Y-Comm: A Global Architecture for Heterogeneous Networking

Jon Crowcroft, David Cottingham
Computer Laboratory
University of Cambridge
Cambridge CB3 0FD
Email: jon.crowcroft, david.cottingham@cl.cam.ac.uk

Glenford E. Mapp, Fatema Shaikh
Networking Research Group
Middlesex University
London, NW4 4BT
Email: g.mapp, f.shaikh@mdx.ac.uk

July 25, 2007

Abstract

In the near future mobile devices with several interfaces will become commonplace. Most of the peripheral networks using the Internet will therefore employ wireless technology. To provide support for these devices, this paper proposes a new framework which encompasses the functions of both peripheral and core networks. The new architecture is called Y-Comm and is defined in a layered manner like the OSI model.

Keywords: Heterogeneous Networking, Mobility, Architectural Framework.

1 Introduction

Developments in various areas of computer communications are progressing at a tremendous rate. Mobile computing devices such as PDAs, cell phones and ultra-thin laptops are becoming essential for human social environments. They are also becoming more versatile enabling us to communicate, play music, do work and play games while on the move. Key applications now being used are phoning, emailing, texting and instant messaging. However, consumers now want even more demanding applications including high quality video streaming, interactive video conferencing and real-time interactive video games. The enthusiasm connected with the recent release of the iPhone confirms this trend. To deliver such services will require new networking infrastructure.

In addition, these devices will have several wireless interfaces. This is clearly seen in the development and deployment of several wireless networks such as 3G, WLAN, WiMax and Ultrawideband. New and faster wireless networks such as 802.11n will also allow more demanding applications to be supported on mobile devices. We define the networking issues associated with this type of device as **heterogeneous networking** and the devices themselves as **hetnet** devices. To facilitate seamless operation of these networks as users move around, will require the development of new mechanisms to support vertical handover as well as policy management systems that decide when handovers should occur. Another major requirement will be the development of Quality-of-Service (QoS) techniques to ensure that applications can still provide the necessary functionality even though the characteristics of the underlying network channels are continuously changing. Finally, the expansion of the number of wireless networks at the physical and link levels is contrasted with the convergence on the use of the Internet Protocol (IP) at the network level to build global multi-service networks. So a fan-out at the lower level is being met by a fan-in at the network level, making the integration of these wireless systems a serious challenge.

The last few years have also seen significant network evolution with regard to the Internet. When the Internet was young, the peripheral networks were primarily Ethernet and Token Bus systems which were similar in terms of performance and technology to the systems used in the core network. However recently there has been a radical divergence. The core network is actually getting faster in terms bandwidth as well as latency with the use of single-mode optics and Multiple Label Switching (MPLS) technologies. In contrast, new peripheral networks are being predominantly built using wireless technologies including WLAN, Bluetooth and WiMax systems. The characteristics of these systems are totally different in terms of latency and bandwidth as well as error distribution properties compared with those in the evolving core network. This weakens the end-to-end arguments which has been a key part of the design framework for the early Internet [SRC84]. We believe that this means that we need to think of the Internet as a global network which should be divided into two key components. The first is the Peripheral Network and the second would be the Core Network which also includes access networks. This division highlights the fact that the challenges in the Peripheral Network will be different from those in the Core Network.

The observations above all indicate that there is a need to develop a new framework for global communications. The challenges that must be met are beyond the current frameworks such as the OSI model [Zim88]. This paper looks at the development of a new global architecture for heterogeneous networks. The architecture, known as the Y-Comm Architecture comprises two frameworks: the Peripheral Framework for peripheral networks and the Core Framework for the Core Network. The rest of this paper is organised as follows: Section 2 looks at the OSI model and outlines key architectural concerns leading to the need for the new Y-Comm design. Section 3 looks at the Peripheral Framework while Section 4 describes the required functionality needed in the core network to support the Peripheral Framework. Section 5 explores the layers of the Core Network and Section 6 combines both frameworks to form the Y-Comm architecture. Section 7 looks at previous and current work being done in this area of research while the paper finishes with a section on conclusions and future work in Section 8.

