

Application of Multiple-Wireless to a Visual Localisation System for Emergency Services

Costas Efthymiou¹, Sedat Gormus¹, Zhong Fan¹, Andrew Calway², Walterio Mayol-Cuevas² and Angela Doufexi²

¹*Toshiba Research Europe Ltd.*

Telecommunications Research Laboratory, 32 Queen Square, Bristol, BS1 4ND, UK

{costas, sedat.gormus, zhong.fan}@toshiba-trel.com

²*University of Bristol*

Merchant Venturers Building, Woodland Road, Bristol, BS8 1UB, UK

{andrew, wmayol}@cs.bris.ac.uk, a.doufexi@bris.ac.uk

Abstract—In this paper we discuss the application of multiple-wireless technology to a practical context-enhanced service system called ViewNet. ViewNet develops technologies to support enhanced coordination and cooperation between operation teams in the emergency services and the police. Distributed localisation of users and mapping of environments implemented over a secure wireless network enables teams of operatives to search and map an incident area rapidly and in full coordination with each other and with a control centre. Sensing is based on fusing absolute positioning systems (UWB and GPS) with relative localisation and mapping from on-body or hand-held vision and inertial sensors. This paper focuses on the case for multiple-wireless capabilities in such a system and the benefits it can provide. We describe our work of developing a software API to support both WLAN and TETRA in ViewNet. It also provides a basis for incorporating future wireless technologies into ViewNet.

I. INTRODUCTION

Recently context-aware and location-aware services have received a lot of attention among the mobile and pervasive computing research community. Various novel applications have emerged in both academia and industry, e.g. [1], [2]. Context awareness enables mobile systems and devices to sense their physical environment and adapt their behaviour accordingly. This can lead to enhanced user experience and new services. In particular, location awareness is now widely recognised in the industry as a key enabler for a broad range of value-added services and new business opportunities.

In this paper we describe our implementation of multiple wireless capabilities in a context-enhanced networked system based on the fusion of mobile vision and location technologies, the ViewNet system [3]. ViewNet aims to address some of the practical requirements of emergency and security services. Emergency service personnel often enter a scene without adequate maps. They report movements and surroundings to a control centre and this is used to direct the response. In ViewNet, sensor networks are deployed to provide distributed localisation and mapping. By fusing visual mapping with RF localisation (e.g. UWB), each operative will be located absolutely and in relation to key features. Visual mapping is done using the visual SLAM (simultaneous localisation and mapping) technology [4]. The objective of the project is to develop a demonstrator that operates over both TETRA

(Terrestrial Trunked Radio) and WLAN (802.11), and integrates GPS absolute positioning with vision-based localisation, 3D mapping and data from inertial sensors.

This paper presents justifications for the inclusion of multiple-wireless capabilities in ViewNet and its implications for the overall network control architecture. This is enabled by ULLA [5], the Unified Link Layer API that is proposed for use with ViewNet. A brief overview of ULLA is presented, together with the development and integration details of 802.11 and TETRA technologies into the ULLA framework.

The rest of the paper is organised as follows. Section II gives a brief background of the ViewNet project and the wireless interfaces involved. Section III discusses the benefits of providing a system such as ViewNet with multiple-wireless capabilities. Section IV describes how multiple wireless capability may be added to the system in a future-proof way through ULLA, while our implementation work of ULLA interfaces for TETRA and WiFi, as well as a RAT (Radio Access Technique) switch module are presented in section V. Some testing results are also reported. Finally, conclusions are drawn in section VI.

II. BACKGROUND

A. The ViewNet Project

To facilitate the quick mapping and information sharing of an incident area among security service operatives, the ViewNet project is developing technologies to allow the deployment of ad hoc wireless networks capable of supporting distributed localisation of users and user-assisted geometric mapping of the area. Users will be linked to a central control station, and their movements and actions will facilitate the automatic construction of an incident map, providing detailed information about spatial geometry, access points, critical areas, suspicious objects and evidence [3]. There will be a wide range of benefits such a system can offer, e.g. improvements on response times, evidence gathering, public safety and efficient resource deployment. It will also have an impact on other application areas employing sensors and context-aware networks. The key components of the ViewNet system are shown in Fig. 1 [3]. It can be seen that wireless communications and network control are at the centre of the overall system.

