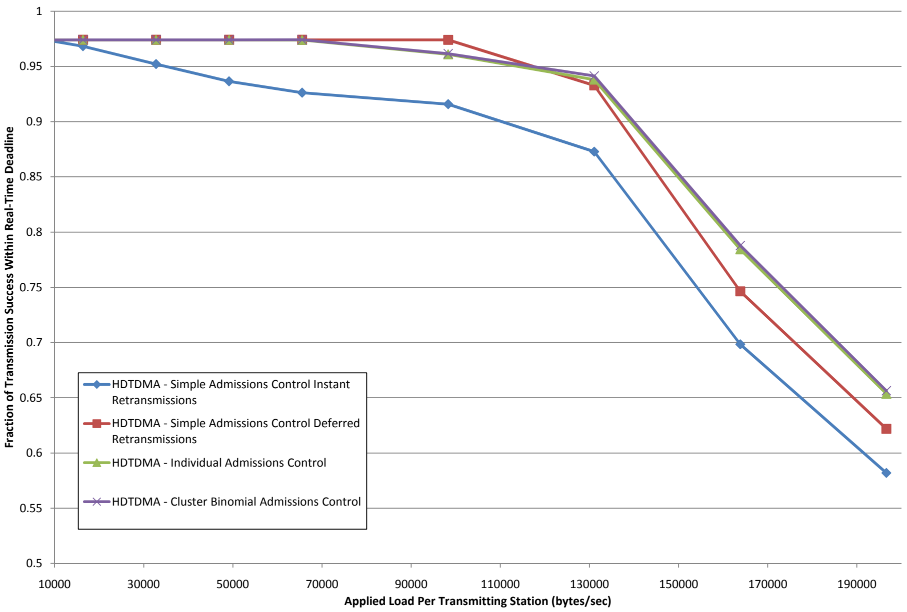

Fig. 6.34: Error Set B - Transmission Reliability Before Admissions

available channel capacity as to minimising queuing delays while ensuring sufficient time is reserved to allow for retransmission owing to the channel error state..

The simple admissions control process ensures that time to transmit at least one frame exists in each HDTDMA slot, thus under low load more than one retransmission will be possible. In admissions ratio terms both simple admissions control with instant retransmission and simple admissions control with deferred retransmission will offer the same admissions rate, as shown in figure 6.35 but will offer different transmission reliability as shown in figure 6.36. The deferred approach offers improved performance by not performing the retransmission instantly in order to minimise the impact of burst errors of duration greater than a frame length.

As shown in figures 6.34 and 6.36 both the clustered and individual binomial admissions control approaches offer the best performance under heavy load, exceeding the performance of the simple admissions control processes. The performance difference between the individual approach and the clustered approach is marginal but the clustered approach offered a higher admissions ratio and a higher throughput of frames in both pre and post admissions results.

Under low load the simple admissions control processes have a minor advantage by always reserving sufficient time to make at least a single retransmission where as the binomial approaches will reserve zero additional time if the estimate of channel reliability is greater than or equal to that of the frames required reliability. This could be compensated by modifying the binomial approach to ensure that time exists for at least a single retransmission regardless of the assumptions made of channel reliability.

6.8 Summary

This chapter presented the evaluation of the HD-TDMA together with the probabilistic admissions control and retransmission protocol. Each element of the communications sub-system was evaluated separately before a evaluation of the fully integrated communi-

