The ORL Radio ATM System, Architecture and Implementation

J. Porter* A. Hopper*† D. Gilmurray* O. Mason* J. Naylon*† A. Jones*

16 January 1996

Abstract

This paper presents a solution to the problem of connectivity of portables to an ATM wired network in the in-building environment. The approach to the support of ATM reduces the mobility load on the wired ATM network and is compatible with standard signalling protocols. The system is picocellular with a large number of base stations. The base stations are designed to be deployed in large numbers each covering a short range with partially overlapping coverages. This increases the aggregate throughput and reduces some of the problems specific to a radio physical layer. The MAC layer is optimised to provide efficient use of bandwidth and support guarantees for ATM traffic classes.

An experimental prototype system based on low-cost fixed ATM switches and software controlled base stations has been developed and is outlined.

1 Introduction

Asynchronous Transfer Mode (ATM) networking is advocated as a key technology for the provision of scalable, high performance networking for telecommunications, the local area, and even the home [1]. ATM, unlike competing technologies, provides true support for multiservice traffic through its ability to give Quality of Service (QoS) guarantees to individual connections. Standardisation in the wide area by the ITU-T with Broadband ISDN (B-ISDN) and in the local area by the ATM Forum has provided the catalyst to the introduction of interoperable low-cost ATM solutions in the marketplace.

Portable computers, including notepads, palmtops and PDAs, are becoming increasingly pervasive. As the functionality of these computers approaches those on the desktop, there is an ever greater requirement for wireless networks to support their connectivity. Multimedia communication, involving text, images, video and audio, which is now commonplace at the desktop, places very strict demands on the wired network. If mobile devices are to be seamlessly connected, then these demands must also be met by the wireless network. Many current and proposed wireless local area networks (WLANs) either do not support multiservice traffic or do so using schemes where traffic is statically divided into asynchronous and isochronous classes [2, 3]. Commercially available radio-based WLANs typically provide a low data-rate Ethernet-like service using either Direct-Sequence or Frequency-Hopping Spread Spectrum access.

By extending the use of ATM into the wireless network, we hope to provide the advantages of ATM to mobile devices along with integration with the local area network and beyond. There has been recent work on a number of different wireless ATM architectures [4, 5, 6, 7]. In this paper, we describe an wireless ATM networking system possessing the following features:

- Radio base stations connected to the wired ATM network provide the equivalent of a wireless *drop cable* from this network to the mobile.
- An ATM mesh of switches running standard ATM protocols is used for base station interconnection.
- ATM traffic characteristics, such as low latency, statistical multiplexing, and QoS reservations are preserved in the wireless network.
- Both fixed and mobile hosts use standard ATM adaptation and signalling protocols.
- Efficient use is made of radio bandwidth.
- Mobiles of varying complexity are supported.
- Active ATM virtual circuits (VCs) are preserved during handover with minimal loss of data.
- The small scale mobility of mobile devices is hidden from the wired ATM network.

The next section addresses the many issues that arise in the construction and deployment of a wireless radio LAN, and in the use of ATM as the networking technology for these LANs. Our proposed architecture for a radio ATM LAN is presented in Section 3. Section 4 describes in detail how handover of ATM connections is supported by the architecture. A Medium Access Control (MAC) protocol designed to efficiently share radio bandwidth between mobiles and support ATM traffic classes is outlined in Section 5. Section 6 describes a prototype implementation of the architecture built at Olivetti Research Ltd (ORL). Finally, Section 7 concludes with a summary of our architecture and implementation.

2 Issues for Wireless LANS

There are a number of issues concerning the design of a wireless network which do not arise in a wired system.

* Olivetti Research Ltd.
† University of Cambridge Computer Laboratory

These are primarily associated with the characteristics of the physical medium and the lack of precise control over the population using the medium.

2.1 Indoor Radio Environment

The indoor radio environment poses problems which do not occur in wired systems and are much less significant in larger scale wireless systems. These difficulties are mainly due to multi-path effects and high attenuation caused by the cluttered environment.

Multi-path effects are caused by the interference of multiple reflected rays from a source as they arrive at a receiver. Multi-path causes inter-symbol interference and fading. RMS delay spread [8] is a good measure of the multi-path dispersion and [9] shows that there is correspondence between the receive and transmit antenna separation, or pico-cell size, and the RMS delay spread.

There have been a large number of attenuation models studied for in-building radio propagation [10] and attenuation in most of these cases is considerably greater than line-of-sight. Walls provide the greatest attenuation in the radio system and in many cases the number of walls in a path is a reasonable metric for range.

Reducing the radio transmission cell size reduces the problems of multipath and high attenuation and thereby simplifies the design of the physical layer of the system.

2.2 Radio Transmission Range

As discussed in the previous section the physical layer problems in an indoor radio environment can be mitigated by a reduction in the radio transmission range. The choice of radio cell size has a significant impact on some of the general aspects of the system. A smaller cell size, improves frequency re-use and hence aggregate throughput, at the expense of the number of base stations and the frequency at which handover occurs [11].

Increasing the throughput using frequency reuse rather than increasing the bit-rate of each mobile reduces the cost because each physical interface can be much simpler.

The expense associated with a large number of base stations is constrained because each base station is a simple cell-relay and their network connectivity is provided by low cost connections to the existing ATM infrastructure.

Reduction of the cell size will increase the rate of handovers. The system, however, is only required to support mobility at human speeds which limits the rate of handover for a particular mobile to approximately one handover every ten seconds in a 10 metre pico-cellular system.

2.3 Access Point vs Peer-to-Peer Networking

There are two fundamentally different approaches to wireless networking: access point based, and peer-to-peer. Access point networking is where all communication for a set of mobiles in a transmission cell is routed through a central base station. Peer-to-peer networking or ad-hoc networking allow mobiles to communicate directly without a fixed infrastructure. The two approaches are two ends of a spectrum, but most wireless networks are optimised towards one of these alternatives.

In general it is probably not possible to make hard guarantees about reservations made in a mobile radio environment. The quality of a link to a mobile can not be guaranteed, as the mobile may move to another transmission cell with no available capacity. In an ad-hoc network the problem of making QoS guarantees is more difficult due to the uncontrolled interaction between sets of users.

In an ad-hoc system the number of devices which are available for communication is restricted to those with the same wireless networking interface and compatible protocol stack. In vertical markets this may well be successful where a large number of users are equipped with the same hardware and software. There is even less likelihood of compatibility when users buy their equipment independently. These problems are considerably reduced for an access-point system where a user needs only to be compatible with a single base station.

In an access point system the base station can be used to provide more sophisticated control over contention within the transmission cell.

One of the major motivations for a wireless ATM system is the ability to connect to the existing ATM infrastructure. Together with the performance advantages to be gained from an access point based system this has led us to conclude that peer-to-peer communication should not be a major design goal.

2.4 Colouring in a Wireless LAN

Colouring refers to the allocation of independent channels to adjacent pico-cells to prevent interference between them. Colouring can be implemented by a number of techniques, such as FDMA, CDMA or TDMA. Although these various forms of colouring are possible for an in building system with high data-rates, the design of the modem is typically limited by multi-path and sensitivity problems. A TDMA colouring approach requires that the modem operates at a multiple of the available throughput and requires synchronisation between adjacent base stations. A CDMA system requires an impractically large spreading bandwidth. FDMA therefore seems the most appropriate colouring technique for this type of system.

For full coverage in a cellular radio system there must be considerable overlap between adjacent transmission cells. In a fading environment the attenuation typically varies over a large range due to multi-path effects, even when movement is on the order of fractions of a wavelength. To give reasonable guarantees of coverage at the periphery of a cell there must be a significant fading margin in the power budget. Given this requirement for coverage at a cell's boundary, a mobile will be within range of multiple base stations when it is moving towards the limit of a transmission cell.

Some form of colouring is necessary in a practical system in order to successfully satisfy QoS guarantees. This is required so that an estimate of the utilisation of a pico-cell can be made in isolation from its neighbours. If there is no colouring then mobility of users within adjacent cells, or variations in signal fading between cells will effect the bandwidth currently allocated within a cell.

2.5 ATM for Wireless Networking

As mentioned in the introduction, there are many important advantages to be gained from using ATM as the networking technology for a wireless network. This section details these advantages and describes some of the associated problems.

2.5.1 Advantages

ATM performance: ATM is designed to provide integrated support for a wide variety of different types of traffic. If the ATM approach is extended to the wireless network, this integration can be available to mobile devices.

ATM applications: It is desirable that the same ATM applications can be run on both fixed and mobile hosts. This requires that the wireless network uses standard adaptation layers together with standard QoS reservation techniques and signalling protocols.

ATM internet: ATM networks will be providing both wide and local area connectivity. Wireless ATM can provide seamless interconnection for mobiles to this infrastructure.

Robustness to high BER: The radio channel typically has a high bit error rate with deep fades. A short PDU with acknowledgements makes it possible to maintain reasonable throughput with bit error rates as high as 10^{-3}.