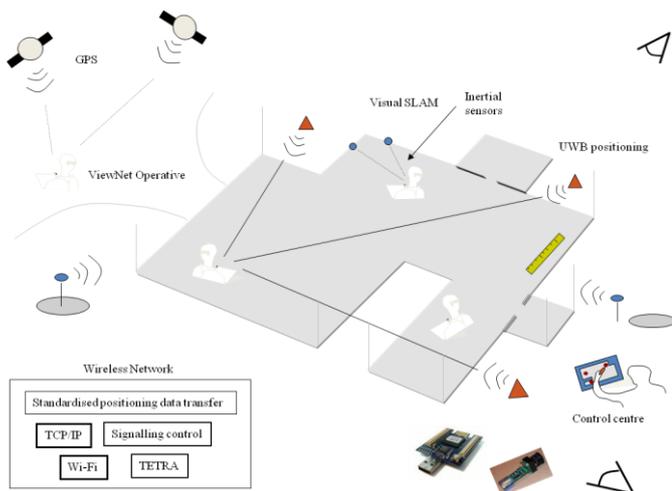


Fig. 1. The ViewNet system

The underlying wireless network in ViewNet is critical to the successful operation of the system. It is envisaged that this wireless network will operate over multiple technologies such as TETRA, WiFi, and GSM. Further, with future 4G high-speed wireless technologies (e.g. LTE, WiMax) on the roadmap, an efficient control architecture is required to enable coexistence and smooth switching among different wireless interfaces as and when they become available.

The development of the ViewNet project was guided by the basic scenario of a number of ViewNet-equipped operatives entering a location (e.g. a building) with an unknown internal layout and with unknown internal features and objects. This was intended to mimic the situation encountered by emergency service personnel in scenarios such as fire, forensic analysis, military activities, etc. As shown in Fig. 1, a number of operatives would enter an unknown area and begin observing the environment using visual SLAM technology [4]. Each operative is aware of their location through a combination of GPS and UWB positioning methods, and a number of position and mapping updates are sent back to the control centre and the other operatives, as necessary.

B. Wireless Interfaces in ViewNet

Central to the ViewNet concept is the idea that a ViewNet system should be able to use whichever wireless technology is available to it at any time – the available technologies to consider would depend on whatever the ViewNet user has available. The two main technologies that are currently being considered in the ViewNet project are the IEEE 802.11 wireless local area networks (WLANs) and TETRA. The reasons for this choice are as follows.

IEEE 802.11 is a very widely used wireless standard and as such, components for it are abundant and cheap. In addition to this, 802.11-based infrastructure is widespread in most offices and public buildings and it is extremely easy to deploy where there is no infrastructure already in place. In tandem with this, 802.11 devices perform quite well in many indoor wireless scenarios with high data rates being available.

TETRA is a digital trunked mobile radio standard developed by the European Telecommunications Standards Institute (ETSI) [8] and it is the PMR (Professional Mobile Radio) system of choice for most public services in Europe, such as the Police and the Fire Services. These are the kind of potential ViewNet users that have been identified and therefore the ViewNet system will have maximum applicability if it can be readily deployed over communications systems that are already in widespread use.

802.11 and TETRA can be considered to be complementary technologies for the purposes of ViewNet. 802.11 can provide high bitrate communications over relatively short distances (possible tens of megabits/s over a range of less than 100 m), whereas TETRA coverage can extend to multiple kilometres, albeit with achievable bitrates of less than 10 kbps.

Although there are just these two wireless technologies being considered for ViewNet use at the moment, there is a clear requirement for ViewNet to have the capability to use any wireless technology that may be available to a user at deployment time. Such a capability would greatly increase the utility of the ViewNet system and open it up to more potential future users of the system. This issue will be addressed in more detail in the next section.

C. Related Work

Most existing work in multiple-wireless has concentrated on the use of heterogeneous wireless resources as fallback or alternative network access options [6]. For example, in existing smartphones and laptops, 802.11, GPRS, HSDPA and Bluetooth can regularly be found. There is a wealth of methods being used for heterogeneous network access and for inter-technology handovers, but their use is usually limited to one data connection at a time. UMA is one of the technologies that operate in this way [7]. This paper discusses the *concurrent* use of heterogeneous wireless resources to increase the performance of the ViewNet system.

