The Intelligent Networked Airport (TINA)

Jon Crowcroft, Malcolm Scott
University of Cambridge
Computer Laboratory

Paul Brennan
UCL
Electronic Engineering

Jaafar Elmirghani
University of Leeds
Electronic Engineering

Richard Penty, Ian White
University of Cambridge
Electrical Engineering

ABSTRACT

This paper is about a very large sensor network system being researched, designed and constructed for next generation advanced airports such as Heathrow T5. There are at least three novel aspects to the network. Firstly it has a radio-over-fibre backhaul, with a high level of resilience both in the physical layer and in the control plane. Secondly, we are deploying *active* RFID tags to track people and devices, as well as examining other prospects for such devices, and alternative tag types. Thirdly, the system is massive in terms of number of devices and capacity, and this poses major challenges to current technologies. In the paper we concentrate on one particular design that we are prototyping.

Keywords

sensor networks, RFID, radio-over-fiber

1. INTRODUCTION

Diverse applications are expected to appear in the future with complex and often variable service requirements, traffic profiles and user expectations requiring extremely advanced adaptive computing and communication systems to provide users with mobile, secure and automatic means of conducting business. A prime application area is in international travel which continues to grow supported by a significant investment in infrastructure with projects such as Heathrow Airport Terminal 5 and Dubai International Airport, two of the largest construction projects on the planet, both undertaken by Laing O'Rourke, one of the industrial collaborators of this project.

An intelligent, adaptive, self-organising wired wireless infrastructure is essential in this environment. It is anticipated that the considerable growth in the complexity of this infrastructure will not just be due to the proliferation of established fixed equipment such as wireless base stations, surveillance cameras (which will number in the thousands), security detection equipment and terminal equipment. However the requirements will also be for much wider deployment of more compact portable equipment, for example, location and control equipment on a wide range of transportation equipment.

Radio frequency identification (RF-ID) tags — supported by a transparent optical-RF network — can be used to sense, locate and track an array of objects including luggage, mobile assets and commercial goods and can provide additional features such as boarding pass auto-tags, access control tags and tagging support for airport commercial zones. These RF-ID tags are expected to operate at low data rates, typically 64 kbit/s, however an airport environment can contain a few million tags. Mobile biometric sensors will be widely deployed in this environment providing advanced features. A range of fixed and mobile terminals will provide additional security measures such as chemical detection and analysis, while other terminals, fixed and mobile, will support passenger information and entertainment services on transit.

The infrastructure will support an array of personal passenger and staff wireless media-rich devices. The wired wireless network envisaged will thus be huge and complex supporting perhaps 10 million tags at 64 kbit/s per tag, 1000 fixed cameras and 500 mobile cameras each operating at 6 Mbit/s MPEG2, 500 mobile biometric units operating at 10 Mbit/s, and 20 Gbit/s passenger and facility data communications, an aggregate data rate of 764 Gbit/s in a relatively local access environment.

This is beyond the capability of any current network and research is needed to understand the principles upon which an effective system could be constructed.

This paper is about the design and initial prototype system being carried out by the Universities of Cambridge, London and Leeds, together with a consortium of companies building the terminals in the airports mentioned above.

The application we are targetting is a smart tracking system to allow passengers and their bags to be located to within 1 metre at all times (around 1 second granularity) in any 10 metre space within the airport. The economics of missing take-off slots due to non-arrival of passenges and the necessity to de-plane their bags, and seperately, the economics of returning bags delivered to the wrong location, to passengers arriving at their destination, are each sufficient to justify the deployment of a tracking system. Obviously, passenger assistance (e.g. guidance from nearby screens in event of gate changes) and security applications also are indicated with the technology.

First we describe the radio tagging system which is very novel, then the organisation of flows and use. The subsequent sections discuss the use of a protocol agnostic, two layer architecture, comprising a physical layer radio-over-fibre network for carrying all airport radio signals back to central data processing resources, and a decentralised topology control protocol based on an extension to ideas in layer two bridged networks. The overall network can be thought of as in some sense analogous to a multi-protocol λ switched system where the segments of spectrum in the air are swtiched analog over the fibre backhaul, but the control plane of the backhaul is digital. The system is highly futureproof due to no dependence on any specific data plane or control. Of course, it is highly optimised to the deployment scenario, but would be applicable to many large building sites.

2. ACTIVE RF-ID TAG ARCHITECTURE

The primary wireless system that we are building anew (as opposed to pre-existing demands from GSM, WiFi, etc) is the location tagging system, based on active RF-ID, attached to all bags, vehicles and on passenger boarding passes. We next describe this technology in some detail, as it is significantly different from previous well known passive RF-ID systems.

2.1 Description

The TINA tag location system is based on a unique concept comprising active "far-field" [3] tags that radiate a pulsed, coded, wideband FM signal, which is used both to identify and locate each tag by time-of-flight measurement. Several time-of-flight (and hence range) estimation strategies have been considered including time of arrival (TOA), requiring an active tag transponder that re-radiates the reader signal in order to measure the round trip delay, and time difference of arrival (TDOA) in which the tag operates in transmit-only mode and the delay between the arrival of the signal at a number of pairs of readers is measured. The TDOA approach, shown in figure 1, offers simplicity of tag and reader design and avoids the need for synchronisation between reader units; it does, however, require one additional reader unit in order to determine location. Other possible location strategies are based on angle of arrival estimation (AOA), which is vulnerable to error due to propagation artefacts [2] and signal strength estimation, which offers a low-cost but crude indication of range. Combinations of these techniques may also be attractive, for instance TOA combined with AOA is able to perform location with just one reader unit, but is demanding in terms of tag and reader hardware and processing.