2.5.2 Problems

Physical layer: There is inevitably some overhead associated with a radio system which is not experienced in a wired system. These are typically due to amplifier switching, frequency stability and synchronisation issues. These overheads present challenges for the efficient use of bandwidth with the transmission of small PDUs.

MAC layer: Unlike a conventional wired ATM system there are multiple end-points contending for the same physical layer. A more complex MAC layer is therefore required. Furthermore, the set of contending users changes as users move between pico-cells or as the radio propagation conditions change. Radio bandwidth is a precious resource and efficient bandwidth utilisation is important.

ATM layer: The shared channel entails the use of a meta-signalling protocol in order to multiplex signalling channels between the mobiles served by the same access-point. The VCI/VPI space is shared among a number of end-points within a pico-cell. Furthermore, the allocation must remain valid as mobiles move between pico-cells.

ATM is connection oriented and this requires state in the network which must be managed as a mobile migrates.

Adaptation layers: ATM adaptation layers are designed with the requirements that there will be low cell loss and no cell reordering. These are not easily achieved in a mobile environment.

2.6 Quality of Service Management

As defined by the ATM Forum there are four traffic classes CBR, VBR, ABR and UBR [12]. One of the main attractions of ATM is its ability to give QoS guarantees to these different traffic classes, based on traffic parameters such as cell loss ratio, cell delay and cell delay variation. These guarantees are achieved by preventing the establishment of new connections which exceed the capabilities of a route between two end-points using call admission control. Subsequently, sources are required to conform to a traffic contract.

In a mobile ATM network the assumptions used for making guarantees in the wired network are not necessarily valid. The bandwidth available to a set of mobiles served by a base station may vary due to changes in the link quality caused by channel fading. The set of mobiles served by a particular base station may vary due to the mobility of mobiles. The route of virtual circuits within the wired network will also change as a mobile moves between base stations.

There are a number of mechanisms by which QoS guarantees may be satisfied in a mobile network, including:

- Closing connections whose QoS cannot be maintained by a base station on handover.
- Prioritising the connections to a particular mobile so that, on handover, some connections will be preferentially retained.
- Prioritising the connections between mobiles so that handover can cause reclamation of bandwidth allocated to connections for other mobiles.
- Redistributing connections to balance the load between base stations.

A mobile network will certainly require a more sophisticated approach to QoS than that currently implemented in the ATM UNI standard. There will be a requirement for renegotiation of QoS during the lifetime of a connection as the service available from the network changes.

3 Radio ATM LAN Architecture

This section describes the architecture of the wireless ATM system implemented at ORL.

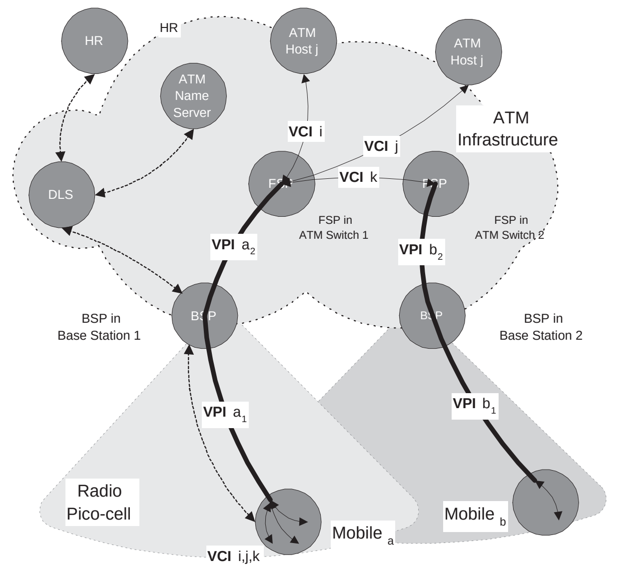

Figure 1: Radio ATM Architecture

The wireless network is divided into a large number of small pico-cells each served by a base station connected to the wired ATM network. Base stations are cell relays which translate the ATM cell headers from the radio ATM format to that used in standard ATM. Each base station supports a single radio frequency with adjacent base stations operating on different channels. The existing wired ATM network provides all connectivity between base stations.

The wireless network is organised into *domains*. A domain consists of a set of base stations which are under the management of a single entity. The domain corresponds to both a logical partition of the wired network and the physical location of the base stations.

3.1 Architectural Elements

This section outlines some of the management entities involved in the control of mobility.

HR (Home Register): This a software entity which is used to record the current proxy ATM address of a mobile. It also provides mobility services such as paging or notification.

DLS (Domain Location Server): The DLS registers the arrival of a mobile within a domain. It allocates the proxy ATM address and FSP for the mobile.

MSP (Mobile Switching Point): The MSP communicates with the current BSP and maintains tables of the link qualities and occupancy of neighbouring base stations.

BSP (Base Switching Point): The BSP handles the management of the virtual path element between the base station and the mobile.

FSP (Fixed Switching Point): The FSP is the point in the network through which all virtual circuits to the mobile are routed.

BM (Base Manager), MM (Mobile Manager): The BM and MM handle meta-signalling between the mobile and the base. They provides the parts of the interface to the radio ATM device which are different from the wired ATM device.

3.2 Scales of Mobility

The system is characterised according to the scale of mobility: Small-scale mobility is the mobility of devices between base stations within a management domain. Large-scale mobility is that between domains. The architecture is designed to isolate the small-scale mobility of mobiles within a very localised part of the wired ATM network and it is assumed that the dynamic nature of the large-scale mobility can be served by standard ATM protocols and services.

This division between large and small scales of mobility is achieved by splitting the route of a connection to a mobile into two sections. The first section of the connection is routed using standard ATM routing and addressing protocols. This part of the route terminates at a nominated point in the network, termed the FSP for the mobile. From the FSP to the appropriate base station, connections are routed using a *virtual path* (VP) mechanism. The base station relays the virtual paths for mobiles via a software entity called the BSP.

As the mobile moves from base station to base station within a domain, the part of its virtual circuits between the ATM end-points and the FSP remains fixed. The ATM routing protocols are not used to manage the dynamic section of the connections, as this is handled by the dedicated mobile software infrastructure.

3.2.1 Large Scale Mobility

In the current ATM standards, as defined by the ATM Forum [12], ATM addresses are organised hierarchically, where the hierarchical nature of the address is used for wide-area routing. Therefore, a single fixed address cannot be used to route connections to a mobile as it moves between domains. We avoid this problem by allocating a proxy ATM address to the mobile when it is registered within a domain. This ATM address is used to route the connection establishment message as far as the mobile's FSP. From there to the base station the routing is determined by the mobile network management.

3.2.2 Small Scale Mobility

In order to isolate the small scale mobility of the mobile as it moves between base stations a virtual path mechanism is used to route the virtual circuits from the FSP to the BSP. This has the advantage that the routing of connection requests to the appropriate base station is not dependent on the ATM address of the destination (whether the destination address indicates a mobile or fixed ATM host). We can therefore decouple the small scale routing between the FSP and the BSP from the PNNI [13, 14] protocols used in the private ATM network between switches.

When a mobile is registered, or on mobile handover, a virtual path is established between the FSP and the BSP. Connection requests destined for the mobile are routed via this virtual path to the BSP. The BSP then acts as a virtual path switch relaying the virtual circuits in the path to the mobile. Each virtual circuit has the same Virtual Circuit Identifier (VCI) at the FSP, BSP and mobile. This is so that, on handover of a mobile, the FSP can independently re-create the virtual circuits for the mobile.

There are likely to be a number of VCs active to a mobile end-point at any time. For example in a multi-media system, such as the ORL Medusa system [15], there are typically a number of control streams which do not necessarily have strict QoS requirements combined with a number of real-time or continuous media streams. Using the virtual path mechanism, virtual circuits can be controlled collectively with reference only to the VP. This greatly reduces the management overhead associated with the handover process.

During handover of a mobile the currently active VCs are copied via the new route in the wired network to the new base station. These VCs will not be made active until the handover process completes but the allocation of resources within the wireless network and within the new base station will be made at this time. This process of VC creation uses the standard signalling mechanisms where the routing is determined by the VP mechanism. Full support is provided for the QoS of connections, although during the handover there is excess reservation due to the duplication of the virtual circuits. The old and new virtual circuits to the mobile have QoS allocations at the same time. Compared to many handover architectures, this over-allocation of resources is well constrained, as it is limited to a factor of two for a single mobile during handover and there is no excess allocation up-stream of the FSP.

3.2.3 ATM Name Service

An ATM name service is required to decouple names and address within the ATM network. There are a number of proposals for a global ATM name service. The most likely candidate for acceptance as a standard is a system based on the Domain Name Service (DNS)[16, 17].

Unlike the IP concept of an address, an ATM address is associated more with the connection of an ATM end-point to the ATM network than with the end-point itself. The ATM name server must be more dynamic than DNS in order to support the re-allocation and re-binding of ATM addresses when ATM end-points are moved between switches.