III. WHY MULTIPLE WIRELESS?

The primary aim of ViewNet is to provide security services with the capability to map and search incident areas as quickly as possible, yielding a ‘picture’ of the scene that indicates spatial layout and the location of critical areas, objects and evidence. ViewNet users are linked to a central control station, and their movements and actions within an area or building of interest facilitates the automatic construction of an ‘incident map’, providing detailed information about spatial geometry, access points, critical areas and suspicious objects.

ViewNet components and services could operate over a single wireless network – provided each of the network nodes are in range of the wireless access point and can transfer their data back to the control centre. Deploying a single 802.11 wireless network and having all of the communications go over that would be a simple and cheap solution; however, such an approach would suffer a number of disadvantages:

- **Reliability** – Reliability is a key requirement for emergency service systems. A single network would be unprotected against any failures, e.g. if the radio environment

the network is operating in is unfavourable, then there aren't many options available that can solve the problem.

- **Coverage** – Whilst 802.11 has relatively good coverage performance indoors in small buildings or single rooms, it is not so great in more complex/large buildings unless multiple access points are deployed throughout the building.

- **Range** – The quoted values for the range of 802.11 networks certainly sound impressive (over 100 m), but in reality these values can only be valid in ideal radio conditions, i.e. line of sight communications in open space and with no adverse multipath effects.

- **Energy** – 802.11 radios consume relatively small amounts of power in indoor environments. However, a large operating environment may require increased transmission power from the mobile nodes involved to ensure adequate connectivity, decreasing their useful battery life.

- **Device mix** – An organisation deploying a ViewNet system may have a mix of devices that are routinely used and not all of those devices may contain 802.11 interfaces.

- **Interference** – The frequency bands that 802.11 devices operate in are unlicensed spectrum bands, which can result in plenty of interference from other equipment.

How critical any of the above potential issues would be in a particular scenario depends of course on a number of factors such as the environment in question, the scale of building or area that needs to be mapped, energy levels of the mobile nodes, etc. One of the core ideas of ViewNet is to use multiple wireless technologies in order to mitigate some of the problems that have been discussed here, and as such, TETRA has been identified from the outset as the technology of choice to go along with 802.11. Although TETRA offers data rates that are up to four orders of magnitude lower than 802.11g, it does offer the significant benefit of already being used by most of the public and security services in the UK today. It is a European standard that is also being used by a large number of governments across Europe and so there is quite a high level of device penetration. Moreover, TETRA offers superior reliability and coverage as compared to 802.11 systems, with a single basestation being able to serve an area many times larger than 802.11 (several square kilometres).

Extending the idea of multiple-wireless further, i.e. encompassing other wireless technologies also (cellular, WiMAX, other future technologies), it becomes evident that the issues discussed above can be mitigated and that further flexibility and communications possibilities can be introduced. Indeed, a truly flexible ViewNet system would be one that could use any wireless technology that is available to the system operator.

IV. ULLA AND MULTIPLE WIRELESS

A. The Unified Link Layer API (ULLA)

There have been several attempts in the past to develop generic solutions addressing the challenges arising from the heterogeneity of link technologies (see [9] for an overview). However, these approaches are quite limited when it comes to technology or platform independence. They are also not

scalable to support devices with different capabilities and not extensible to cater for future technologies.

The Unified Link Layer API (ULLA) [5], provides a simple and uniform way to access link layer information independently of the targeted technologies, whilst addressing important requirements such as platform independence, scalability and extensibility. The features of ULLA allow it to fit in well with some of the multi-wireless requirements of ViewNet, providing a mechanism for any ViewNet software or agents to compare heterogeneous wireless technologies, gather information about them and even to set their parameters and control them.