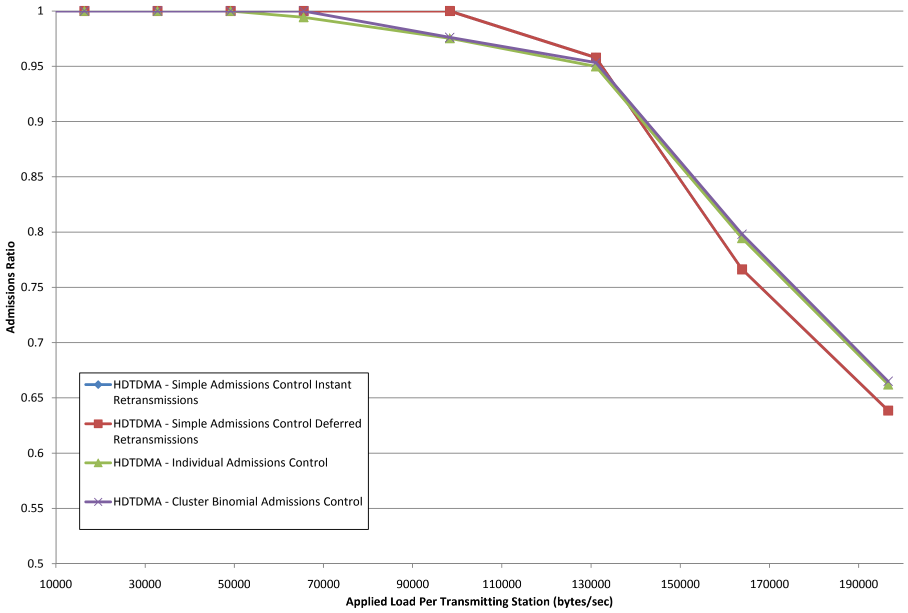

Fig. 6.35: Error Set B - Admissions Ratio

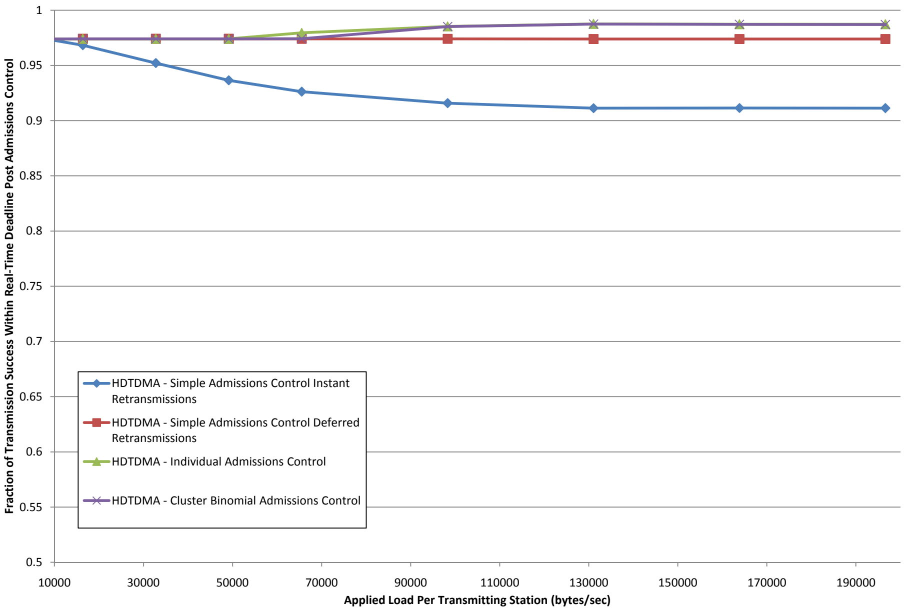

Fig. 6.36: Error Set B - Transmission Reliability Post Admissions

cations protocol was performed. The results from several sets of experiments conducted in order to evaluate the performance of the HD-TDMA MAC as well as the related admissions control and retransmission protocol were presented. Each set of experiments builds on the previous scenario until finally all the complete real-time communication sub-system was considered.

The HD-TDMA protocol supported by a localised autonomous decision making process provides a balance: The slotted structure ensures that transmissions are independent thus failure of packet transmission on one link does not impact on others. It is possible to dynamically reorder the transmission schedule in each slot to minimise the impact of bursty error conditions which in turn reduces the overall number of retransmissions. It is also possible to support a small number of retransmissions or sporadic messages.

The first set of experiments detailed in section 6.4 focused solely on response to single randomly distributed bit errors. HD-TDMA outperformed all protocols under all scenarios and applied loads. Section 6.5 repeated the evaluation with the channel now suffering from a combination of both bursty and single bit errors. HD-TDMA offered the best performance of all protocols supporting acknowledgements.

The analysis of the probabilistic admissions control protocol performance in section 6.6, demonstrated a clear and consistent benefit through the application of a clustered binomial distribution to determine the number of transmissions required for a group of transmissions while meeting the real-time reliability requirements of those frames.

Under aggressive channel conditions in particular, the individual admissions control approach offers marginally better delivery rate at the price of lower admission rates. By allocating more time on average to each transmission than the clustered approaches, the individual admissions control approach offers a greater capacity to cope with unexpected drops in communication. The key outstanding challenge is determining an accurate estimate of the channel reliability, this effort is frustrated by the very nature of the channel. A bursty channel is said to constant and reliable with relatively little bad bursts [Wang et al., 2007] short or long term channel statistics suggest a high level of channel reliabilty, with details of bursty periods lost.

Chapter 7

Conclusions and Future Work

This thesis discussed the issue of supporting firm real-time communication in wireless networks through a greater understanding of the nature of the behavour of the physical communications medium enabling the development of a medium access control protocol which incorporates awareness of medium characteristics.

This thesis presented the Hierarchal Distributed Time Division Multiple Access (HD-TDMA) protocol together with a probabilistic admissions and retransmission control protocol, which combined address the challenge of providing firm real-time communications support over wireless communication links with time varying reliability.

This chapter reviews the significant contributions of the work described in this thesis and reviews the capabilities of HD-TDMA and the associated probabilistic admissions control and retransmission protocol with respect to the requirements set out previously in chapter 2. The chapter concludes with a brief discussion of issues which may form possible future works.

7.1 Summary

The HD-TDMA protocol is based on a cellular wireless structure in which each TDMA slot may contain a number of transmissions. Each allocated slot begins with a beacon-frame containing information concerning the current slot allocation state and proposed future updates. Dynamic cell membership is supported through a time-bounded fault-tolerant membership management service.

A combination of both flexible and fixed structures supports communication. A fixed length TDMA cycle with a fixed number of slots, but where each HD-TDMA slot permits a station to make multiple transmissions, supporting both acknowledged and unacknowledged communication as well as a polled transmission mode to overcome the scalability issues inherent in TDMA.

Acknowledging the variability of transmission reliability, a combination of a probabilistic admissions control and adaptive scheduling was employed. Frames are assigned to slots to ensure transmissions occur within their real-time deadlines, while also ensuring that time has been reserved to allow for possible retransmissions. A clustering algorithm employing binomial distribution is employed to estimate the number of transmissions required to satisfy the real-time reliability requirements of each frame. The clustering approach reduces the estimated number of retransmissions and in doing so reduces the transmission time which must be reserved for retransmissions, while continuing to meet real-time reliability requirements.

7.2 Contribution

The contribution of this thesis is threefold: the Hierarchal Distributed TDMA protocol, the probabilistic admissions control system and the localised scheduling process to react to transmission failures.

7.2.1 Hierarchal Distributed TDMA

HD-TDMA represents a new approach in wireless TDMA style communication by providing support at the MAC level for retransmissions through the ability to make multiple transmissions per slot. This recognises the inherent variability in transmission reliability in the wireless domain. In addition to addressing the need for transmission reliability HD-TDMA provides a flexible transmission sub-system permitting multiple transmissions per slot, flexibility in frame size, acknowledgments as well as a polled transmission mode.

Support for propagation of TDMA slot allocation cycle data and membership updates is provided for by HD-TDMA through the periodic beacon-frame which is transmitted at the beginning of each allocated slot, by that slots owner. This functionality supports a distributed time-bounded fault tolerant membership service.