Our architecture makes some extra demands of the dynamic nature of the name server. These demands are mainly due to the increased frequency of re-mapping requests, rather than a requirement for a reduction in the delay before updates become valid. In this architecture, the name-server is only updated when a mobile is allocated a new proxy ATM address, that is, when it is registered in a new domain. This new binding only needs to be available the next time an an ATM host attempts to make a connection to the mobile.

The mobile has a domain name which is used by ATM end-points to communicate with the mobile using the name server to determine the appropriate proxy ATM address. All mobiles are also assigned a globally unique fixed length ID, this is used to identify a mobile before it has established a signalling channel to a base station.

3.3 Wireless ATM Signalling Protocols

There are a number of well defined protocols standardised for signalling in ATM networking in a wired system. The ATM Forum have defined a set of protocols for user-network interfaces (UNI) and network-node interfaces (NNI). These protocols are not designed with mobility in mind: there is no direct support for mobility; connection setup is heavyweight due to the complexity of the signalling protocols and there is no support for a shared signalling channel.

These problems are solved by introducing a further set of protocols which run in parallel with the standard ones. This enables us to use completely standard protocols in fixed ATM end-points and run a superset of the protocols in mobiles. These protocols are meta-signalling, handover signalling and an RPC mechanism.

3.3.1 Meta-signalling

Meta-signalling refers to the protocols used to establish signalling channels in a situation where well-known permanent virtual circuits cannot be used directly. In the wireless network many mobiles may share the same broadcast channel to a base station. In order to use the standard signalling protocols, which have no support for multiplexing signalling requests over the same virtual circuit, a signalling channel must be established between the base station and each mobile. In this architecture this is achieved by allocating a virtual path to each mobile using the meta-signalling protocol. The signalling VCs then use well-known VCIs within this virtual path.

3.3.2 Handover Signalling for VC Switching

Handover signalling performs the in-band switching of active virtual paths. In order to reduce the problems associated with the use of UNI protocols, we adopt a two stage process for handover. The creation of new VCs operates on the same time-scale as signalling, but the switching of the active set of circuits must be significantly faster to minimise cell loss and delay.

The state machines which use the protocol do not assume reliability, so the protocol is simple. The on-the-wire encoding is also simplified, as we can choose an appropriate encoding depending on the end-points of the connection. Multiplexing and demultiplexing of handover signalling for different mobiles is handled efficiently by making use of separate VCIs.

3.3.3 RPC Between Management Entities

The RPC mechanism is based on a CORBA architecture [18] and is used to provide machine independence for the location of management entities. CORBA is a standard platform for the development of higher layer distributed systems implementations of which are available for workstations and embedded systems. We have a lightweight implementation for base stations and switches.

3.4 Wireless Virtual Circuits

3.4.1 Virtual Path Tunnelling

In many private ATM networks switched virtual paths are not available. Furthermore, the UNI ATM cell header only provides 8 bits for the VPI, which is a serious restriction on the number of mobiles which can be supported simultaneously within a domain. We circumvent these problems by using a *tunnelled* virtual path mechanism. At the end-points of the connection we use a programming interface which provides virtual path like semantics, but does not require the implementation of real virtual paths. This is achieved by *tunnelling* connection setup messages through the network.

Tunnelling is a term derived from mobile IP [19] in which a datagram can be encapsulated within a wrapper which has a different IP address. The datagram is routed through the network using this visible address. Outside the tunnel the datagram is stripped of the wrapper and is passed to higher layer protocols with the original addresses visible.

In the wireless ATM network an out-of-band connection mechanism is used to indicate the real end-points of a connection. The two actual end-points of a tunnelled VP share a signalling virtual circuit which is used for, among other things, the transmission of the tunnelling information. The tunnelling request includes the network addresses of the end-points of the connections and returns a private port allocated by the destination. When the connection request arrives from the source of the tunnelling request the port field is used to distinguish the original connection end-points.

3.4.2 Wireless VCI and VPI Allocation

Because the wireless network is a broadcast channel, a contention mechanism is required to share the bandwidth among the users of a single channel. The combined VCI and VPI space for all mobiles within range of a base station is therefore shared. It would be possible for a base station to allocate VCIs to mobiles from a flat VCI space but this has a number of disadvantages:

- The flat VCI space must be large enough to support the maximum number of VCIs in use by mobiles within range of a single base.
- Each mobile requires a lookup table in the interface to determine which VCIs it should receive.
- There are some limitations on the MAC protocol unless a short identifier is available which can identify a mobile within the pico-cell. In this architecture the VPI can be used for this functionality.
- Assuming that VCs are to be maintained on handover to a new base station the VCI allocation must be renegotiated. The new VCIs must then be rebound to the existing VCs in the wired network. This requires some means of identifying each circuit and re-binding the end-points both at the mobile and the point of connection to the existing VCs in the wired network. A handover signalling mechanism could certainly do this, but as will be discussed the delays inherent in this approach may be large.
- Signalling requests must be multiplexed for different end-points.

An alternative mechanism for dividing the VCI/VPI space is to allocate a VPI to each mobile within range of a base station. On handover the VPI is modified at the new base station. This provides a natural multiplexing of VCs to different mobiles and allows us to re-map VCs during handover by modifying only the VPI.

The same VCI is used to identify a virtual circuit at either end of the VP between the FSP and the mobile. The VCI is allocated at the base station and is then used in the wireless part of the link directly.

3.5 Multicast and Wireless ATM

Multicast is a desirable feature in a wired ATM network for reducing the load on the system and reducing the resources occupied by a point to multipoint connection. In a radio system it would seem that multicast comes for free, as the medium is inherently broadcast. However, the probability of broadcast being useful for multicast is proportional to the number of users sharing a pico-cell. This is likely to be small in a pico-cellular system.

Support of multicast adds considerable complexity in the wired network and this is likely to be even worse for a mobile system involving handover. As a mobile moves between pico-cells the multicast tree for a connection involving that mobile must be updated. The point at which the multicast is implemented must then be moved in and out of the wired and wireless part of the system.

An Automatic Repeat reQuest (ARQ) scheme is considered a necessity if reliability levels similar to those on the wired network are to be reached in the radio link. If there are potentially multiple recipients of a single transmission, as happens in a multicast, then a simple ARQ mechanism cannot be used. The solution is, either to increase the complexity of the MAC to allow multiple acknowledgements for the same transmission, or to implement reliability at a higher layer.

Our conclusion is, that radio broadcast is not useful for multicast in the wireless network. Broadcast is only used within a pico-cell for distribution of management information to the mobiles within that pico-cell. A mobile can of course take part in a multicast connection if the FSP is on a leaf of the connection.

3.6 Switching Point Implementation

As mentioned, the FSP is located in an ATM device through which connections to a mobile are routed, via its ATM proxy address. Connections are established as far as the FSP using standard signalling protocols. There are a number of choices for the location of the FSP: it can reside at an ATM host, a base station or an ATM switch.

The FSP can be located in an ATM end-point attached to a switch in the wired ATM network. In this case, the FSP terminates the UNI to its switch and the connections are forwarded using UNI along the dedicated VPI to the appropriate base station. In some cases VPs may not be available via the UNI, see Section 3.4.1. This approach has the important advantage that we can add mobility support to any ATM network without modification to the network itself. However, an ATM host is unlikely to be able switch cells efficiently although in most cases it will be capable of supporting the maximum data rate of a single base station. The route to the mobile, through the switch, is clearly non-optimal as connections traverse the same path twice.

If the FSP is located in a base station and the mobile is within the pico-cell for that base, then connections to the mobile are routed directly to the mobile. If the mobile has moved, however, then the connections are routed back to the current base station. No additional fixed infrastructure is required to support mobility as FSPs can run in the same hardware as base stations. However, when the mobile has moved between a number of base stations then FSP migration is required. (see Section 3.6.1).

The natural location for the FSP is in an ATM switch. If switched virtual paths are available, then the FSP has the same functionality as a VC/VP switch. A low level control interface to the switch is required, including access to the switch routing tables.

In our implementation, the FSP is located in a suitable ATM switch in the wired network. We have access to switches developed in-house over which we have control of the software. There is some scope for splitting the FSP management functionality from the switching functionality. In this case the switch has a simple control interface, which is used to manipulate the VCI tables. The higher level software can be run on a suitable workstation.

3.6.1 Switching Point Migration

When the DLS allocates an FSP to a mobile, its location within the wired network is chosen to optimise the use of resources. However, as the mobile moves between base stations the FSP may no longer be in the optimum location. This is usually not a serious problem due to the greater capacity of the wired network relative to the wireless network, but in some cases it will be desirable to migrate the FSP between physical ATM devices.

FSP migration is different from mobile movement as there is no physical layer handover requiring all routed virtual circuits to switch simultaneously. For simplicity, and to reduce the management cost, the migration process uses a two-stage mechanism similar to that used for connected handover (see Section 4.4). Individual virtual circuits are controlled collectively using messages which deal with virtual paths and the associated set of circuits.