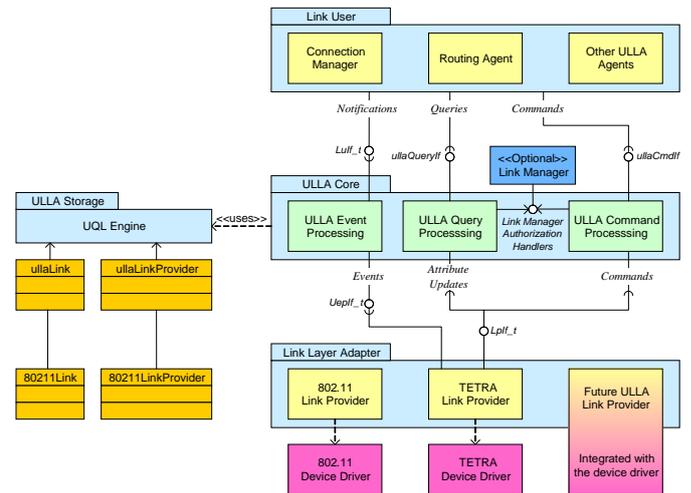


Fig. 2. The Unified Link Layer API (ULLA)

As shown in Fig. 2, ULLA defines the concepts of Link User (LU, any application or other entity which makes use of wireless link information) and Link Provider (LP, an abstraction of a wireless communications device, usually the driver, that provides links to be used by Link Users). The ULLA Core may be used by multiple Link Users at the same time. It includes Command Processing, Event Processing and Query Processing blocks. A database back-end may be used to store link related information according to the ULLA schema, also referred to as ULLA Storage. An optional Link Manager (LM) block is responsible for handling potential conflicts among multiple LUs according to pre-defined policies. A suitable interface is defined to allow the insertion of third party Link Managers.

Link Layer Adapters (LLAs) are software modules that are loaded by the ULLA Core in a platform-dependent way. They are responsible for translating ULLA commands into driver specific methods as well as exporting driver specific events towards the ULLA Core. LLAs also fill the link and link provider tables in the ULLA Storage by properly manipulating the proprietary statistics exported by the driver. A full explanation of the ULLA architecture can be found in [5].

B. Benefits of ULLA for ViewNet

There are a number of benefits associated with using ULLA specifically within the ViewNet project. It is important to

note that ULLA itself is ‘dumb’ in the sense that it won’t provide any means for making decisions regarding communications of the type that have been described in this paper, e.g. determine which to use between WLAN and TETRA. It does however provide:

- **A well defined framework and API** for retrieving information on a range of wireless link technologies. This single API can be used for gathering information on, and controlling, all the relevant wireless technologies that may be used with ViewNet.

- **A way of comparing heterogeneous wireless technologies** for the purposes of link selection and optimisation. The *ullaLink* base class that all Links in ULLA must conform to ensure that basic link characteristics such as bitrate, signal strength, mutual exclusivity and cost are available for all kinds of wireless links and this class can be used as a basis for comparison.

- **Notifications of important wireless events** – ULLA provide a notification mechanism, where interested LUs can ask to be notified when certain conditions are met or triggers activated.

- **Extensibility** – Incorporating a new wireless technology into ULLA is just a matter of implementing the appropriate LLA for the technology in question. Since the ULLA Core sits in between any Link Users and the Link Providers that are registered, no changes would have to be made to any Link User code in order to accommodate the addition of new wireless technologies into a system.

V. IMPLEMENTATION DETAILS

A. WLAN LLA for SoftMAC Wireless Devices

Most of the implementation work carried out in ViewNet has been on the Linux platform. Within Linux, SoftMAC devices allow for a finer control of the hardware by implementing 802.11 frame management functions in software for both parsing and generation of 802.11 wireless frames [10]. Most 802.11 devices today tend to be of this type. The *mac80211* module implements the *cfg80211* callbacks for SoftMAC devices, *mac80211* then depends on *cfg80211* for both registration to the networking subsystem and for configuration of the hardware devices. Although Linux Wireless Extensions (WE) can be used for managing a SoftMAC device, the set of functions that are currently supported by these devices using WE are limited. A better approach is to utilise the *nl80211* module, which together with *cfg80211*, replaces WE to interact with a SoftMAC device. Therefore, a LLA for SoftMAC devices (*sLLA*) has been developed for enabling the use of 802.11 WLAN devices in the ULLA framework.

The radiotap header format [11] provides a means to collect additional information about received frames. It enables the collection of link level information from the driver for user-space applications, and lets user-space applications pass extra information to the driver for transmission. *sLLA* takes advantage of this information embedded in the radiotap header and registers the relevant link level statistics with the ULLA storage using ULLA API calls. The user applications (LU)

accesses the link information using LU APIs and makes efficient use of wireless technologies exposed to the ULLA Core by the LLA modules.