7.2.2 Probabilistic Admissions Control

This thesis presented a probabilistic admissions control approach which ensures that sufficient transmission resources are allocated to ensure the successful transmission of frames in line with both real-time reliability and real-time timeliness requirements.

Transmission resource requirements are based on the iterative application of the cumulative density function of the binomial distribution to determine the number of transmissions required based on channel conditions and real-time requirements.

In order to minimise the transmission resources required a clustering algorithm combines destinations with similar probabilities to exploit the characteristics of the binomial distribution to minimise the overall transmission resources required. Evaluations show that this approach results in a consistent reduction in the transmission resources required while satisfying real-time requirements.

7.2.3 Localised Scheduling to Overcome Transmission Failures

The localised scheduling approach exploits the understanding of the burst nature of errors in the wireless domain to optimise the transmission order to maximise the probability of transmission success. By deferring transmissions to destinations which experienced transmission failures the impact of burst error conditions are reduced while maximising throughput and minimising the impact on other transmissions.

This scheduling approach is key, as the localised scheduling process provides independence between transmission attempts thus allowing the application of the binomial distribution to determine the transmission resources required.

7.3 Review of Requirements

In section 2.5, ten requirements for a real-time wireless MAC protocol where outlined. This section reviews the satisfaction of those requirements by the combined HD-TDMA and probabilistic admissions control and retransmissions process.

A dynamic decentralised membership protocol is provided within HD-TDMA, which allows stations to join and leave a cell, incorporating failure detection and release of allocated slots of failed stations through consensus. Membership operations occur independent of real-time data.

An admission control process ensures that frames are only admitted when both their real-time reliability and timeliness requirements can be satisfied. All admitted frames will be transmitted within the specified deadline or within Δ HD-TDMA cycles whichever is soonest. A binomial clustering algorithm is employed to ensure sufficient transmission time is reserved to accommodate retransmissions in order to satisfy the reliability requirement. The binomial approach minimises the resources required while satisfying real-time reliability requirements which in turns leads to a greater throughput of data frames.

Flexibility is offered by supporting both acknowledged and unacknowledged data frames. To overcome the inflexibility of the the TDMA slot structure a polled transmission mode is offered. Transmissions within each HD-TDMA slot are by default prioritised by real-time priority, ensuring that under failure conditions the highest priority frames are sent in preference to lower priority frames.

Data throughput is maximised through the combination of several elements, firstly the contention free TDMA structure, secondly by the probabilistic admissions control and retransmission protocol which reduces the resources reserved for possible future retransmissions thus allowing a greater number of data transmissions.

Through a combination of MAC level flexibility provided by the ability to make several transmissions per HD-TDMA slot, the admissions control process which ensures that each frame allocated to a slot have a real-time deadline beyond the end of this slot time results in scheduling flexibility. Recognising the bursty nature of errors on the wireless medium it is possible to reschedule transmissions employing the flexibility within to maximise reliability and therefore throughput.

7.4 Future Work

Future work should investigate a number of issues, with the aim to further improve the performance of the HD-TDMA and the probabilistic admission control and retransmission protocols.

The presented clustered binomial approach provides a significant reduction in resource requirements compared to other considered approaches. However, it remains a sub-optimal approach. The Beta Binomial approach while promising, cannot be applied as the number of frames transmitted within each slot is small and does not provide a sufficiently large space to ensure the correct application of the Beta Binomial distribution. Further study may lead to a approach which generates an optimal result.

The admissions control system relies on an estimate of channel reliability, which in this thesis was based on a short term weighted history of channel reliability. Many current works focus on long term measurements of marginal statistics e.g. bit error rate, whereas the wireless medium, a bursty channel is constant and reliable with relatively little bad bursts [Wang et al., 2007]. In order to fully exploit the scheduling functionality within HD-TDMA and to maximise timely frame delivery, an reliability projection approach with memory is required, perhaps including elements of machine learning to adapt intelligently.

The queue assignment strategy employed by HD-TDMA selects the soonest queue which supports the real-time requirements of the submitted frame. This approach potentially, results in inefficient use of the transmission slot resources available to a station. In a lightly loaded system, it may occur that some slots are fully utilised but others poorly utilised. Under poor channel conditions there is no way to employ the free time in other available slots to ensure where transmissions cannot be successfully made in their assigned slots that they could, potentially be made subject to real-time deadlines in later slots which are not fully utilised. A more effective assignment algorithm should distribute frames across all available slots balancing not only the number of frames assigned to each slot but also the breakdown of priority of frames within each queue. In an ideal implementation the estimate of transmission success probability would be also a function of time to ensure the optimum time reservation for retransmissions in each slot.

Glossary

- **ATM (Asynchronous Transfer Mode)**: A packet oriented network protocol. A connection-oriented protocol which establishes a virtual circuit between receiver and destination before data is exchanged.

- **BTMA (Busy Tone Multiple Access)**: A channel access method using two channels, the control channel contains a busy tone when a station is transmitting on the data channel. A station wishing to transmit checks the control channel status before transmitting.

- **CAN (Controller-area network)**: A wired real-time communications protocol based on dominant and recessive bits, with priority encoded addresses providing deterministic and time bounded contention resolution.

- **CBR (Constant Bit Rate)**: Data traffic sent at a constant rate, e.g a video stream, telemetry.

- **CRC (Cyclic Redundancy Check)**: A polynomial function applied to a piece of data, which provides a signature of the data, such that if the data is damaged a very high confidence exists that the polynomial function will generate a different value.

- **CFP (Contention Free Period)**: A period within a communications cycle where transmissions are scheduled.

- **CP (Contention Period)**: A period within a communications cycle where random access is permitted, typically to allow stations to compete to gain access to the medium.