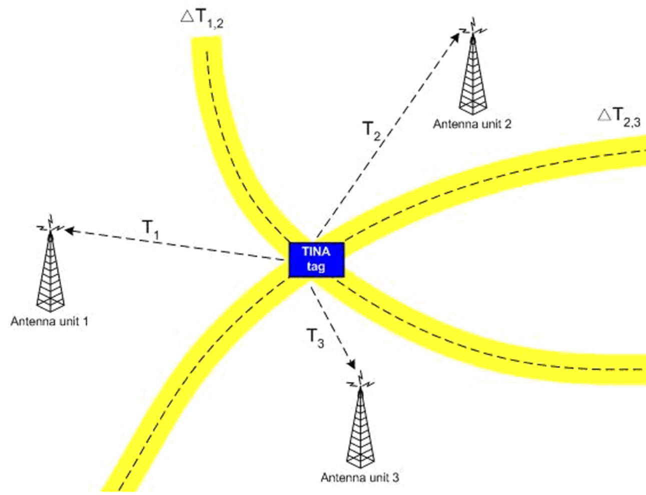

Figure 1: The TINA TDOA location approach

After careful consideration, the TDOA option has been chosen for TINA since it represents a good overall compromise between accuracy and ease of implementation, particularly lending itself to economical tag realisation. Further economies may be made, for instance, using recent advances in RF-ID tag antenna design [11]. The TINA tag uses a novel protocol based on the coded transmission of a set of chirped-FM pulses, to fulfil both identification and tracking functions where, for instance, a logic '1' may be represented by an upchirp and a logic '0' by the absence of a chirp. Figure 2 shows the general arrangement in which each tag transmits a short, coded chirped-FM burst that is processed at adjacent pairs of reader units to derive TDOA information by means of de-ramping the chirp signals, at a 480 MHz IF frequency, in a similar manner to the technique commonly used in FM-radar systems. However, in contrast to radar systems, the time delay may be either positive or negative and it also may be very close to zero resulting in a de-ramped pulse well below one cycle of a sinusoid. These complications can be readily resolved using a complex de-ramp arrangement, as indicated in figure 2. The parameters for TINA are as follows: 2.4 GHz (ISM band), 1 μs chirp pulse length, 85 MHz chirp sweep bandwidth, 144 μs data burst duration (128 bits plus pre-amble) and 0.5 s mean burst repetition interval. This approach offers the benefit of good operating range, high accuracy (dependent on the sweep bandwidth) and ease and economy of tag implementation since there is a relaxed frequency tolerance requirement. It also allows many hundreds or thousands of tags to operate in the same location with frequent updates, typically at least once per second.

2.2 Location considerations and resolution

The transmitted pulses consist of a sequence of linear FM chirps of duration T_c and swept-frequency bandwidth B, as shown in figure 3. The instantaneous frequency variation may therefore be expressed as

$$f(t) \quad = \quad \frac{Bt}{T_c} \quad \text{for} \quad -T_c/2 \le t \le T_c/2 \tag{1}$$

This signal arrives at a given pair of elements with a small differential time delay, Δt, corresponding to the differential path difference between the tag and the pair of elements. The signals are then de-ramped at the receiver, producing a pulse at a much lower frequency indicative of the time delay,

$$\Delta f \quad = \quad \frac{B\Delta t}{T_c} \tag{2}$$

Measurement of this frequency therefore allows very straightforward calculation of the differential time delay and hence differential range, as follows

$$\text{Differential range } R \quad = \quad \frac{cT_c\Delta f}{B} \tag{3}$$

Also, the number of cycles of sinusoid present in the de-ramped pulse is

$$N_{\text{cycles}} \quad = \quad \Delta f T_c \quad \equiv \quad B\Delta t \quad \equiv \quad \frac{R}{c} \tag{4}$$

and this may be a very small quantity in the case of low differential range. For the TINA parameters, assuming a spacing of 20 m between tag reader units, the de-ramped frequency range is ± 5.67 MHz. The number of cycles of sinusoid in the de-ramped pulse is 0.28 cycles per metre of differential range, which confirms the need for a processing technique capable of accurately estimating the frequency of a short sample of sinusoid.

The de-ramped pulse duration is very slightly reduced as a result of the delay between the received pulses, by a fractional amount $\Delta t/T_c$. Again for the TINA parameters, the maximum reduction in pulse width is some 6.7%. The de-ramped pulse may be processed using standard FFT techniques, giving a sinc response with a -3.9 dB width of $1/T_c$, resulting in the classic range resolution limit of c/B – which is 3.5 m in the case of the TINA parameters. However, it is possible to improve on this rather modest figure by judicious processing of the de-ramped signals. For instance, zero-padding FFT interpolation of the IQ de-ramped signal should yield an order of magnitude improvement under reasonable signal-to-noise conditions, and pulse-to-pulse averaging may provide further resolution improvement and greater noise immunity. The challenge is to achieve consistent, high range resolution under less favourable