3.7 Mobile Registration

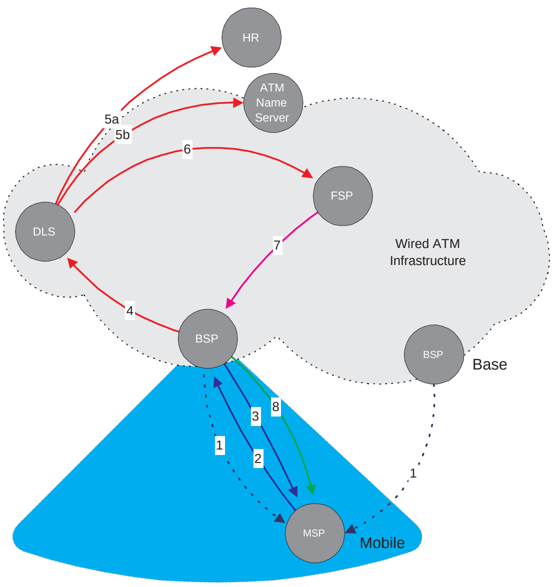

Figure 2: Mobile registration

Figure 2 illustrates the protocol used for registration of a mobile when it arrives in a domain. The mobile monitors the available radio channels listening for transmissions from base stations. It builds up a table of the link qualities to the detected base stations (1). Once the mobile has detected a suitable base station it transmits a *handover_request* meta-signalling message containing its unique identifier (2). When a base station receives such a message it allocates a VPI for the mobile and returns it in a meta-signalling *handover_request_ack*(3) message. The BSP can now use the handover signalling protocol to communicate with the mobile. The first message over the signalling channel from the BSP to the MSP is a query to determine the domain name of the mobile.

The base station registers the mobile with the DLS using RPC (4). The DLS then allocates a new proxy ATM address for the mobile and registers the new address with the HR and the ATM name service.$(5a, 5b)$. The DLS now assigns an FSP (6) for the mobile and informs it that it should create an initial virtual path to the appropriate base. The FSP creates the path and handover signalling channel to the BSP and transmits a *handover_complete* message to the BSP. The BSP forwards this message to the mobile, which completes the registration process. The mobile is now available for communication.

3.8 Pico-cell Topology

When a base station is initially turned on it is registered with the DLS which informs it of the available radio channels. The base station then monitors the adjacent base stations by cycling through the channels and measuring the signal strength of receptions. After monitoring for a period of time the list of neighbouring base stations is returned to the DLS. The DLS uses this information to maintain a connectivity map between the base stations under its control. This connectivity map is used by the DLS to implement frequency colouring with the allocation of a channel to the base station. If a large number of base stations are powered up simultaneously the DLS allocates channels in turn as the connectivity map is constructed.

In some cases re-allocation of a base station channel may be necessary. This re-allocation is similar to the handover of all mobiles which are active in the cell. This uses the connected handover model in which all mobiles are told to handover to the new channel at the current base station. This causes all the mobiles to perform a *handover_request* to the current base station. The base station then issues the *handover_ready* message after a timeout or when all mobiles have acknowledged. The mobiles then switch radio channels as in normal handover.

A BSP is informed of its neighbours by the DLS. The list of neighbouring cells is broadcast to the mobiles within the current pico-cell. This information includes the occupancy of the adjacent base stations so a mobile can make a reasonable decision about handover in situations where the link qualities are not decisive. Each BSP maintains its own current QoS allocation plus a cached version of its neighbours allocation. This information is periodically exchanged between neighbouring BSPs using RPC.

Each mobile monitors the quality of the links between itself and the available base stations. The link quality to the current base station is monitored by measuring the received strength of each incoming cell and the number of repetitions required for reliable transmission. This in-band link quality information is only available between the mobile and its current base station. The quality of the links to adjacent base stations is measured by periodically switching the mobile's physical layer to the channels for the neighbouring pico-cells and monitoring for receptions. This does not interfere with the data traffic because the channel is monitored at the end of the base station allocation frame (see Section 5).

4 Handover

We have already assumed that in a pico-cellular ATM environment handover of VCs between adjacent base stations is a requirement. We will now discuss this issue in more detail. The system supports real-time communications between mobile devices. There is a requirement that, as the mobile moves, continuous media connections, such as video or audio phone connections, maintain the appropriate QoS guarantees.

Handover may be necessary in a number of situations with different degrees of justification:

Load balancing between adjacent cells.: In this case the system is making a decision that a handover is necessary for reasons which do not necessarily benefit the mobile. It is clear that in this case the service to the mobile should be as continuous as possible.

Fading of a link to a stationary mobile.: The radio channel will suffer fades due to the motion of other objects, and again it seems reasonable that in this case the mobile should be given continuous service.

Movement of a mobile.: Here the mobile, or its attached user, is responsible for causing the handover and it is less obvious that the active virtual circuits need to be maintained. However, we consider that there will be many situations where this is desirable.

4.1 Rerouting of ATM Virtual Circuits

Some alternative mobile ATM schemes require the rerouting of active virtual circuits within the wired ATM network. We do not consider this to be a viable approach for a number of performance and pragmatic reasons:

- The PNNI protocols for the establishment and routing of ATM VCs between switches are already very complex. It seems unlikely that switch manufacturers will be willing to increase this complexity and impose further timing constraints on their switch software to support dynamic reconfiguration of VCs for mobile use.
- It is unlikely that any practical system would perform well enough to reroute active connections whilst providing adequate QoS guarantees. The time taken to re-route connections will be at least of the order of the current connection setup time, which is a large interruption in the context of an active VC.
- The management overhead of routing multiple circuits is considerable. Each VC is likely to require individual rerouting within the ATM mesh.

4.2 Soft Handover and Hard Handover

Hard handover, seamless handover and soft handover are terms used to describe different types of handover in cellular communication systems [20]. These terms have some relevance to a wireless ATM radio system.

In hard handover, the mobile device switches from one base station to another with active data being forwarded on only one path at a time.

In a seamless handover system, a path to a new base station is established, so that during handover two paths exist between the mobile and the base stations. However, data is only flowing on one path and the mobile switches paths based on the link quality.

In a soft handover two paths may be active simultaneously and the network makes a decision about which path should be used for extraction of the data. This scheme is used by the CDMA cellular system [21].

A design requirement of the wireless ATM network that it is an extension of the existing wired ATM network. The existing wired ATM infrastructure can then be used to support base stations. Furthermore, the wireless ATM network should support the same adaptation layers as the wired ATM network. There is therefore no way of including sequencing information in each ATM cell and no way of filtering out duplicate cells arriving via different paths within the wired network. [1]

The wireless ATM network uses a hard handover mechanism because of the complexity of a dual channel receiver. The system behaves in some ways as if it was using seamless handover as the new virtual path is established before the physical handover is initiated.

4.3 Handover Decision

The MSP receives information from the Mobile Manager about the quality of the communication links to adjacent base stations. Other information is combined with link quality measurements in order to make a decision about handover. This information concerns: resource utilisation, topological physical and logical information, historical communication and location information, and management information such as billing, maintenance and reliability. This information is generally of use to all mobiles within a pico-cell and radio broadcast is used for its distribution.

4.4 Connected Handover Protocol

While the mobile can still communicate with the current base station the handover protocol shown in Figure 3 is used. This is a form of backward handover in which the old base station is used for control until the new path is established.

[1]The GFC field in the UNI ATM cell header is not guaranteed to be carried between switches and thus cannot be used in a general system

The mobile initiates handover by sending a request via the current signalling channel to BSP_1 indicating that it wishes to handover to BSP_2. This message is then forwarded to the FSP. The $handover\_request$ message is acknowledged immediately as the path creation phase is relatively slow.

On receipt of a $handover\_request$ the FSP creates a virtual path to BSP_2 and copies all the virtual circuits. Once this process has completed the FSP sends a $handover\_ready$ message to the mobile via BSP_1. The mobile acknowledges this message and switches to the appropriate radio channel for the new base station. Receipt of the acknowledgement causes the FSP to switch completely to the new virtual path and free the old path.

The final stage of handover is similar to mobile registration, except the path to the FSP is already established. The mobile transmits a meta-signalling message to BSP_2 and is allocated a new VPI which BSP_2 then associates with the previously created path to the FSP.

4.5 Disconnected Handover Protocol

If the mobile has already lost contact with the current base station, it uses the protocol illustrated in Figure 4. This is broadly similar to the connection establishment which occurs when a mobile is first registered. Having decided which base station to choose for handover, the mobile switches to the appropriate radio channel and transmits a meta-signalling *handover_request* message. Since in this case the BSP does not have a signalling channel to the FSP it uses RPC to communicate with the DLS. The DLS then instructs the FSP to handover the mobile to the new base station, which creates the new path to BSP_2 and creates the virtual circuits. The completion of the handover is indicated by the *handover_complete* message which is forwarded to the MSP.

4.6 Cell Re-ordering and Cell Loss

In normal operation, ATM guarantees that within a single virtual circuit cell sequence integrity will be preserved. Under normal operation of the radio network, where circuits are subject to handover, this requirement is a challenge. There is a tradeoff between cell loss and cell delay; a scheme which guarantees no cell loss will introduce large delay. We have considered a number of potential schemes for controlling the re-ordering of cells during handover:

Relaying Cells.: In this approach cells are copied between the relevant base stations using a separate virtual circuit. The virtual circuit used for relaying cannot have any specific QoS allocation as it is used for the set of circuits at the mobile.