- **CTS (Clear to Send)**: A control packet used in many medium access control protocols, to indicate that the receiving station received a valid RTS packet. The CTS packet ensures all stations within range of the receiver are aware of an imminent data transmission and not to attempt to transmit until after that transmission has completed.

- **CSMA (Carrier Sense Multiple Access)**: a medium access control protocol which before each transmission attempt checks first to see if data is being transmitted by another station. If so, this station defers its transmission attempt to some time in the future and checks again. If the medium is determined to be idle the packet is sent.

- **CSMA/BA (Carrier Sense Multiple Access with Bitwise Arbitration)**: a medium access control protocol, relying on dominant and recessive bits, used by CAN Bus.

- **CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance)**: a wireless medium access control protocol where the transmitter attempts to avoid collisions through the use of both virtual carrier sensing (RTS-CTS) and physical sensing.

- **CSMA/CD (Carrier Sense Multiple Access with Collision Detection)**: a medium access control protocol. Employed by Ethernet.

- **GSM (Global System for Mobile communications)**: Originally from Groupe Spécial Mobile, a digital cellular based mobile phone network architecture, includes support for receiving and making voice calls when on a different network through roaming.

- **HD-TDMA-D (Hierarchal Distributed Time Division Multiple Access)**: A variant of the HD-TDMA protocol where transmissions to failed destinations are delayed in an attempt to overcome the bursty error phase.

- **HD-TDMA-I (Hierarchal Distributed Time Division Multiple Access)**: A variant of the HD-TDMA protocol where retransmissions to failed destinations are sent immediately.

- **MAC (Medium Access Control)**: The medium access control protocol implements two distinct functions, First that of arbitration of access to the communications medium and secondly that of transmission control.

- **MACA (Medium Access Collision Avoidance)**: a medium access control protocol, which exchanges RTS and CTS frames before data transmissions to overcome the hidden terminal problem.

- **RTS (Request to Send)**: A control packet used in many medium access control protocols, sent by the station which wishes to transmit to inform all stations in range not to transmit during this stations data transmission. Destination of the packet will reply with CTS.

- **SNR (Signal to Noise Ratio)**: The power ratio formed between the signal and background noise, units db.

- **SRMA (Split channel reservation multiple access)**: CSMA or Aloha medium access control principles using a separate control channel for signaling.

Bibliography

[Abramson, 1970] Abramson, N. (1970). "The Aloha System - Another Alternative for Computer Communications". In *Proceedings of 1970 Fall Joint Computer Conference*, volume 37, pages 281–285, Texas, USA.

[Agba et al., 2008] Agba, B. L.; Gagnon, F.; and Kouki, A. (2008). "Global Approach Of Channel Modeling In Mobile Ad Hoc Networks Including Second Order Statistics And System Performances Analysis". *Journal of Systemics, Cybernetics and Informatics*, 6(3), pp. 40–48.

[Aldridge and Ghanbari, 1995] Aldridge, R. and Ghanbari, M. (1995). "Bursty error model for digital transmission channels". *Electronics Letters*, 31(25), pp. 2144–2145.

[Ali-Rantala et al., 2003] Ali-Rantala, P.; Ukkonen, L.; Sydanheimo, L.; Keskilammi, M.; and Kivikoski, M. (2003). "Different kinds of walls and their effect on the attenuation of radiowaves indoors". In *Antennas and Propagation Society International Symposium*, volume 3, pages 1020–1023, Ohio, USA.

[Amouris, 2001] Amouris, K. (2001). "Space-time division multiple access (STDMA) and coordinated, power-aware MACA for mobile ad hoc networks". In *Global Telecommunications Conference (GLOBECOM)*, volume 5, pages 2890–2895, Texas, USA.

[Andersen et al., 1995] Andersen, J.; Rappaport, T.; and Yoshida, S. (1995). "Propagation measurements and models for wireless communications channels". *IEEE Communications Magazine*, 33(1), pp. 42–49.

[Balasubramanian et al., 2006] Balasubramanian, K.; Gathala, S. A.; Manimaran, G.; and Wang, Z. (2006). "A Novel Real-Time MAC Protocol Exploiting Spatial and Temporal Channel Diversity in Wireless Industrial Networks". In *Proceeding of the IEEE International Conference on High Performance Computing*, pages 534–546, Bangalore, India.

[Baldwin et al., 1999] Baldwin, R. O.; Nathaniel J. Davis, I.; and Midkiff, S. F. (1999). "A real-time medium access control protocol for ad hoc wireless local area networks". *SIGMOBILE Mobile Computing Communication Review*, 3(2), pp. 20–27.

[Basagni and Bruschi, 1999] Basagni, S. and Bruschi, D. (1999). "A logarithmic lower bound for time-spread multiple-access (TSMA) protocols". *Wireless Networks*, 6(2), pp. 161–163.

[Berber, 2003] Berber, M. S. (2003). "Accurate and time efficient estimation of the probability of error in bursty channels". *International Journal of Communication Systems*, 16(7), pp. 593–603.

[Bhagwat et al., 1997] Bhagwat, P.; Bhattacharya, P.; Krishna, A.; and Tripathi, S. K. (1997). "Using channel state dependent packet scheduling to improve TCP throughput over wireless LANs". *Wireless Networks*, 3(1), pp. 91–102.

[Bharghavan et al., 1994] Bharghavan, V.; Demers, A.; Shenker, S.; and Zhang, L. (1994). "MACAW: A media mac protocol for wireless LAN's". *SIGCOMM Computing Communication Review*, 24(4), pp. 212–225.

[Caccamo et al., 2002] Caccamo, M.; Zhang, L. Y.; Sha, L.; and Buttazzo, G. (2002). "An Implicit Prioritized Access Protocol for Wireless Sensor Networks". In *Proceedings of the 23rd IEEE Real-Time Systems Symposium (RTSS)*, pages 39–48, Texas, USA.