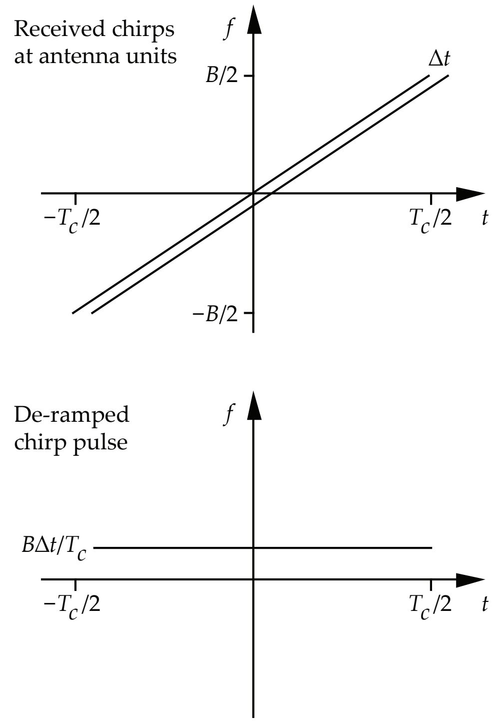

Figure 3: Illustration of the FM chirp de-ramp process

signal-to-noise conditions, though we are confident that sub-metre resolution is entirely feasible.

The FM chirp linearity requirement may be found as follows. Considering the more general case of a chirp waveform of slope k:

$$f \quad = \quad kt \tag{5}$$

for a differential range, R, with corresponding differential time delay, $\Delta t = R/c$, we have

$$R \quad = \quad c\Delta t \tag{6}$$

$$\quad = \quad \frac{c\Delta f}{k} \tag{7}$$

Now, if the chirp slope is nonlinear, so that k varies by amount Δk over the duration of the pulse, then the variation in the indicated range will be

$$\Delta R \quad = \quad -\frac{c\Delta f}{k^2}\Delta k \tag{8}$$

If this variation is small relative to the range resolution, of c/B, then there will be little adverse effect; however if this variation is comparable to the range resolution then performance will be compromised. Equating ΔR with the range resolution provides a convenient measure of the acceptable linearity,

$$-\frac{c\Delta f\Delta k}{k^2} \quad = \quad \frac{c}{B} \tag{9}$$

$$\Rightarrow \qquad \frac{\Delta k}{k} \quad = \quad -\frac{k}{B\Delta f} \quad \equiv \quad -\frac{1}{T_c\Delta f} \tag{10}$$

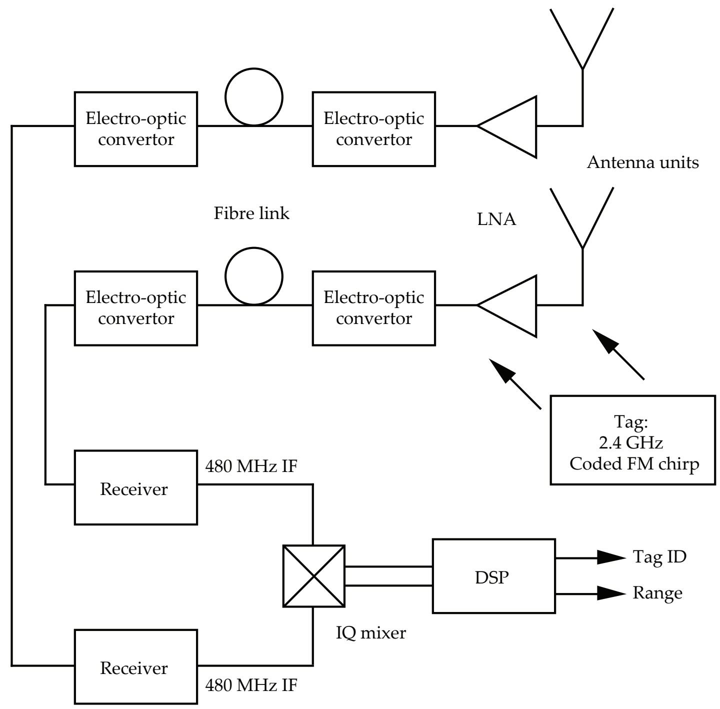

Figure 2: The TINA tag reader/location system

This indicates that the required linearity is equal to the reciprocal of the time-bandwidth product of the de-ramped pulse. Since the de-ramped pulse bandwidth is relatively small, this requirement is very modest. For the TINA parameters, the required linearity is 17.6%, which may be very easily achieved with a free-running, uncompensated voltage-controlled oscillator.

2.3 Tag population analysis

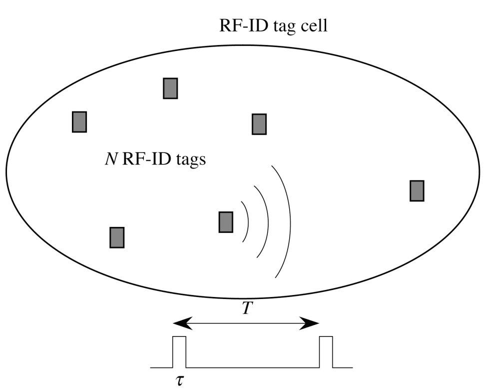

Figure 4: Illustration of RF-ID tag arrangement in a given cell

An issue of direct relevance to this system is the possibility of data loss due to clashes between the tag bursts. Some such data loss is inevitable, depending on the number of tags in use in a given area and the duration of the data bursts. This has been carefully modeled [9] as follows. The basic arrangement is shown in figure 4 and consists of a number, N, of far-field RF-ID tags in a single reader cell, each radiating a short data packet of duration τ and with repetition interval, T. The system is asynchronous so that the timing of these data bursts is unrelated and varies randomly from tag to tag.