Relaying Virtual Circuits.: Here the virtual circuits are routed via the original base station to the new base station. Cells queued at the old base station can then

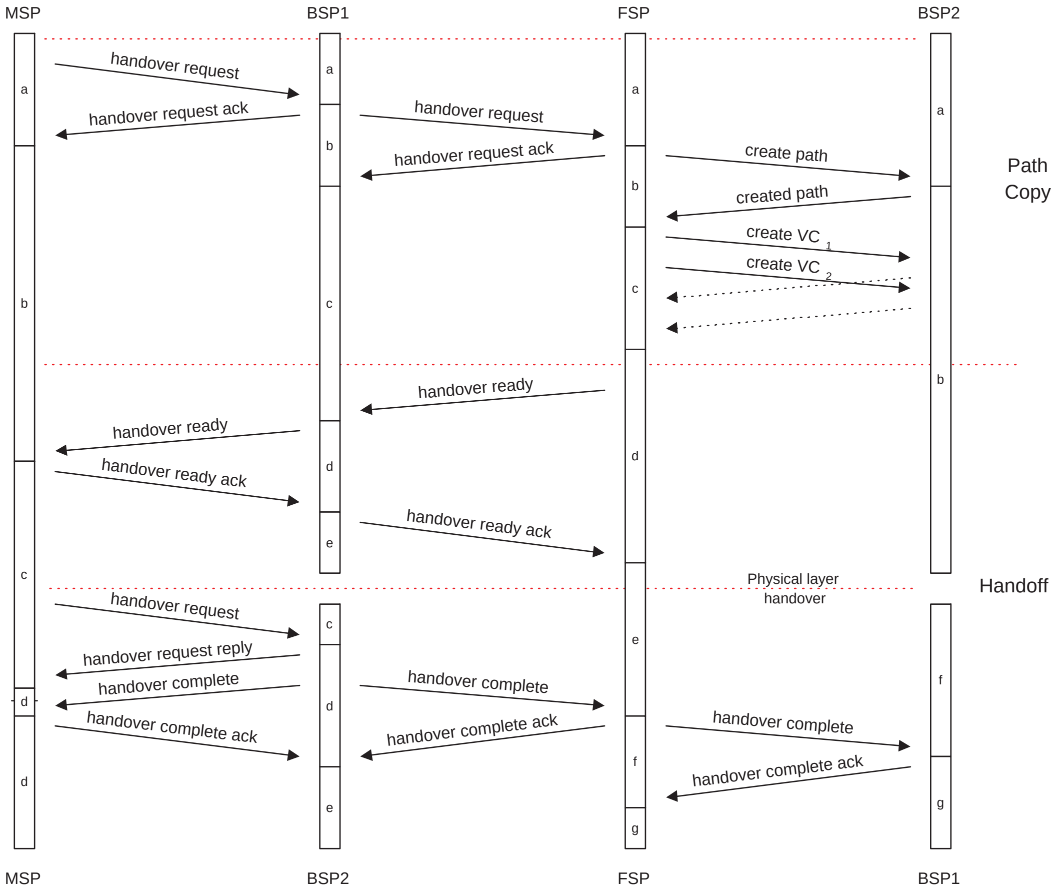

Figure 3: Connected Handover

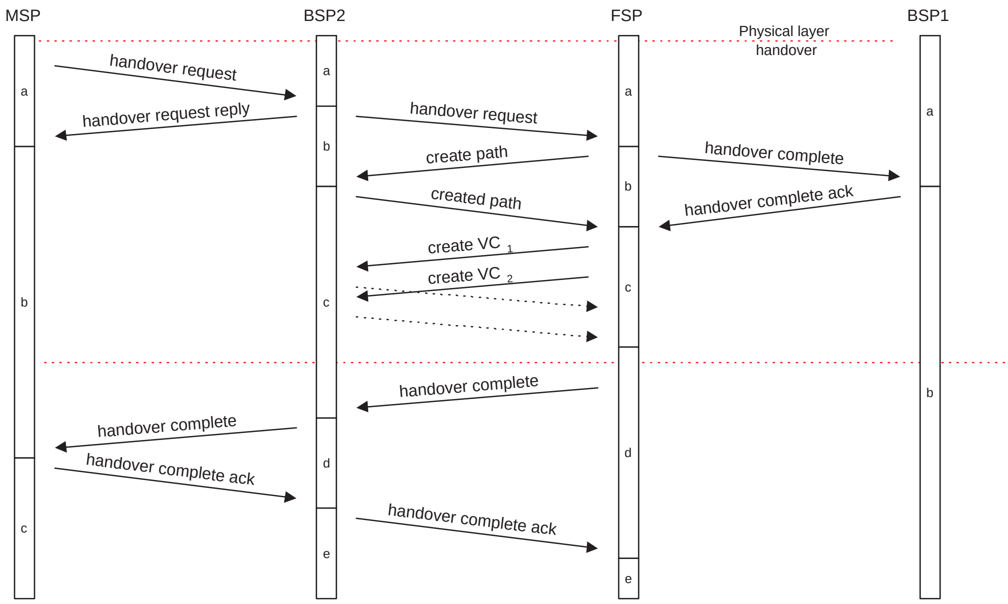

Figure 4: Disconnected Handover

be forwarded along the new path with no loss or re-ordering. Unless there is some reclamation of resources or re-routing of active VCs within the ATM network this may result in virtual circuits being routed many times through the same base station. Solving this problem by re-routing the virtual circuits requires further mechanisms to prevent cell loss and re-ordering.

Flushing Virtual Circuits.: Circuits are flushed using an in-band message which is transmitted on each virtual circuit. In this scheme the cell stream is monitored for a flagged cell which indicates when handover is occurring. This is probably the ideal approach, but the management overhead is high and a field is required in the ATM header to identify a flush message.

Flushing Virtual Paths.: This is applicable to a scheme in which all the virtual circuits are routed via a single virtual path. The set of circuits is flushed collectively based on the virtual path through which it is routed. A statistical approach is used to reduce loss of data based on the expected delays within the system.
It is assumed that in the system the bottleneck is almost always the interface to the radio network. Most of the buffering of cells transmitted over the radio will therefore be at either the mobile or the base. The flushing of virtual paths only needs to flush these buffers to result in adequate performance.

In the current prototype radio ATM system we chose to use the virtual path flushing mechanism although in the future we will consider flushing by virtual circuit.

5 Medium Access

In a wired ATM network the MAC protocol can be relatively simple as there is normally a point-point link between the entities sharing the physical layer. This is not the case in a radio system where the channel is shared by a set of users. Furthermore, the channel is unreliable and bandwidth is constrained.

The MAC protocol is optimised to support a small number of mobiles whose primary traffic loads are ABR and VBR. Support for CBR is provided, but it is assumed that the CBR traffic load will be a smaller proportion of the overall data-rate.

5.1 Variable Frame Reservation Protocol

The MAC protocol is a frame based reservation protocol similar to a number of previously described protocols [22, 23, 24, 25]. The protocol is designed to support the three main ATM traffic classes (ABR, VBR and CBR, see Section 5.7) and has a number of novel features to achieve this. The allocation of bandwidth to CBR traffic uses a fixed length reservation frame independent of the MAC frame which gives periodic slot reservations to mobiles. All other transmissions are made with explicit reservation. The number of cells transmitted in a MAC frame is variable, with the length of the MAC frame determined by the number of outstanding reservation requests and the recent contention history.

The allocation of bandwidth to up-stream and down-stream traffic in the MAC protocol is completely flexible due to the variable frame structure. This is important in a mobile ATM system where it is very difficult to predict what the actual division of traffic will be. Some applications, such as multi-media phone systems, have symmetrical traffic whereas many applications will involve higher down-stream traffic loads.

The MAC frame (Figure 5) consists of a down-stream part and an up-stream part. The frame fields are interpreted as follows:

Preambles (P): for each element which can be received independently,

Acknowledgements (A1,A2,A3): for the previous up-stream cells and reservation requests.

Frame Descriptor Field:

- **D:** number of down-stream cells
- **R1,R2:** VPIs for up-stream reservation indications
- **C:** number of contention slots

Down-stream MAC PDUs: The PDU containing ATM cell and MAC protocol wrapper (see Figure 6).

Acknowledgements (A3,A4): acknowledgements for the down-stream cells in the same MAC frame.

Up-stream MAC PDU: The PDU in the up-stream direction (see Figure 6) contains the sequence number, reservation request, ATM cell including condensed header and the CRC error check. The reservation request indicates the size of the reservation as a number of cells and the traffic type (ABR or VBR).

The frame contains at least one contention interval, which is used in a slotted ALOHA mode to make a reservation in the subsequent MAC frame. Reservation requests can also be piggybacked onto a cell transmission.