[Carley et al., 2003] Carley, T. W.; Ba, M. A.; Barua, R.; and Stewart, D. B. (2003). "Contention-Free Periodic Message Scheduler Medium Access Control in Wireless Sensor / Actuator Networks". In *Proceedings of the 24th IEEE International Real-Time Systems Symposium (RTSS)*, pages 298–307, Cancun, Mexico.

[Cheong, 2007] Cheong, Y. (2007). *Multivariate beta binomial distribution model as a web media exposure model*. PhD thesis, The University of Texas at Austin.

[Chlamtac et al., 1997] Chlamtac, I.; Farag, A.; and Zhang, H. (1997). "Time-spread multiple-access (TSMA) protocols for multihop mobile radio networks". *IEEE/ACM Transactions on Networking*, 5(6), pp. 804–812.

[Ci and Sharif, 2000] Ci, S. and Sharif, H. (2000). "Adaptive approaches to enhance throughput of IEEE 802.11 wireless LAN with bursty channel". In *25th Annual IEEE Conference on Local Computer Networks (LCN)*, pages 44–45, Florida, USA.

[Cristian, 1990] Cristian, F. (1990). "Synchronous atomic broadcast for redundant broadcast channels". *Real-Time Systems*, 2(3), pp. 195–212.

[Crow et al., 1997] Crow, B.; Widjaja, I.; Kim, L.; and Sakai, P. (1997). "IEEE 802.11 Wireless Local Area Networks". *IEEE Communications Magazine*, 35(9), pp. 116–126.

[Cunningham, 2003] Cunningham, R. (2003). *Time Bounded Media Access Control for Ad-Hoc Networks*. PhD thesis, University of Dublin, Trinity College.

[Cunningham and Cahill, 2002] Cunningham, R. and Cahill, V. (2002). "Time Bounded Medium Access Control for Ad Hoc Networks". In *Proceedings of the 2nd ACM International Workshop on Principles of Mobile Computing*, pages 1–8, Toulouse, France.

[Davis et al., 2007] Davis, R.; Burns, A.; Bril, R.; and Lukkien, J. (2007). "Controller Area Network (CAN) schedulability analysis: Refuted, revisited and revised". *Real-Time Systems*, 35(3), pp. 239–272.

[Decotignie, 2008] Decotignie, J.-D. (2008). “Real-Time Wireless Sensor Network: why do we need another view at them?”. In *Keynote, 7th International Workshop on Real-Time Networks (RTN)*, Prague, Czech Republic.

[Demarch and Becker, 2007] Demarch, D. D. and Becker, L. B. (2007). “An Integrated Scheduling and Retransmission Proposal for Firm Real-Time Traffic in IEEE 802.11e”. In *Proceedings of the 19th Euromirco Conference on Real-Time Systems (ECRTS)*, pages 146–158, Pisa, Italy.

[Doufexi et al., 2002] Doufexi, A.; Armour, S.; Butler, M.; Nix, A.; Bull, D.; McGeehan, J.; and Karlsson, P. (2002). “A comparison of the HIPERLAN/2 and IEEE 802.11a wireless LAN standards”. *IEEE Communications Magazine*, 40(5), pp. 172–180.

[Eckhardt and Steenkiste, 1996] Eckhardt, D. and Steenkiste, P. (1996). “Measurement and analysis of the error characteristics of an in-building wireless network”. In *Conference proceedings on Applications, technologies, architectures, and protocols for computer communications, (SIGCOMM)*, pages 243–254, California, USA.

[Eckhardt and Steenkiste, 1998] Eckhardt, D. and Steenkiste, P. (1998). “Improving wireless LAN performance via adaptive local error control”. In *Sixth International Conference on Network Protocols*, pages 327–338, Texas, USA.

[Egan, 2005] Egan, D. (2005). “The Emerence of Zigbee in Buildings Automation and Industrial Controls”. *IEE Computing and Control Engineering*, 16(2), pp. 14–19.

[Elliott, 1963] Elliott, E. (1963). “Estimates of Error Rates for Codes on Burst-Noise Channels”. *Bell Systems Technical Journal*, 42, pp. 1977–1997.

[Erceg et al., 1999] Erceg, V.; Greenstein, L.; Tjandra, S.; Parkoff, S.; Gupta, A.; Kulic, B.; Julius, A.; and Bianchi, R. (1999). “An empirically based path loss model for wireless channels in suburban environments”. *IEEE Journal on Selected Areas in Communications*, 17(7), pp. 1205–1211.

[Fritchman, 1967] Fritchman, B. (1967). “A binary channel characterization using partitioned Markov chains”. *IEEE Transactions on Information Theory*, 13(2), pp. 221–227.

[Fu et al., 2008] Fu, B.; Bernath, G.; Steichen, B.; and Weber, S. (2008). “Wireless Background Noise in the Wi-Fi Spectrum”. In *4th International Conference on Wireless Communications, Networking and Mobile Computing (WiCOM)*, pages 1–7, Dalian, China.

[Garcia Luna Aceves and Fullmer, 1999] Garcia Luna Aceves, J. J. and Fullmer, C. L. (1999). “Floor acquisition multiple access (FAMA) in single-channel wireless networks”. *Mobile Networks and Applications*, 4(3), pp. 157–174.

[Garg et al., 2003] Garg, P.; Doshi, R.; Greene, R.; Baker, M.; Malek, M.; and Cheng, X. (2003). “Using IEEE 802.11e MAC for QoS over wireless”. In *Performance, Computing, and Communications Conference*, pages 537–542, Arizona, USA.

[Garren, 2004] Garren, S. T. (2004). *Handbook of Beta Distribution and Its Applications*, chapter Goodness-of-Fit Testing of the Beta-Binomial Model, pages 237–253. Marcel Dekker.