Considering one particular tag sending a single data packet; a collision of some degree will occur if any other tag sends a data packet within $\pm\tau$ s of this transmission. The probability that this data packet will avoid the effects of such collisions from any of the other $(N-1)$ tags, thus ensuring there is no resulting data loss, is therefore given by

$$P(\text{no collision}) = \left(1 - \frac{2\tau}{T}\right)^{N-1} \tag{11}$$

The mean update interval, over a population of tags, is clearly increased by this finite probability of interference from neighbouring tags, and taking the worst-case situation in which all collisions result in loss of the data packet, the mean update interval may be expressed as

$$T_{\text{mean}} = \frac{T}{\left(1 - \frac{2\tau}{T}\right)^{N-1}} \tag{12}$$

These results are illustrated in figures 5 and 6, for the TINA system parameters of 1 Mbit/s data rate, a data packet duration of 150 μs and a tag population of between 500 and 2000. It is clear, from figure 5, that the probability of avoiding data collisions increases with longer repetition intervals and with smaller tag populations; however, from figure 6, for a given tag population an optimum repetition interval may be chosen to give minimum mean update interval and hence maximum data throughput. This optimization process may be formulated as follows:

$$\frac{\partial T_{\text{mean}}}{\partial T} = \frac{\left(1-\frac{2\tau}{T}\right)^{N-1} - \frac{2\tau}{T}(N-1)\left(1-\frac{2\tau}{T}\right)^{N-2}}{\left(1-\frac{2\tau}{T}\right)^{2N-2}} \quad (13)$$

$$= 0 \text{ at } T = T_{\text{opt}} \quad (14)$$

resulting in

$$T_{\text{opt}} = 2N\tau \quad (15)$$

and

$$T_{\text{mean}}(\text{min}) = \frac{2N\tau}{\left(1-\frac{1}{N}\right)^{N-1}} \rightarrow 2eN\tau \quad (16)$$

$$\equiv eT_{\text{opt}} \text{ for large } N \quad (17)$$

Taking again the values used in figures 5 and 6, for a tag population of 1000, the optimum choice of repetition interval is thus 0.3 s and the corresponding minimum update interval is 0.82 s. For an optimised repetition interval such as this, substituting equation 15 into equation 11, the probability that a given data packet suffers no collision is

$$P(\text{no collision, optimised}) = \left(1-\frac{1}{N}\right)^{N-1} \quad (18)$$

$$\rightarrow \frac{1}{e} \text{ for large } N \quad (19)$$

and so, remarkably, maximum data throughput coincides with some 63% loss of data packets due to collisions. The prototype system is actually designed with a 0.5 s repetition interval equating to a mean update interval of 0.9 s — indicating that the position of all tags can be determined and updated on a second-by-second basis. Thus the system can easily accommodate 1000 tags in any given cell, which is probably close to the limit of the number of people who can possibly be squeezed into a 10 m radius area!

3. USER FLOW DYNAMICS

In this section, we present initial thoughts on demand and organisation of the network.

The airport requirements for a complex but flexible converged wired/wireless network, which is easy to use, cost effective and robust, call for a fresh look at network design rules. Of particular interest is the development of new design rules for hybrid wired-wireless self-organising intelligent networks that minimise human intervention, simplify the network and increase its capacity, throughput, resource utilisation efficiency, reliability and resilience. As self-organisation can affect