5.2 Physical Layer Error Control

The ATM adaptation layer protocols are designed with specific assumptions about cell loss rate. The corruption or loss of a cell is assumed to be rare. In the radio environment the raw radio link has a much higher error rate than that required for ATM. Some effort must therefore be made to detect and correct data errors. FEC (Forward Error Correction) is a common technique especially for unidirectional channels. For bidirectional data networks with bounded PDU and acknowledgement times ARQ is often used.

In our MAC each cell is covered by a 16-bit CRC which is used for error detection. The CRC covers both the header and payloads of the ATM cell and the MAC header. Cells

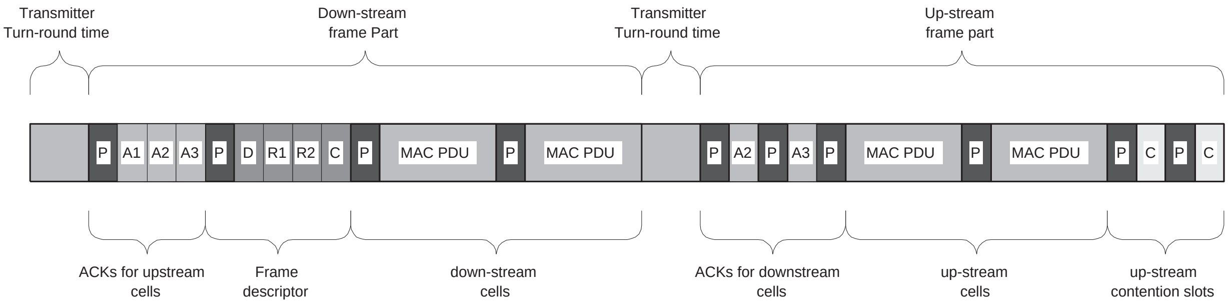

Figure 5: MAC Frame

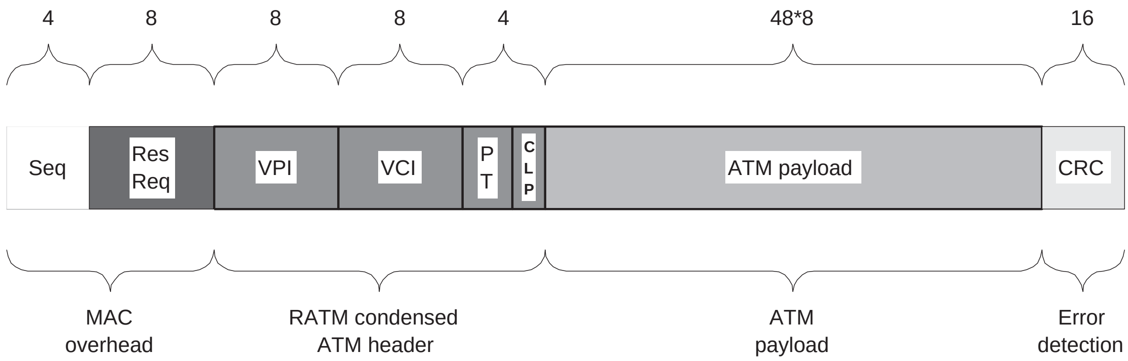

Figure 6: MAC PDU

which are received correctly are acknowledged in the subsequent acknowledgement interval. The correspondence between the acknowledgement and the cell transmitted is determined by including the VPI of the cell in the acknowledgement. The interface repeats an unacknowledged cell a number of times before it is discarded. The number of repeats is dependent on the traffic class for the cell. If a cell is not correctly received, the reservation is carried over to the next MAC frame.

5.3 MAC Frame Size

The number of up-stream and down-stream cells in a MAC frame is limited to ten cell times in order to place an upper bound on cell delay variation for CBR streams.

The overhead imposed by the physical constraints of the radio is reduced by transmitting cells in bursts. However, each cell transmission is independent, so that each cell in a burst can be received if other cells are lost. For that reason each cell has its own preamble and CRC . This helps retain the advantages of transmitting short PDUs in a high BER channel.

Between adjacent cells in a burst from the same transmitter, there is no overhead due to amplifier settling, phase-lock loop disturbance and clock synchronisation. The inter-cell preamble in a burst can therefore be reduced compared to the preamble used between cells from different sources.

5.4 MAC Interface Queueing

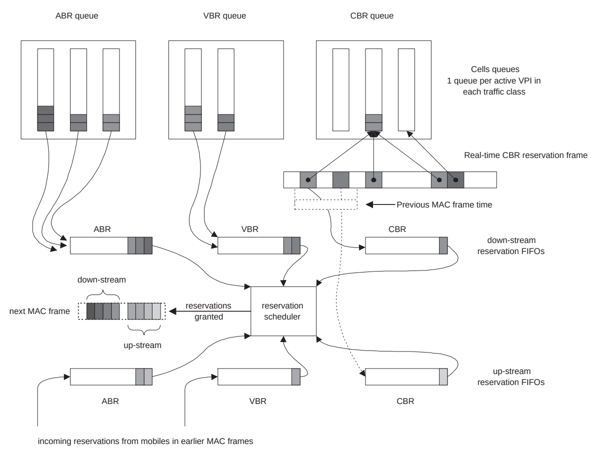

Figure 7: MAC Implementation

The MAC interface maintains queues in SRAM comprising linked lists of cells to be transmitted and received. The base maintains three sets of transmission queues for the three supported traffic classes. Each traffic class queue has a circular linked list of cell queues per VPI. The reservation scheduler generates the structure for the MAC frame by combining cells from the base station queues with previous requests from mobiles for reservations.

Each cell contains a sequence number which is used to resequence cells which can be re-ordered due to the lossy channel. The sequence numbers are incremented independently for each VPI.

5.5 ATM Traffic Policing Call and Admission Control

Traffic policing is the enforcement of a traffic contract which prevents ATM end-points from using more than the agreed allocation. In many ATM systems this has not yet been implemented. Traffic policing requires that the particular traffic distribution can be measured on an individual virtual circuit.

Call admission control is the decision made at connection setup time to either grant or refuse a QoS contract based on the available resources. The decision is made based on the resources in each switching element in turn, although routing takes into account the distribution of resources in other parts of the network [13].

Although these mechanisms concern the traffic patterns flowing on an ATM link, each is concerned with processes which happen on very different time scales from the medium access control. In the MAC architecture there is interaction between the call admission control and the allocation of bandwidth to CBR. For other traffic classes the relationship is statistical. There is no direct allocation or reservation of resources within the MAC protocol corresponding to ABR or VBR circuits.

5.6 ATM Traffic Shaping

In wired ATM, ABR traffic is rate paced so that the peak throughput of a VC is controlled. This is achieved by spacing out the transmission of individual ATM cells and this function is normally performed by the ATM interface. The host writes complete blocks into the ATM interface which are then transmitted at the currently allowed peak rate.

The reservation MAC protocol performs optimally when mobile hosts do not rate pace ATM cells within blocks. Traffic shaping, if required, is performed in the base station as the cells are relayed onto the wired ATM network. There is some tradeoff here between the efficiency of the bandwidth utilisation and the traffic shaping required by ATM. However, the wireless network will still provide interleaving at the cell level if multiple blocks are available for transmission.

5.7 Support for ATM Traffic Classes

The current set of QoS classes defined by the ATM Forum are designed to support the anticipated mix of traffic in ATM networks. The traffic classes are Available Bit Rate (ABR), Constant Bit Rate (CBR), Variable Bit Rate (VBR).

5.7.1 ABR Support

ABR is used for bursty LAN type traffic where there is no fixed cell rate. It is designed for best effort service which makes use of bandwidth not assigned to the other traffic classes.

5.7.2 VBR Support

VBR service class is used for connections such as variable bit rate video which have specifications of mean rate, peak rate and burst size. There is some debate about how VBR should be supported in an ATM system. Our view is that it is treated as ABR, with high priority, as far as transmission is concerned and a statistical approach is used for call admission control. Policing for VBR traffic operates over a much longer time-scale than the MAC and should be implemented in the end-points. In our system, the call admission control uses an effective bandwidth calculation for VBR allocation within the pico-cell [26]. This is a statistical measure which is used to determine whether the requirements for a particular call can be met with a predictable probability of cell loss.

5.7.3 CBR Support

The CBR service class is used for traffic such as voice calls which have a constant bit rate requirement with typically a low latency requirement. A fixed length allocation frame is used to provide CBR reservations. The fixed allocation frame is clocked in real time, with each entry in the frame providing a guaranteed bandwidth for mobile originated communication.

The frame repeats every 5 ms giving one cell transmission (assuming 40 bytes of user data per ATM cell) every 5ms. This gives a resolution of 64Kbit/s for each unit of allocation. These reservations are allocated during call setup for a connection. If the requested CBR bandwidth is available then the VPI for the mobile is stored in the CBR reservation frame.

A real-time counter is used to index the reservation frame. The time of the beginning of the previous frame and the current time is used to determine the reservations for the current frame. This means that it is possible to determine the CBR reservations required in the current frame at the start of the calculation of the MAC frame.