A simple sequence diagram for *sLLA* is given in Fig. 3. During the initialisation, *sLLA* registers itself with the ULLA Core and creates a monitor interface through which it forwards all incoming frames to the monitor module. The monitor module extracts the link level statistics from the radiotap header of each received frame. The link level statistics are then averaged over the update interval and passed to the adapter module which is responsible for registering the link provider and managing links belonging to this class with the ULLA Core as introduced in [5].

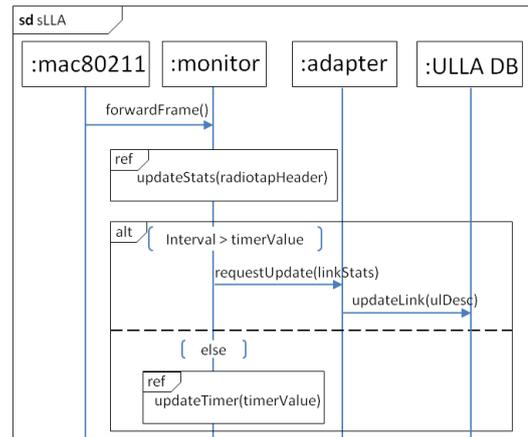


Fig. 3. *sLLA* sequence diagram

B. TETRA LLA

The incorporation of TETRA into ULLA requires the development of an appropriate Link Layer Adapter which can interface with the TETRA hardware that is being used within ViewNet. Typical TETRA handsets include AT command functionality and we have identified this as the basis for the TETRA LLA (*tLLA*) implementation. The AT command interface provides a universal end-point that is manufacturer-independent and easy to use. This enables the development of a wrapper software module that uses the AT commands for retrieving the link layer information from the TETRA radio.

TETRA radios can typically be controlled from a host terminal using a null modem cable connected to the Peripheral Equipment Interface (PEI) of the handset. Unlike *sLLA*, TETRA radios need to be queried for link level information regularly over the PEI. The query commands are sent to the radio as AT commands, responses to which are received as ASCII encoded text at the host terminal, generally in a solicited fashion. However, some unsolicited responses are generated by the TETRA terminal when an asynchronous notification is required. For example, it is important to notify the host terminal when a disconnect event is received by the TETRA handset.

All of the notification responses are parsed by the monitor module using a threaded routine that waits for the AT command events from the PEI of the TETRA terminal as shown in Fig. 4. When a query is made using the

queryAttribute function from the adapter module which implements functions to query and retrieve link attributes, the response is registered by the monitor routine through the private setAttribute function. When the link attribute is set, the monitor module sends an event to notify the adapter module for a successful attribute update. Then, the adapter module can access the updated attribute using the getAttribute function. The unsolicited notifications are handled using a parser module. If the unsolicited notification is successfully decoded, the respective attribute is updated by the AT class and an event is sent to the adapter module for the updated attribute. When all of the queried link attributes are available at the adapter module, they are registered with the ULLA Core via the ulDesc structure using the LP API.

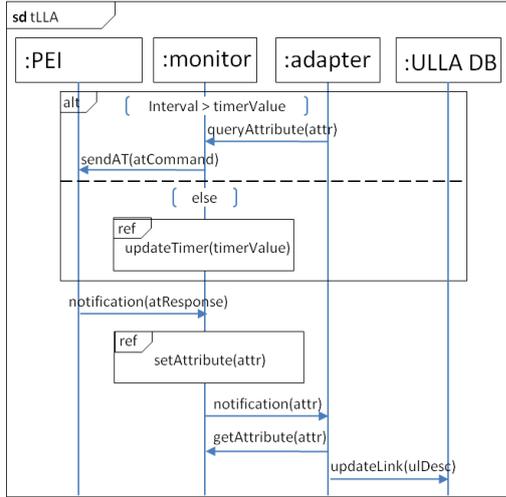


Fig. 4. tLLA sequence diagram

The TETRA and WLAN LLAs enable information gathering and control of both network interfaces using the ULLA APIs. These form the basis for the communications decision making part of ViewNet, which is implemented by the link user module.

C. Radio Access Technology (RAT) Switch

Any application or software agent making use of ULLA for discovering connectivity opportunities and for managing, creating, tearing-down links is referred to as a Link User or LU [5]. The Link Users can monitor a set of link metrics by querying ULLA objects through a subset of Structured Query Language (SQL), called the ULLA Query Language (UQL). This enables the link user application to use a universal API to control any underlying link technology.