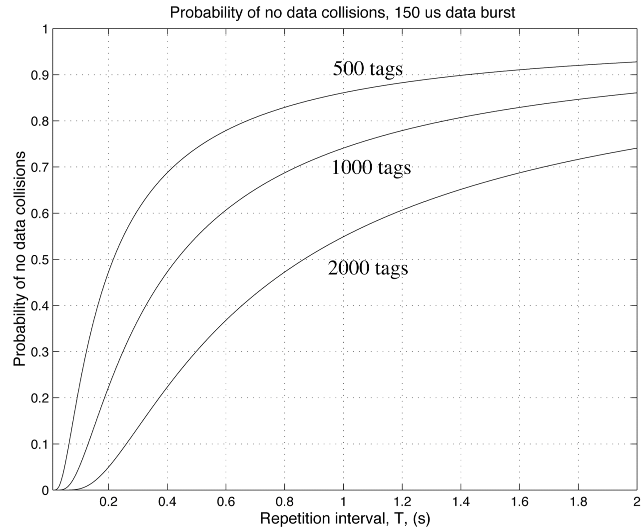

Figure 5: Modeled variation of tag clash probability with tag population

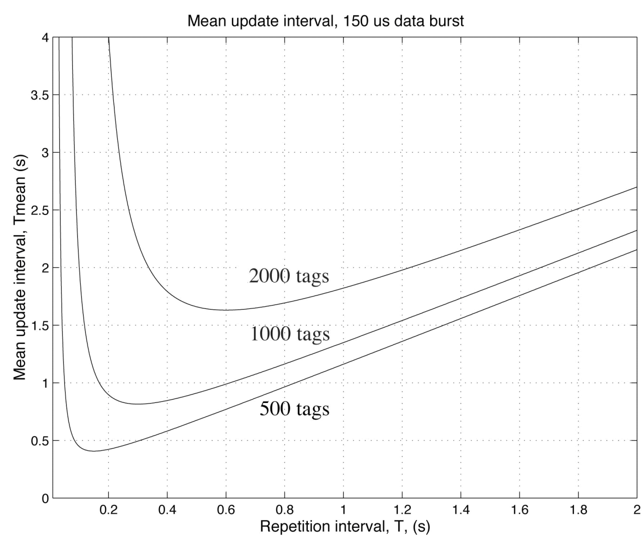

Figure 6: Modeled variation of mean update interval with tag population

all the network layers, the design rules deserve special consideration. Self-organisation has been studied to some extent in ad-hoc networks, however the airport scenario under consideration must combine indoor multi-cell networking (for coverage) and mobile mesh networking within and between the cells to introduce new degrees of freedom, enable nodes self organisation, improve capacity and enhance resource (eg. frequency) utilisation efficiency.

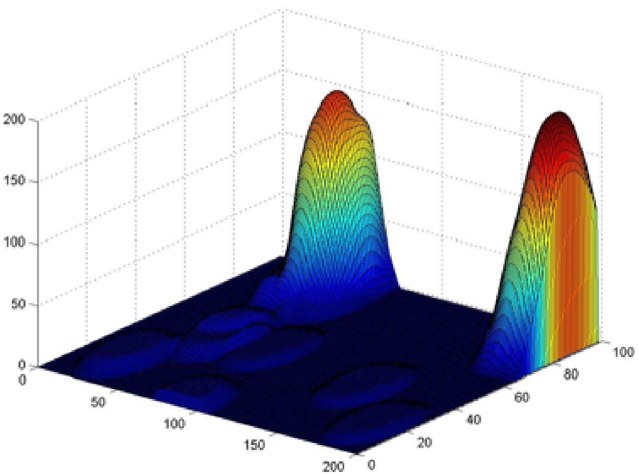

Figure 7: Space-time distribution of bandwidth requirements in airport section at time T minutes following departure from a gate

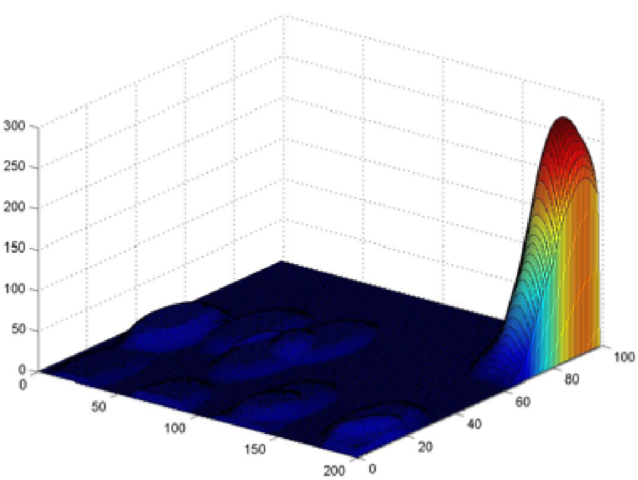

Figure 8: Space-time distribution of bandwidth requirements in airport section at time $T + 15$ minutes following departure from a gate

We have modelled an airport section to gain an appreciation of the passenger spatial flow dynamics and the variation expected in data rate and bandwidth requirements. This will inform the network architecture design and the higher layer self organisation functionalities. The airport section considered has dimensions of 200 m × 100 m, there is a single entry point to this section through a security check point. There are 10 gates and 10 shops of varying size. The node (passenger) speed is 1 m/s. Passenger arrivals are Poisson distributed at the entrance with a mean arrival rate of 4 to 5 passengers per minute. Departure is through one of the 10 gates with 1 flight leaving the system of 10 gates every 20 minutes carrying 100 passengers, a departure rate of 5 passengers per minute, Poisson distributed. The gates are equally likely and the choice of destination gate is uniformly distributed among the 10 gates. Passengers make a number of stops at shops after entry. The number of stops is assumed Gaussian distributed with a mean of 3 stops and a standard deviation of 0.5 (ie most passengers do 1.5 to 4.5 stops at the shops). All the shops are equally likely for a stop. Passengers spend 1 to 10 minutes in each shop, uniformly distributed. Passenger motion is graph based with corridors and shop entry points representing branching points (with different branching probabilities). Passenger motion within a shop is assumed to follow a random walk. Passengers make voice calls (64 kbit/s, Poisson distributed, 5 minute mean call duration), video calls (6 Mbit/s Pareto distributed, mean duration 5 minutes, mean interarrival time 50 minutes and data calls (10 Mbit/s Pareto distributed, mean duration 15 minutes, mean interarrival 40 minutes) . This study has shown that Antenna access points experience large variation in data rate from few kbit/s to over 200 Mbit/s in a future airport design, as illustrated in figures 7 and 8.

Such a new structure provides new challenges in scalability, optimum routing, mobility, resilience and security. We are developing novel network architecture and topology design rules inspired by lessons learned in natural systems where such complex systems are found to self-organise into *scale-free* formations that obey *small-world* principles. A small-world network architecture will ensure that the separation of nodes (in hops) is always small throughout the network and a scale-free formation will ensure that key network performance parameters such as delay are little affected as the wired-wireless network scales.