CBR reservations are reallocated on handover. When the existing virtual circuits are copied to the new base station, the MAC re-creates the reservations if possible.

5.8 MAC Protocol Performance

5.8.1 Throughput Performance

As the required bandwidth of a single source approaches the available bandwidth, multiple cells will be queued in the interface. This will cause the hardware to continuously maintain outstanding reservation requests. Thus, the MAC protocol behaves like slotted ALOHA for low load and approaches TDMA as the load is increased.

Variable frame length means that contenders are removed from contention and given a reservation slot as soon as possible.

5.8.2 Bandwidth Utilisation

Turnaround time between transmission and reception is one of the limiting performance factors of a radio system. The turnaround time is dependent on the design of the physical layer but is to some extent independent of the bit-rate. For higher bit-rate channels and channels with a large latency it is clearly important to minimise the number of turnaround times per cell transmission. The MAC protocol minimises this overhead by allowing multiple sequential cell transmissions between acknowledgements.

Many slot reservation schemes have a static allocation of bandwidth to different sources or traffic classes. This is inefficient for an ATM network in which statistical multiplexing is used to support bursty ABR or VBR traffic. In a fixed frame system static allocation is often made to both.

In our MAC the worst case bandwidth efficiency is experienced when a single source is generating cells at a rate marginally too low to maintain continuous reservations. In this case a separate reservation request is required in the contention slot for every transmission. This is an unlikely situation and since the source is not demanding excess bandwidth this not a significant problem.

5.8.3 Contention Performance

The contention slot consists of fields indicating the VPI and traffic class for which a reservation is required and a checksum. A transmission of a reservation request is made in a contention slot with probability of $1/c$, where c is the number of contention slots. When a collision is detected by the base station the number of contention intervals is increased and this has the effect of stabilising the ALOHA access [27]. When idle contention intervals are detected the number of contention slots is reduced until there is one left.

In practice it is unlikely that there will be much contention. Much of the published data on slotted ALOHA assumes a large population and does not include the probability of capture. In this system there is a small population which only uses ALOHA access to acquire a reservation for the first cell of a burst. As soon as a successful reservation is made the source is removed from contention.

5.9 MAC Simulated and Measured Performance

For the final draft of the paper simulated results will be available for the MAC performance.

The final version of the MAC is currently being implemented and it is hoped that performance figures will also be available before final submission.

6 Experimental Wireless ATM Network

6.1 ATM Framework

An experimental radio system has been built with the aim of investigating the performance of the wireless ATM system. The prototype is based around a range of low-end ATM switches which use the Advanced RISC Machines ARM processor as a central switching unit. This switching is either directly through shared memory or by using broadcast backplane fabrics. Where a backplane is used, the ARM CPU sets up DMA transfers between input and output ports on the switch. There is a control interface to the switch which can be used to communicate to another system which evaluates VCI routes and performs signalling and other control and management functions. If performance permits these two functions can be combined on a single CPU.

This approach makes possible a flexible strategy for handling control and data in the switch. At some cost in performance it is possible to derail the ATM cell and execute priority or other QoS mechanisms in software. The range of switches includes 4x4, 8x8 and 16x16 designs and typical throughput performance for the backplane is around 4M cells per second. At this speed the performance of the ATM switches is sufficient to support the traffic requirements in the local area.

6.2 ATM Direct Peripherals

The basic hardware design of the switch is based on a building block, the ATMos card[2], which can also be used for attaching peripherals directly to the ATM network. The design consists of line modules on the network side typically operating at 100Mbps and based on TAXIchip[3] technology. On the data side the hardware presents a general purpose data interface to which peripherals (or switch fabrics) can be attached. Thus the radio module becomes a direct peripheral which in this role performs the functions of a bridge between the fixed and wireless networks.

A software kernel, also called ATMos, runs on this standard platform and is designed to facilitate the writing of control programs for small to medium size applications. The kernel is a lightweight real time operating system for the support of ATM networked applications. It consists of a boot ROM, applications libraries and initial core modules. There is a UNIX[4] -based development and debugging envi-

[2] ATMos is a trademark of Advanced Telecommunications Modules Ltd

[3] TAXIchip is a trademark of Advanced Micro Devices

[4] UNIX is a registered trademark of UNIX System Laboratories

ronment as well as interfaces to POSIX threads and general purpose distributed platforms based on the CORBA specification. A typical embedded ATMos application will consist of the scheduler, a number of device handlers, as well as the application related processes. The base station for the support of radio ATM consumes around 25 percent of CPU servicing the ATM network.

6.3 Radio Interface

The system is designed to minimise the complexity of the RF interface while providing as much flexibility as possible for experimentation with prototype access protocols.

The system operates in the 2.4–2.5GHz ISM band with a bandwidth of 10MHz per channel. Transmissions conform to the direct-sequence spread spectrum rules for power density within the European ETSI regulations [28]. Transmission power is less than 10dBm and this gives sufficient power for pico-cells of a radius of around 10 metres. The modulation scheme is QPSK with a bit-rate of 10Mbit/s; QPSK is used to reduce the symbol rate and allow us to keep within a 10MHz frequency bandwidth. Figure 8 outlines the radio front end architecture.

One of the major design goals was to minimise the transmission overhead per cell. This requires that the radio turnaround time, between transmission and reception, must be minimised. Currently this is of the order of 40 bit-times or 4 micro-seconds including preamble.

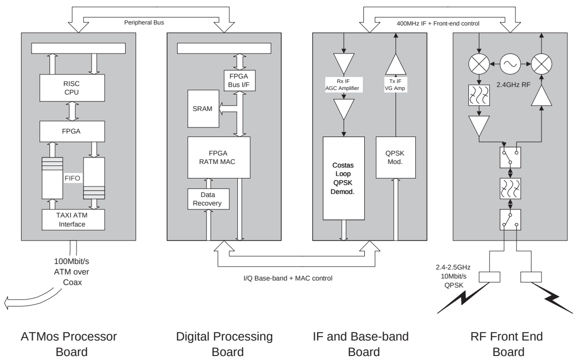

Figure 8: Hardware

6.3.1 RF Implementation

The receive LNA amplifier is a GaAs FET and a silicon MIMIC. An active down-converter is used to convert to an Intermediate Frequency (IF) of 400MHz with a variable frequency 2.0–2.1 GHz first local oscillator. The 400MHz IF has an AGC with a fast attack and slow decay. The fast attack gives a quick response to the beginning of the cell and when the end of the cell is detected by the MAC layer, the gain of the amplifier is forced low.

QPSK phase recovery uses an analogue Costas Loop. When there is no detected signal the QPSK demodulator is driven from a fixed frequency phase-locked 400MHz oscillator. When a received signal is detected the phase error signal from the Costas Loop is used to lock the oscillator. The recovery oscillator, therefore, only has to track the total frequency error between transmitter and receiver and is already centred at close to the correct frequency.

The front-end antenna switch selects a particular antenna before transmission of each cell. The choice of antenna is determined by the success rate of transmissions.

A variable gain amplifier stage is used at the IF to provide isolation when receiving. The output power is varied to optimise the pico-cell coverage.

The transmit side uses a separate fixed frequency 400MHz oscillator so that the switch between reception and transmission can be as fast as possible.

The total switching time between reception and transmission is of the order of 2 micro-seconds. Given that the total cell transmission time is around 50 micro-seconds, this is considered to be a reasonable overhead.

6.3.2 Host Interface

Figure 8 shows the prototype host interface to the radio network interfaced to the 32-bit ARM host.

The MAC protocol is implemented in a re-programmable Xilinx[5] gate array allowing experimentation with the details of the MAC protocol. A digital phase-locked loop is used to recover the incoming symbol phase and this is also used to generate the transmission clock.

6.4 Coverage Trials

In order to make measurements ahead of the full hardware design of the mobile end-points, radio base stations are used to implement mobiles. A base station provides an ATM interface for the attachment of peripherals which can emulate mobile devices. In this configuration, we have conducted coverage trials around our building, and tested our support for multimedia streams.

A preliminary set of trials was conducted to investigate the propagation characteristics of the prototype radio. A mobile was programmed to transmit cells continuously and was moved around the building. A fixed base station was then used to monitor reception strength and cell integrity. The ORL Active Badge system [29] was used to record the mobile's position.

One of the results of the coverage trials are shown in Figure 9 with antenna diversity enabled. It can be seen that the radio physical layer supports the original requirement for 10 metre pico-cells. Reception is more a function of obstructing walls and floors than antenna range. The location of the base station is such that the signal propagates well along the central corridor through the length of the building.

[5] Xilinx is a trademark of Xilinx Inc.