[Gilbert, 1960] Gilbert, E. N. (1960). “Capacity of a burst-noise channel”. *Bell Systems Technical Journal*, 39, pp. 1253–1265.

[Gleeson, 2004] Gleeson, M. (2004). “A Real Time implementation of TBMAC using IEEE 802.11b”. Master’s thesis, University of Dublin, Trinity College.

[Gleeson et al., 2009] Gleeson, M.; Hughes, B.; Cunningham, R.; Weber, S.; and Cahill, V. (2009). “Implementation and Evaluation of Time-Bounded Medium Access Control for Wireless Networks”. Technical report, University Of Dublin, Trinity College.

[Gobriel et al., 2008] Gobriel, S.; Cleric, R.; and Mosse, D. (2008). “Adaptations of

TDMA Scheduling for Wireless Sensor Networks". In *Proceedings of 7th International Workshop on Real-Time Networks (RTN)*, pages 11–16, Prague, Czech Republic.

[Hamann et al., 2007] Hamann, C.; Roitzsch, M.; Reuther, L.; Wolter, J.; and Hartig, H. (2007). "Probabilistic Admission Control to Govern Real-Time Systems under Overload". In *Proceedings of the 19th Euromicro Conference on Real-Time Systems (ECRTS)*, pages 211–222, Dresden, Gemany.

[Hassanein et al., 2005] Hassanein, H.; You, T.; and Mouftah, H. T. (2005). "Infrastructure-based MAC in wireless mobile ad-hoc networks". *Ad Hoc Networks*, 3(6), pp. 717–743.

[Haykin, 2001] Haykin, S. (2001). *Communication Systems.* Wiley, 4th edition.

[Huber and Elmenreich, 2004] Huber, B. and Elmenreich, W. (2004). "Wireless Time-Triggered Real-Time Communication". In *Proceedings of the 2nd Workshop on Intelligent Solutions in Embedded Systems*, pages 169–182, Graz, Austria.

[Hughes, 2007] Hughes, B. (2007). *Hard Real-Time Communication for Mobile Ad Hoc Networks.* PhD thesis, University of Dublin, Trinity College.

[Hughes et al., 2006] Hughes, B.; Gleeson, M.; Karpinski, M.; Cunningham, R.; and Cahill, V. (2006). "Real-Time Communication in IEEE 802.11 Mobile Ad hoc Networks A Feasibility Study". Technical Report TCD-CS-2006-55, Dept of Computer Science, Trinity College, Dublin.

[IEEE, 1999] IEEE (1999). *IEEE std 802.11 Wireless LAN medium access control (MAC) and Physical Layer (PHY) Specification.* IEEE.

[IEEE, 2003] IEEE (2003). *Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wirless Personal Area Networks (LR-WPANs).* IEEE.

[IEEE, 2005a] IEEE (2005a). *IEEE std 802.11 Wireless LAN medium access control (MAC) and Physical Layer (PHY) Specification, Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements.* IEEE.

[IEEE, 2005b] IEEE (2005b). *Part 15.1: Wireless medium access control (MAC) and physical layer (PHY) specifications for wireless personal area networks (WPANs).* IEEE.

[ISO/IEC, 1994] ISO/IEC (1994). *ISO/IEC 7498-1 Information Technology - Open Systems Interconnection - Basic Reference Model: The Basic Model.* ISO, 2nd edition.

[Jiao et al., 2002] Jiao, C.; Schwiebert, L.; and Xu, B. (2002). “On modeling the packet error statistics in bursty channels”. In *27th Annual IEEE Conference on Local Computer Networks (LCN)*, pages 534–541, Florida, USA.

[João L Sobrinho and A. S. Krishnakumar, 1996] João L Sobrinho and A. S. Krishnakumar (1996). “Real-Time Traffic over the IEEE 802.11 Medium Access Control Layer”. *Bell Labs Technical Journal*, 1(2), pp. 172–187.

[Johnson, 1928] Johnson, J. B. (1928). “Thermal Agitation of Electricity in Conductors”. *Physics Review*, 32(1), pp. 97.

[Johnsson, 1999] Johnsson, M. (1999). *HiperLAN/2 - The Broadband Radio Transmission Technology Operating in the 5 GHz Frequency Band.* HiperLAN2 Global Forum, 1st edition.

[Kanal and Sastry, 1978] Kanal, L. and Sastry, A. (1978). “Models for channels with memory and their applications to error control”. *Proceedings of the IEEE*, 66(7), pp. 724–744.

[Karn, 1990] Karn, P. (1990). “MACA - A New Channel Access Method for Packet Radio”. In *Proceedings of the 9th ARRL Computer Networking Conference*, pages 134–140, Ontario, Canada.

[Katragadda et al., 2003] Katragadda, S.; Ganesh Murthy, C.; Ranga Rao, M.; Mohan Kumar, S.; and Sachin, R. (2003). “A decentralized location-based channel access protocol for inter-vehicle communication”. In *The 57th IEEE Semiannual Vehicular Technology Conference, VTC 2003-Spring*, volume 3, pages 1831–1835, Jeju, Korea.

[Kleinrock and Tobagi, 1975] Kleinrock, L. and Tobagi, F. (1975). “Packet Switching in Radio Channels: Part I–Carrier Sense Multiple-Access Modes and Their Throughput-Delay Characteristics”. *IEEE Transactions on Communications*, 23(12), pp. 1400–1416.

[Koopman and Chakravarty, 2004] Koopman, P. and Chakravarty, T. (2004). “Cyclic redundancy code (CRC) polynomial selection for embedded networks”. In *International Conference on Dependable Systems and Networks*, pages 145–154, Florence, Italy.