4. THE RADIO-OVER-FIBRE BACKHAUL

In this section, we describe the backhaul network design for the airport. Instead of assuming an all-IP Internet style of communications architecture, we have to accommodate a very wide range of different protocol architectures from radar from air traffic, through cellular telephone, and on to wireless data. To this end, we have to be compeltely protocol agnostic.

As wireless services, such as mobile telephony, wireless networking and private mobile radio (PMR), become ever more pervasive, there is increasing demand for mobile coverage in indoor locations such as shopping malls, airports and office buildings. As data rates and carrier frequencies become higher, the range and coverage of basestations, particularly in a complicated radio environment such as a building, inevitably decreases. In contrast to the old days (of the 1980s) when a phone rang and you went indoors to answer it, today your cellular phone rings and you go outside to improve reception! Consequently there have been many developments in the field of improving the indoor coverage of such wireless networks.

The network architecture which provides maximum coverage at low installation and maintenance cost is a distributed antenna system with a central distribution unit and several remote antenna units. However, since cellular signals require gigahertz range carrier frequencies, it has been assumed until recently that the fibre DAS systems either had to use expensive, custom-installed single mode fibre or they had to use down-converted IF transmission over the much cheaper installed base multimode fibre (MMF) which has a 3 dB frequency of up to about 1 GHz for the lengths of fibre used in airports. However, the TINA network is envisaged to support a wide variety of radio services — the primary ones of interest being RF-ID in the 850 MHz to 950 MHz, 2.4 GHz to 2.5 GHz and 5.8 GHz Bands, fixed cameras, wireless enabled cameras in the 5.8 GHz band, wireless enabled biometric devices in the 5.8 GHz band, wireless enabled IT devices in the 2.4 GHz and 5.2 GHz IEEE 802.11 bands and cellular communications including TETRA at 450 MHz, GSM at 900 MHz and 1.8 GHz and 3G at 2.25 GHz — and so these approaches are either expensive or not technically viable.

A much better and cheaper solution would be to connect the basestations directly to the antennas using the already installed MMF. Work at Cambridge and UCL sought to develop the SCM technology developed by the Cambridge group. This has resulted in radio over multimode fibre links demonstrated at carrier frequencies up to 20 GHz — 20 times the fibre bandwidth[8]. Whilst allowing multiple radio services[7, 6], alongside conventional baseband ones such as Ethernet, to be transported over the same wired infrastructure, the approach has the added benefit that all service provision hardware can be centralised.

To demonstrate the efficacy of such an approach, figure 9 shows an experimental distributed antenna network installed in a laboratory environment. An 802.11g access point was used to generate WLAN signals, and a Rohde & Schwarz SMIQ vector signal generator to simulate a 3G base station. The 802.11g and 3G signals are electrically combined into a broadband radio channel before splitting the electrical signal to directly modulate 3 lasers in a Zinwave 2700 Primary Hub. Each optical signal is then fed to an RU along 330 m of 62.5 μm FDDI grade MMF. The long fibre length was chosen to represent a worst case fibre run from a basement equipment room to another area of a building. Comparing the performance of this network to one using a single antenna, the number of measurement locations with greater than 19.5 Mbps throughput is increased from 40% with a single antenna to 80% by using the DAS. Figure 10 shows that 10 dB more transmitted power would be required from the single antenna to achieve similar performance to the three antenna DAS.

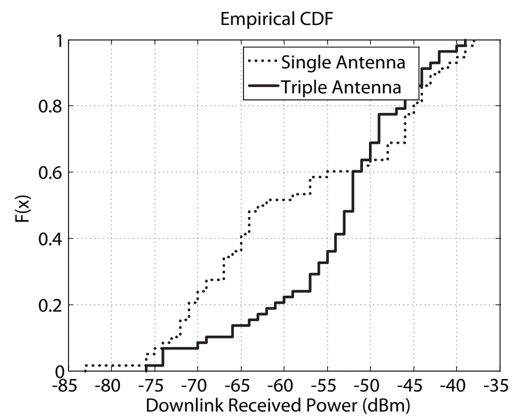

Figure 10: Improvement from one to three antenna DAS

5. NETWORK RESILIENCE

We must also consider the architecture of the network at a higher level than the physical layer. It would be sensible to start with a proven, widely-used protocol; therefore we have chosen Ethernet as a basis for the protocol stack, as it is ubiquitous and has existing wireless extensions in the form of IEEE 802.11. Operating at the data link layer allows many important features to be implemented in a way that is transparent to higher-layer protocols such as IP, and allows the network to be largely protocol-agnostic. (We cannot assume that all devices on the network will wish to use or be capable of using IP.)

However, Ethernet has a number of failings in networks containing a large number of nodes, so we are carrying out research in order to investigate these limitations and to adapt Ethernet in order to improve its scalability. It is important to maintain compatibility with existing Ethernet equipment at the edge of the network, allowing computers to use the network unmodified with either a wired or a wireless connection.