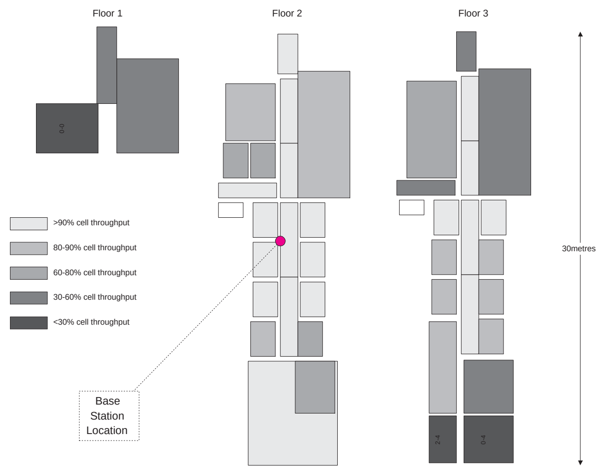

Figure 9: Coverage trial results

It must be emphasised that propagation is dependent on the nature of the building. In a large open-plan office with thin partitions and large windows a base station will have good coverage. In a building which is predominantly of metal construction, or which has internal walls of aluminium-backed plaster board, propagation will be more constrained. The ORL building is of a solid brick and stone construction, where partition walls have a higher attenuation than a more modern building.

The coverage map indicates that a single floor of the building can be covered by two suitably located base stations. Conversely the system bandwidth could be increased by spreading around a larger number of base stations on different frequency channels. Frequency reuse will be possible at opposite ends of the building.

6.5 Multimedia Experiments

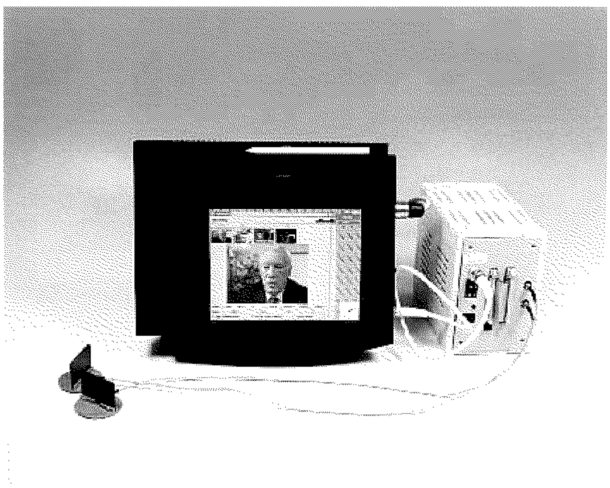

Figure 10: Prototype system demonstration

An important goal of the wireless network is to support streams of continuous media. To demonstrate this capability we have attached an ATM display device, the Video Tile, to the mobile device. The tile has an ATMos ATM processor card, a framestore, and an LCD display with pen input. A number of tiles are deployed on the wired network, as part of the Medusa project [15].

The tile acts as a remote frame buffer which displays an X session relayed from virtual frame buffer running on a workstation on the wired network. TCP connections between the tile and the workstation are used to transmit screen updates, pen movements and button events. Medusa transport connections running over separate VCIs display live video streams on windows on the tile. The creation, deletion and display of the video windows are controlled by the X applications displayed on the tile. Thus, the ATM traffic consists of control data, burst data and continuous data running over a variety of transport protocols.

The application runs successfully with the radio link extending the wired ATM network. As the Medusa framework does not yet make full use of the ATM QoS mechanisms, some of the transport parameters were modified from those used on the wired network. Combined throughputs of up to 6 Mbit/s are realisable for multi-stream video displays, frame buffer updates and pen events. Figure 10 shows an experimental setup running with five video windows. This configuration uses around twenty active virtual circuits.

7 Conclusions

The ORL Radio ATM architecture and prototype system have been developed over a number of years, together with a wired ATM infrastructure. The design decisions have been guided by experience gained from a series of distributed multimedia systems running on our wired ATM network. Experimental installations have shown that, subject to the QoS available, the radio ATM system can transparently support multimedia applications developed for the wired ATM network.

We have presented an architecture for a radio-based ATM local area network that allows mobile devices to benefit from the advantages that ATM has hitherto made possible in the wired network. Standard ATM adaptation layers, signalling and routing are used throughout, entailing no re-engineering of existing infrastructure. Care has been taken in the design of the overall wireless architecture to consider the wide area ATM networking issues. Local mobility is supported efficiently, while global mobility is handled using proxy addressing and global ATM name services as they become available. The combination of overlapping pico-cells, careful MAC protocol design and an efficient handover strategy makes effective use of limited bandwidth and allows realistic multimedia communication with mobile devices.

References

[1] I. M. Leslie, D. R. McAuley, and D. L. Tennenhouse, "ATM everywhere?," *IEEE Network*, March 1993.

[2] W. Diepstraten, G. Ennis, and P. Belanger, "Wireless access method and physical specification, DOC: IEEE P802.11-93/190," November 1993.

[3] J. Kruys, "Hiperlan, applications and requirements," in *PIMRC '92*, IEEE, October 1992.

[4] J. Porter and A. Hopper, "An ATM based protocol for wireless LANS," Tech. Rep. OLR 94-2, Olivetti Research Ltd., 1994.

[5] A. S. Acampara and M. Naghshineh, "An architecture and methodology for mobile-executed handoff in cellular ATM networks," *IEEE Journal on Seletecd Areas in Communications*, vol. 12, October 1994.

[6] D. Raychaudhuri and N. Wilson, "Atm-based transport architecture for multiservice wireless personal communication networks," *IEEE Journal on Seletecd Areas in Communications*, vol. 12, October 94.

[7] K. Eng *et al.*, "A wireless broadband ad-hoc ATM local-area network," *Wireless Networks*, vol. 1, pp. 161–174, 1995. BAHAMA.

[8] A. Saleh and R. Valenzuela, "A statistical model for indoor multipath propagation," *IEEE Journal on Seletecd Areas in Communications*, vol. 5, pp. 138–146, February 1987.

[9] R. Bultitiude, P. Melancon, and H. Zaghoul, "The dependence of indoor radio channel multipath characteristics on transmit/receive ranges," *IEEE Journal on Seletecd Areas in Communications*, vol. 11, no. 7, 1993.

[10] H. Hashemi, "The indoor radio propagation channel," *Proceedings of the IEEE*, vol. 81, July 1993.

[11] A. Saleh, J. Rustako, and R. Roman, "Distributed antennas for indoor radio communications," *IEEE Trans. Comm.*, vol. COM-35, pp. 1245–1521, December 1987.

[12] ATM Forum, *ATM User-Network Interface Specification Version 3.0*. Prentice Hall, 1993.

[13] "P-NNI draft specification," Tech. Rep. 94-0471R7, ATM Forum, March 1995.

[14] A. Alles, "ATM internetworking," tech. rep., CISCO Systems Inc, May 1995.

[15] S. Wray, T. Glauert, and A. Hopper, "Networked multimedia: The medusa environment," *IEEE Multimedia*, vol. 1, pp. 54–63, Winter 1994.

[16] P. Mockapetris, "Domain names—concepts and facilities," *RFC 1034*, November 1987.

[17] P. Mockapetris, "Domain names—implementation and specification," *RFC 1035*, November 1987.

[18] Object Management Group, *The Common Object Request Broker: Architecture and Specification*, December 1991. Revision 1.1, OMG Document Number 91.12.1.

[19] J. Ioannidis, D. Duchamp, J. Maguire, and Q. Gerald, "Ip-based protocols for mobile internetworking," *Computer Communication Review*, vol. 21, September 1991.

[20] R. Steele, J. Whitehead, and W. C. Wong, "System aspects of cellular radio," *IEEE Comm. Mag.*, vol. 33, pp. 80–86, January 1995.

[21] A.J.Viterbi *et al.*, "Soft handoff extends CDMA coverage and increases reverse link capacity," *IEEE Journal on Seletecd Areas in Communications*, vol. 12, October 1994.

[22] I. Rubin, "Access-control disciplines for multi-access communication channels: Reservation and TDMA schemes," *IEEE. Trans. Inf. Th.*, vol. 25, pp. 516–536, Sept 1979.

[23] S. Nanda, D. J. Goodman, and U. Timor, "Performance of PRMA: a packet voice protocol for cellular systems," *IEEE Trans. Veh. Tech.*, vol. 40, pp. 584–598, Aug. 1991.

[24] P. F. M. Smulders and C. Blondia, "Application of the asynchronous transfer mode in indoor radio networks," in *PIMRC '94*, vol. 3, pp. 839–843, IEEE, September 1994.

[25] M. J. Karol, Z. Liu, and K. Y. Eng, "An efficient demand-assignment multiple access protocol for wireless packet (ATM) networks," *Wireless Networks*, vol. 1, no. 1, pp. 267–279, 1995.

[26] C.-S. Chang and J. A. Thomas, "Effective bandwidth in high-speed digital networks," *IEEE Journal on Seletecd Areas in Communications*, vol. 13, pp. 1091–1100, August 1995.

[27] G. Cunningham, "Delay versus throughput comparisons for stabilized slotted ALOHA," *IEEE Trans. Comm.*, vol. 38, no. 11, pp. 1932–1934, 1991.

[28] ETSI, "European Telecommunications Standard prETS 300 328," tech. rep., ETSI TC-RES, 1994.

[29] R. Want and A. Hopper, "Active badges and personal interactive computing objects," *IEEE Trans. Consumer Electronics*, February 1992.