The ViewNet LU module is responsible for forwarding outgoing frames to the selected wireless interface; it is hence called the Radio Access Technology (RAT) switch. The selection of the interface is based on the link layer statistics stored in the ULLA database and priorities assigned to the outgoing frames by the ViewNet application. For example, when the 802.11 link signal to noise ratio (SNR) is below a certain link quality threshold, the high priority frames are forwarded to the TETRA interface to improve the overall transmission reliability.

Fig. 5 summarises the forwarding of a single frame received from the ViewNet application at the control centre. The ViewNet control centre application regularly queries the operative equipment for observation updates using HTTP requests. Upon receiving a frame, the LU module queries the ULLA storage for the available links that are suitable for transmitting the query frame. If at least a suitable link is found, the frame is forwarded to the operative equipment over this link. Otherwise, the query frame is dropped and the ViewNet application is informed about the link failure. When there are multiple suitable links for forwarding ViewNet frames, the link selection mechanism takes the link selection policy into account where the links are categorised according to user preference. For example, in an emergency scenario, the emergency services may want to use TETRA as the default link due its increased security and privacy as opposed to an 802.11 link. The link selection policy can easily be configured by modifying a text based configuration file that contains wireless link priorities as integer values.

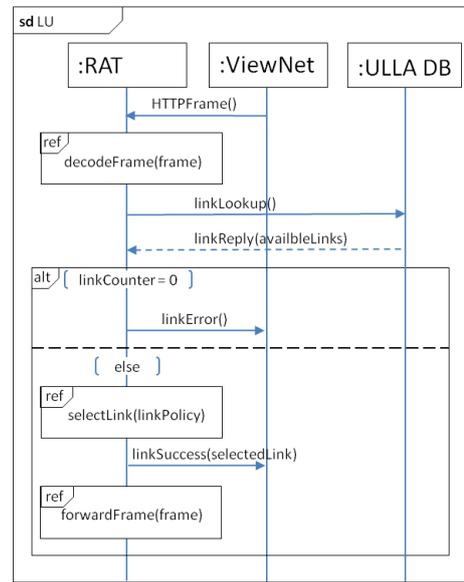


Fig. 5. LU sequence diagram

Selection of the best link for forwarding a ViewNet frame depends on both link level statistics and application preferences. While link level information can be retrieved from ULLA using LU APIs, the application preferences should also be taken into account by the link selection mechanism. This is achieved by inserting a header called the Observation Header in front of each frame handed over to the LU module.

Each item of ViewNet operative equipment registers itself with the control centre when turned on. This operation requires the operative equipment to register IP addresses of its interfaces to the control centre so that the frames can be forwarded to the correct interface using this information. The registration operation is done using the information embedded in the Control Header. Table-I summarises the structures of both the Observation Header and Control Header.

OBSERVATION HEADER	
Field	Description
Priority	Indicates priority of the data contained in the frame. Takes values from 1 to 10, 1 being the highest priority.
isEssential	If true, the observation must be sent, Otherwise, the observation may be discarded depending on the link condition.
isHTTP	This indicates that the frame should be sent using a HTTP POST request. Otherwise, the port, protocol and destination fields are used for forwarding the frame.
Port	Indicates the destination port that the frames should be sent to.
Protocol	Takes values 1, 2, where 1=UDP, 2=TCP.
Destination	This unique ID is given to each operative terminal during the registration process.
CONTROL HEADER	
Field	Description
Destination	Defines the destination ID as defined in observation header.
ifIPList	Structure that contains IP and link technology information of the operative equipment interfaces.