5.1 Ethernet routing

One major limitation of Ethernet is its inability to make efficient use of a meshed network topology. Using RSTP, existing Ethernet networks of arbitrary topology

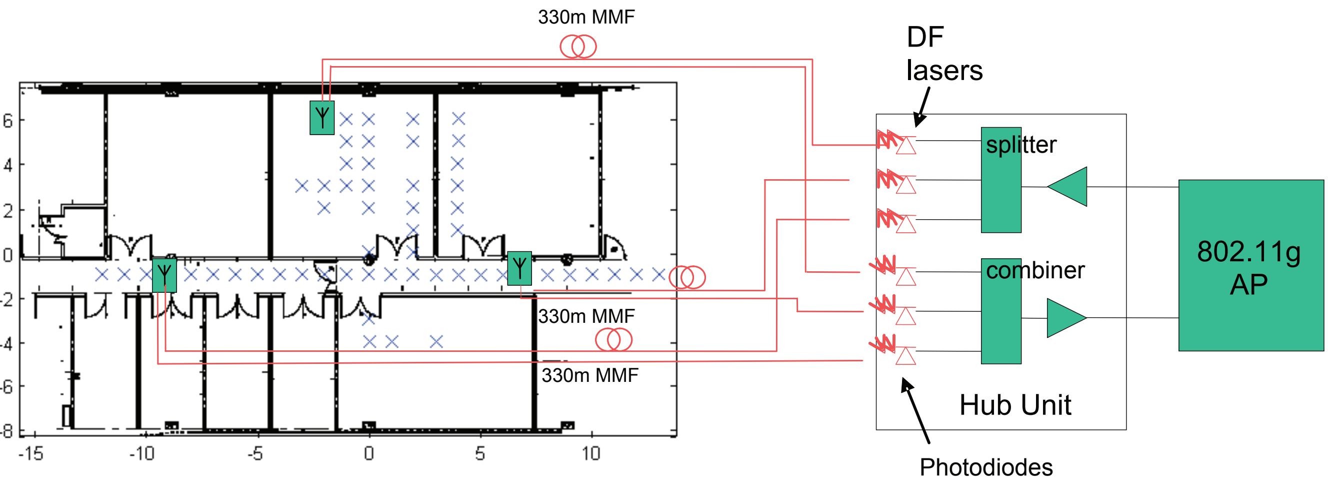

Figure 9: Schematic layout of the DAS showing antenna locations and measurements points (blue crosses) and (right) empirical Cumulative Probability Distribution of the experimentally measured received signal strength

are made to run as a tree by automatically disabling links which form loops. This often results in heavy congestion at the root of the tree and high latencies as frames are forced to take paths which are often considerably longer than necessary. Furthermore, the details of these protocols limit the depth of any spanning tree produced and hence the diameter of the network.

Hence an adaptation of Ethernet is needed which routes frames directly along the best path, whilst retaining RSTP-like ability to rapidly route around failures. Once such adaptation is SmartBridge [14] which attempts to solve this by using a separate spanning tree for each transmitting node; conversely LSOM [4] and Rbridges [13] involve running a link-state routing protocol (IS-IS [12] in the case of Rbridges) at the data link layer. LSOM has no mechanism to prevent the count-to-infinity scenario which may occur during a topology change; SmartBridge and Rbridges both a facility to reliably and consistently update the routes through the network in a loop-free manner and are hence superior in this respect.

We will be conducting a comparison of the approaches taken by these three approaches in order to determine which is the most effective in our scenario. This will involve the simulation of each protocol in operation on a large mesh network, paying particular attention to the effects of different kinds of failures.

5.2 Broadcast

Of the extensions to Ethernet described above, only Rbridges attempts to address the problem of broadcast traffic, which could consume a sizeable proportion of the network's capacity when a large number of nodes are present. Unfortunately broadcasting is fundamental to the interoperation of Ethernet with current higher-layer protocols such as IP via layer-bridging protocols such as ARP. Rbridges address the case of ARP, but not any other broadcasting protocol. A more disruptive approach is suggested by Myers et al. [10] which requires the protocol used by end nodes to be modified so that directory services are used in place of broadcasted queries. This is infeasible for TINA as compatibility with existing systems is of great importance.

We are therefore investigating means of reducing the amount of traffic which must be broadcast on an Ethernet network. One possible approach would be to extend the approach taken by Rbridges to more protocols, or ideally to generalise it to work with as-yet unknown broadcasting protocols.

5.3 Address table reduction

Another barrier to Ethernet's scalability is the MAC address table which every switch must maintain. In current switches the capacity of this table is of the order of 8000 entries [1]; therefore, problems may become apparent if the number of active nodes on the network exceeds 8000.

These problems can be alleviated to an extent by introducing structure and predictability to MAC addresses, which at present form a flat namespace in which any MAC address could appear anywhere on the network. We have begun work on a scheme entitled “Multi-level Origin-Organised Scalable Ethernet” (MOOSE) in which edge switches rewrite the source MAC address of frames entering the network; the new address is hierarchical and comprises at least a switch identifier and a node identifier allocated by that switch. This allows other switches to just maintain a table containing each known switch identifier once, ignoring the node identifier. The effect of this is that the address table would now scale with the number of edge switches rather than the number of nodes.

This goal has been attempted before by Hadžić [5], and to an extent in the MPLS standard [15], but instead of rewriting MAC addresses these previous approaches encapsulate the Ethernet frame intact inside another frame which uses the new addressing scheme. The receiving node does not see the hierarchical address, and hence will send any replies to the original flat-namespace address. Hence that node's switch must somehow convert this back into a hierarchical address. Hadžić suggests two approaches. Each switch could maintain a table of mappings; however this merely substitutes one large table for another. Frames where the correct hierarchical address is unknown could be broadcast; but this just exacerbates the problems with broadcast which we highlighted in the previous section.