[Kopetz, 1997] Kopetz, H. (1997). *Real-Time Systems: Design Principles for Distributed Embedded Applications*. The Kluwer International Series in Engineering and Computer Science. Kluwer Academic Publishers.

[Kopetz and Grunsteidl, 1993] Kopetz, H. and Grunsteidl, G. (1993). “TTP - A time-triggered protocol for fault-tolerant real-time systems”. In *The Twenty-Third International Symposium on Fault-Tolerant Computing*, pages 524–533, Toulouse, France.

[Koubaa et al., 2006] Koubaa, A.; Alves, M.; and Tovar, E. (2006). “i-GAME: an implicit GTS allocation mechanism in IEEE 802.15.4 for time-sensitive wireless sensor networks”. In *18th Euromicro Conference on Real-Time Systems (ECRTS)*, pages 183–192, Dresden, Germany.

[Lin and Gerla, 1999] Lin, C. R. and Gerla, M. (1999). “Real-time support in multihop wireless networks”. *Wireless Networks*, 5(2), pp. 125–135.

[Malcolm and Zhao, 1995] Malcolm, N. and Zhao, W. (1995). "Hard real-time communication in multiple-access networks". *Real-Time Systems*, 8(1), pp. 35–77.

[Manoj and Siva Ram Murthy, 2002] Manoj, B. and Siva Ram Murthy, C. (2002). "Real-time traffic support for ad hoc wireless networks". In *Proceedings of the 10th IEEE International Conference on Networks (ICON)*, pages 335–340, Singapore.

[Metcalfe and Boggs, 1976] Metcalfe, R. M. and Boggs, D. R. (1976). "Ethernet: distributed packet switching for local computer networks". *Communications of the ACM*, 19(7), pp. 395–404.

[Minzer, 1989] Minzer, S. (1989). "Broadband ISDN and asynchronous transfer mode (ATM)". *IEEE Communications Magazine*, 27(9), pp. 17–24, 57.

[Murthy and Manoj, 2004] Murthy, C. and Manoj, B. (2004). *Ad Hoc Wireless Networks: Architectures and Protocols.* Prentice Hall Communications Engineering and Emerging Technologies Series. Prentice Hall PTR, 1st edition.

[Namislo, 1984] Namislo, C. (1984). "Analysis of Mobile Radio Slotted ALOHA Networks". *IEEE Journal on Selected Areas in Communications*, 2(4), pp. 583–588.

[Nelson and Kleinrock, 1985] Nelson, R. and Kleinrock, L. (1985). "Spatial TDMA: A Collision-Free Multihop Channel Access Protocol". *IEEE Transactions on Communications*, 33(9), pp. 934–944.

[Nguyen et al., 1996] Nguyen, G. T.; Katz, R. H.; Noble, B.; and Satyanarayanan, M. (1996). "A trace-based approach for modeling wireless channel behavior". In *Proceedings of the 28th conference on Winter simulation (WSC)*, pages 597–604, California, USA.

[Nicopolitidis et al., 2004] Nicopolitidis, P.; Papadimitriou, G. I.; and Pomportsis, A. S. (2004). "Distributed protocols for ad hoc wireless LANs: a learning-automata-based approach". *Ad Hoc Networks*, 2(4), pp. 419–431.

[Nyquist, 1928] Nyquist, H. (1928). "Thermal Agitation of Electric Charge in Conductors". *Physics Review*, 32(1), pp. 110–113.

[O'Connor, 2003] O'Connor, N. (2003). "Clock Synchronization for Multihop Wireless Networks". Master's thesis, University of Dublin, Trinity College.

[OPNET Technologies, 2008] OPNET Technologies (2008). *Modeling Concepts Reference Manual.* OPNET Technologies, Inc, 13th edition.

[Otani et al., 1981] Otani, K.; Daikoku, K.; and Omori, H. (1981). "Burst error performance encountered in digital land mobile radio channel". *IEEE Transactions on Vehicular Technology*, 30(4), pp. 156–160.

[Parsons, 2000] Parsons, J. D. (2000). *The Mobile Radio Propagation Channel.* Wiley, 2nd edition.

[Peterson and Brown, 1961] Peterson, W. and Brown, D. (1961). "Cyclic Codes for Error Detection". *Proceedings of the Institute of Radio Engineers*, 49(1), pp. 228–235.

[Prentice, 1986] Prentice, R. L. (1986). "Binary Regression Using an Extended Beta-Binomial Distribution, With Discussion of Correlation Induced by Covariate Measurement Errors". *Journal of the American Statistical Association*, 81(394), pp. 321–327.

[Proakis, 2001] Proakis, J. (2001). *Digital Communications.* Mc-GrawHill, 4th edition.

[Prodromides and Sanders, 1993] Prodromides, K. H. and Sanders, W. H. (1993). "Performability Evaluation of CSMA/CD and CSMA/DCR Protocols under Transient Fault Conditions". *IEEE Transactions on Reliability*, 42, pp. 116–127.

[Racciu and Mantegazza, 2006] Racciu, G. and Mantegazza, P. (2006). *RTAI 3.4 User Manual*, 0.3 edition.

[Rahnema, 1993] Rahnema, M. (1993). "Overview of the GSM system and protocol architecture". *IEEE Communications Magazine*, 31(4), pp. 92–100.

[Ramanathan, 1999] Ramanathan, S. (1999). “A unified framework and algorithm for channel assignment in wireless networks”. *Wireless Networks*, 5(2), pp. 81–94.

[Rappaport, 2001] Rappaport, T. (2001). *Wireless Communications: Principles and Practice*. Prentice Hall PTR Communication Engineering and Emerging Technologies Series. Prentice Hall PTR, 2nd edition.