TABLE-I OBSERVATION AND CONTROL HEADERS

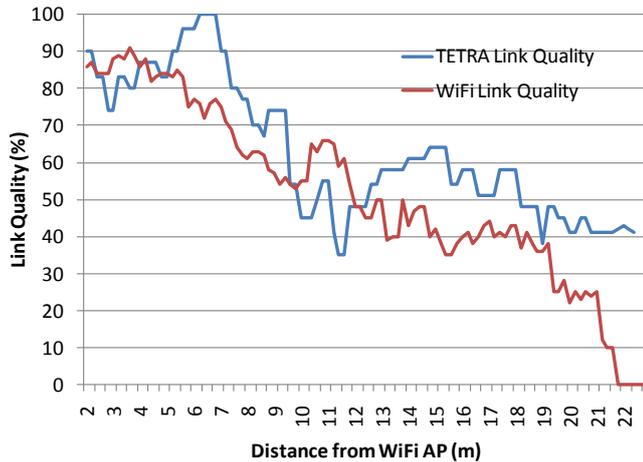


Fig. 6. RAT switch example

Fig. 6 shows an example of how the RAT switch that we have developed works in practice. In the ViewNet demonstrator, a WiFi access point (AP) and TETRA basestation are (almost) co-located in close proximity to the indoor demonstration area – the WiFi AP is on a wall in the demonstration area and the TETRA basestation antenna is located on the roof of the same building. This closely mimics the kind of setup that would be observed in a real ViewNet deployment scenario, i.e. an indoor WiFi AP and an outdoor TETRA cell that the operatives can also connect to. Fig. 6 shows the link quality observed by an operative terminal while moving away from the WiFi AP in the demonstration area. In the first region, the WiFi link provides adequate coverage and signal strength, hence it is used as the default link by the RAT switch, with all packets being routed over this link. In the second region, the WiFi link is only used for

the transmission of low priority packets and the high priority position information is forwarded to the TETRA terminal. The WiFi link fails in the third region, so only the high priority position data is forwarded to the control centre and the low priority data is buffered at the operative terminal.

Although currently every operative in the ViewNet system is equipped with both a TETRA and an 802.11 interface, it is not a strict requirement for any operative to be equipped with either of these interfaces. The ViewNet system can register an operative with a totally different communication interface as long as the control centre supports this interface with an appropriately designed LLA. This creates a reliable and flexible system which requires no further modifications to the existing LU module. The operation of the RAT switch as detailed above can in the future be duplicated for any underlying RAT with little modification. The use of ULLA as the underlying mechanism for obtaining heterogeneous link information is what enables this.

VI. CONCLUSION

In this paper we have discussed the application of multiple-wireless technology to a practical context-enhanced service system called ViewNet. We have proposed the use of ULLA, and described in detail our implementation of ULLA interfaces for both TETRA and WLAN as well as intelligent switching functions between the different radio access technologies. We have tested the system and results show that seamless connectivity can be achieved with the help of the link APIs and RAT switch. This is particularly useful for providing reliable communications with different wireless technologies for mobile ViewNet operatives in a challenging environment.

ACKNOWLEDGMENTS

This work was partly funded by the UK TSB project ViewNet, 3C Research, www.3cresearch.co.uk. We would like to thank the other ViewNet partners for their useful input.

REFERENCES

- [1] P. Bahl and V. Padmanabhan, "RADAR: an in-building RF based user location and tracking system", *IEEE Infocom*, 2000, pp. 775-784.
- [2] S. Helal et al., "Enabling location-aware pervasive computing applications for the elderly", *IEEE Percom*, 2003, pp. 531-536.
- [3] A. Calway et al., *ViewNet project plan*, University of Bristol, 2008.
- [4] A. Gee et al., "Discovering higher level structure in visual SLAM", *IEEE Trans. Robotics*, October 2008, pp. 980-990.
- [5] M. Sooriyabandara et al., "Unified Link Layer API: A Generic and Open API to Manage Wireless Media Access", *Elsevier Computer Communications Journal*, vol. 31, no. 5, pp. 962-979, March 2008.
- [6] M. C. Lucas-Estan, J. Gozalvez, "Real-time computational performance of advanced JRRM policies in B3G heterogeneous wireless systems", *2010 European Wireless Conference (EW)*, pp. 564-571, April 2010.
- [7] <http://www.umatoday.com/>, accessed June 2010.
- [8] *Terrestrial Trunked Radio (TETRA); Voice plus Data (V+D); Part 1: General Network Design*, ETSI Standard, EN300 392-1, 2009.
- [9] T. Farnham et al., "Toward Open and Unified Link-Layer API", in *Proceedings of the IST Mobile and Wireless Summit 2005*, Dresden, Germany, June 2005.
- [10] Linux wireless networking, <http://linuxwireless.org/>, accessed February 2010.
- [11] Radiotap standard, <http://www.radiotap.org/>, accessed February 2010.