Our solution avoids this problem by leaving the new hierarchical address in each frame's source address field when delivering the frame to its destination. Hence that node will use the hierarchical address as the destination of the response and no destination address conversion needs to take place.

5.4 Quality of Service

It will be important for the TINA network to prioritise different types of traffic in order to adhere to different quality of service (QoS) guarantees. In particular, QoS should be used to eradicate or at least alleviate the possibility of a denial of service attack. QoS would also useful for less security-critical aspects of the network such as ensuring that audio and video streams are not affected by less time-sensitive data — the TINA network will carry many such simultaneous streams for a variety of different applications such as VoIP, security, entertainment, and many others.

5.5 Test implementation

A small scale virtual network for testing has been set up using a number of Xen virtual machines running Linux and communicating with each other via a modified Linux kernel bridging driver. This allows us to test the interoperation of new protocols with standard unmodified Ethernet nodes, and will complement simulation of a larger network, allowing us to gain a reasonably complete picture of the effectiveness of our modifications to Ethernet and of the TINA network.

6. CONCLUSIONS, FUTURE WORK AND ACKNOWLEDGEMENTS

In this paper, we have presented the preliminary architecture and design for the Intelligent Networked Airport TINA system. There are two novel technical aspects to the network, namely that it uses a *Radio-over-Fibre* backhaul for the infrastructure which is protocol agnostic, and therefore unifies all of the wireless services in the airport space, and that it uses an *active* RF-ID tag system for location services, which results in high accuracy and timeliness. In addition to this, a layer two control plane provides resilience for the topology.

The system is dimensions for a high traffic load both from conventional radio sources such as cellular and wireless data, and from the sensor system itself. Initial models of mobile user flows are presented based on current analysis of airport use.

Future work will report on deployment experience and performance results. We are also considering ways of using more conventional passive RF-ID, and seeing if we can achieve the same coverage and accuracy. Other devices will be compared for performance and accuracy. We have also built a simulator to analyse the resilience of the network architecture and compare it with more conventional approaches.

We would like to acknowledge the support of the UK EPSRC who funded part of this work under their *Wired and Wireless Intelligent Networked Systems* programme.

7. REFERENCES

[1] 3Com Corporation. SuperStack® 3 Switch 4400 family data sheet. `http://www.3com.com/other/pdfs/products/en_US/400679.pdf`.

[2] R.S. Andrews. Antenna and other systematic effects on amplitude comparison monopulse systems. In *IEE Proc. Electronic Circuits and Systems*, volume 3, pages 103–108, May 1979.

[3] J.L.M. Flores, S.S. Srikant, B. Sareen, and A. Vagga. Performance of RFID tags in near and far field. In *Personal Wireless Communications*, pages 353–357, January 2005.

[4] R. Garcia, J. Duato, and F. Silla. LSOM: a link state protocol over MAC addresses for metropolitan backbones using optical ethernet switches. In *Second IEEE International Symposium on Network Computing and Applications*, 2003.

[5] I. Hadžić. Hierarchical MAC address space in public Ethernet networks. In *IEEE GLOBECOM*, volume 3, 2001.

[6] P. Hartmann, A. Bothwell, K. Cronin, R. Leeson, A. Loveridge, D. Parkinson, R. Ure, R. Penty, I. White, and A. Seeds. Wideband fibre-agnostic DAS using pluggable analogue optical modules. In *IEEE Topical Conference on Microwave Photonics*, 2006.

[7] P. Hartmann, R.V. Penty, I.H White, and A.J. Seeds. Dual-service wireless LAN antenna-remoting system using uncooled directly

modulated DFB laser diodes. In *Conference on Lasers and Electro-Optics*, 2004.

[8] P. Hartmann, P. Qian, A. Wonfor, R.V. Penty, and I.H. White. 1-20 GHz directly modulated radio over MMF link. In *IEEE Topical Conference on Microwave Photonics*, 2005.

[9] T.M. McCoy, P.V. Brennan, and R.J. Bullock. RF-ID receiver techniques for wideband synchronisation. In *IEEE International Conference on RFID*, March 2007.

[10] A. Myers, E. Ng, and H. Zhang. Rethinking the service model: Scaling Ethernet to a million nodes. In *Proc. ACM SIGCOMM Workshop on Hot Topics in Networking*, November 2004.

[11] P.V. Nikitin, S. Lam, and K.V.S. Rao. Low-cost silver ink RFID tag antennas. In *IEEE Antennas and Propagation Society International Symposium*, volume 2B, pages 353–356, 2005.

[12] D. Oran. OSI IS-IS Intra-domain Routing Protocol. RFC 1142 (Informational), February 1990. `http://www.ietf.org/rfc/rfc1142`.

[13] R. Perlman. Rbridges: transparent routing. In *INFOCOM*, volume 2, 2004.

[14] T.L. Rodeheffer, C.A. Thekkath, and D.C. Anderson. SmartBridge: A scalable bridge architecture. In *SIGCOMM*, 2000.

[15] E. Rosen, A. Viswanathan, and R. Callon. Multiprotocol Label Switching Architecture. RFC 3031 (Proposed Standard), January 2001. `http://www.ietf.org/rfc/rfc3031`.