[Rhee et al., 2005] Rhee, I.; Warrier, A.; Aia, M.; Min, J.; and Sichitiu, M. (2005). “Z-MAC: a Hybrid MAC for wireless sensor networks”. In *3rd ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 511–524, California, USA.

[Robert Bosch GmbH, 1991] Robert Bosch GmbH (1991). *CAN Specification version 2.0*.

[Roberts, 1975] Roberts, L. G. (1975). “ALOHA packet system with and without slots and capture”. *SIGCOMM Computing Communication. Review*, 5(2), pp. 28–42.

[Shannon, 1948] Shannon, C. E. (1948). “A Mathematical Theory of Communication”. *Bell Systems Technical Journal*, 27, pp. 379–423,623–656.

[Shepard, 1996] Shepard, T. J. (1996). “A channel access scheme for large dense packet radio networks”. *SIGCOMM Computing Communication Review*, 26(4), pp. 219–230.

[Sobral and Becker, 2008] Sobral, M. M. and Becker, L. B. (2008). “A wireless hybrid contention/TDMA-based MAC for real-time mobile application”. In *Proceedings of the 2008 ACM Symposium on Applied Computing (SAC)*, pages 284–288, Cear, Brazil.

[Stepanov and Rothermel, 2008] Stepanov, I. and Rothermel, K. (2008). “On the impact of a more realistic physical layer on MANET simulations results”. *Ad Hoc Networks*, 6(1), pp. 61–78.

[Sudhaakar et al., 2009] Sudhaakar, R.; Yoon, S.; Zhao, J.; and Qiao, C. (2009). “A novel QoS-aware MAC scheme using optimal retransmission for wireless networks”. *IEEE Transactions on Wireless Communications*, 8(5), pp. 2230–2235.

[Swarts and Ferreira, 1999] Swarts, J. and Ferreira, H. (1999). "On the evaluation and application of Markov channel models in wireless communications". In *IEEE Vehicular Technology Conference (VTC 1999-Fall)*, volume 1, pages 117–121, Amsterdam, the Netherlands.

[Tadokoro et al., 2008] Tadokoro, Y.; Ito, K.; Imai, J.; Suzuki, N.; and Itoh, N. (2008). "Advanced transmission cycle control scheme for autonomous decentralized TDMA protocol in safe driving support systems". In *IEEE Intelligent Vehicles Symposium*, pages 1062–1067, Eindhoven, The Netherlands.

[Takai et al., 2001] Takai, M.; Martin, J.; and Bagrodia, R. (2001). "Effects of wireless physical layer modeling in mobile ad hoc networks". In *Proceedings of the 2nd ACM international symposium on Mobile ad hoc networking & computing (MobiHoc)*, pages 87–94, California, USA.

[Tanenbaum, 2002] Tanenbaum, A. S. (2002). *Computer Networks.* Prentice Hall PTR, 4th edition.

[Tobagi and Kleinrock, 1975] Tobagi, F. and Kleinrock, L. (1975). "Packet Switching in Radio Channels: Part II–The Hidden Terminal Problem in Carrier Sense Multiple-Access and the Busy-Tone Solution". *IEEE Transactions on Communications*, 23(12), pp. 1417–1433.

[Tobagi and Kleinrock, 1976] Tobagi, F. and Kleinrock, L. (1976). "Packet Switching in Radio Channels: Part III–Polling and (Dynamic) Split-Channel Reservation Multiple Access". *IEEE Transactions on Communications*, 24(8), pp. 832–845.

[Wang et al., 2007] Wang, Y.; Yin, J.; Fu, W.; and Agrawal, D. (2007). "Performance Enhancement Scheme for Wireless LAN under Bursty Channel". In *Wireless Communications and Networking Conference (WCNC)*, pages 2155–2160, Kowloon, Hong Kong.

[Willig, 1997] Willig, A. (1997). "A MAC Protocol and a Scheduling Approach as Elements of a Lower Layers Architecture in Wireless Industrial LANs". In *Proceedings of 2nd IEEE International Workshop on Factory Communication Systems (WFCS)*, pages 139–148, Barcelona, Spain.

[Willig et al., 2002] Willig, A.; Kubisch, M.; Hoene, C.; and Wolisz, A. (2002). "Measurements of a wireless link in an industrial environment using an IEEE 802.11-compliant physical layer". *IEEE Transactions on Industrial Electronics*, 49(6), pp. 1265–1282.

[Ye et al., 2004] Ye, W.; Heidemann, J.; and Estrin, D. (2004). "Medium access control with coordinated adaptive sleeping for wireless sensor networks". *IEEE/ACM Transactions on Networking*, 12(3), pp. 493–506.

[Zhang, 2008] Zhang, S. (2008). "Real-Time Communication in Mobile Ad Hoc Networks". Master's thesis, University of Dublin, Trinity College.

[Zhou et al., 2005] Zhou, T.; Sharif, H.; Hempel, M.; Mahasukhon, P.; and Ci, S. (2005). "Performance of IEEE 802.11b in mobile railroad environments". In *IEEE 62nd Vehicular Technology Conference (VTC)*, volume 4, pages 2527–2531, Texas, USA.

[Zhu et al., 1991] Zhu, W.; Hellmich, T.; and Walke, B. (1991). "DCAP, a decentral channel access protocol: performance analysis". In *41st IEEE Vehicular Technology Conference - Gateway to the Future Technology in Motion*, pages 463–468, Missouri, USA.

[Zorzi, 1998] Zorzi, M. (1998). "Outage and error events in bursty channels". *IEEE Transactions on Communications*, 46(3), pp. 349–356.

[Zou et al., 2003] Zou, S.; Wu, H.; and Cheng, S. (2003). "A new mechanism of transmitting MPEG-4 video in IEEE 802.11 wireless LAN with DCF". In *International*

Conference on Communication Technology (ICCT), volume 2, pages 1226–1229, Beijing, China.