2 The Architectural Need for Change

2.1 The OSI model

The OSI model has been the dominant reference model in computer communications over the last two decades. While this model has helped in the development and deployment of mechanisms that are widely used in the Internet today, the authors believe that the OSI is inadequate for modelling heterogeneous networking. This is mainly because in the OSI model, the network is essentially being used to forward packets. This is specified in the first three layers of the architecture. However, in heterogeneous networking, more functionality needs to be supported which require intimate communication between the mobile node and the peripheral network. An example of this is the set of mechanisms required to support vertical handover which is a key requirement in the deployment of 4G networks [MZ04]. These mechanisms will include the ability to reconfigure network parameters to allow the reservation of resources to ensure that the required quality-of-service is maintained after vertical handovers.

2.2 The Layered Approach

Though the OSI model may be superseded, the layering model adopted in the OSI model is a good approach for defining the new architecture. This approach specifies the hierarchy of functionalities but leaves the detailed interfaces between the layers to be optimized by the implementors of the architecture. In addition, whereas the model details the ordering of the necessary functionality, it does not discount the use of cross-layer implementations. For efficiency, there is also scope for implementors to integrate one or two layers while keeping the general ordering of the layers intact. This effort therefore stresses the conceptual use of layering as an effective way of specifying network architectures while allowing implementors the flexibility to pursue efficient networking platforms.

2.3 Asymmetric vs Symmetric Architecture for End and Intermediate Systems

Y-Comm proposes a break with tradition in terms of the classical Internet model, which was designed to use the same protocols in end systems or hosts and in intermediate systems or routers. Most network architectures (OSI, PTSN, B-ISDN, DECNET, SNA) made this distinction, although some systems, namely Novell Netware, XNS and TCP/IP, attempted to unify the model of nodes that were terminals and nodes that were switches. This is in recognition of the computer-centric view of data networking where networking infrastructure devices, such as routers, were general purpose computers just as much as end-user systems (clients, servers, handsets).

Mobile computing is predicated on wireless access and introduced new protocol elements in the wireless hop to reach the fixed Internet. Some of these elements or middle boxes, had already started to appear but these were due to other reasons such as security through firewalls, address space management through DHCP and NATs, etc. These have combined with the additional support needed for wireless networking including roaming support, access control and split TCP mechanisms to produce services which are now indispensible to the modern Internet.

In addition, a key operational goal of the early TCP/IP implementations was that any device with multiple interfaces could act as a router, even if only to choose the first hop. In practice, the evolution of the Internet to a more commercially-oriented environment has meant that today, few network providers allow an end system to participate in the routing exchanges to influence the way packets are delivered other than by simply enabling and disabling interfaces. Wireless networks are under even more pressure to make the separation since there are different types of wireless networks being deployed – from essentially free WiFi to expensive, operator-dominated cellular wide area services. These systems have very different characteristics including service provisioning, AAA (Authorisation, Access Control and Accounting) concerns and QoS parameters. Presently, there exists no commonly-agreed framework to allow these networks to operate together seamlessly – even with input from the user of the mobile node. Thus, at time of writing, it is common to find that handset operating systems have removed the functionally of choosing which network to use. We believe this will eventually evolve into a situation where choices will be allowed based on economic considerations, but much more is needed to provide a seamless environment.

Future developments at the physical layer of het-

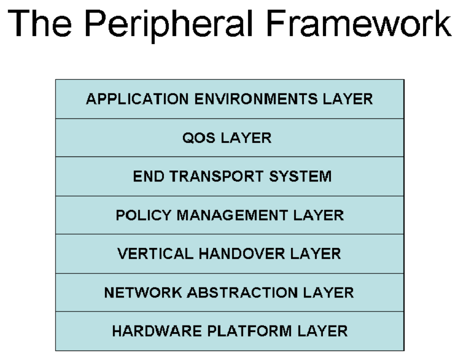

Figure 1: The Peripheral Framework

net devices, such as cognitive radio [ZS07], will necessitate further client-side control of what network and technology is best to use at any one time or location. Indeed, cognitive radios are likely to employ collaborative sensing in order to ascertain what spectrum may be utilised, and hence a network provider will not be the sole source of such information. Hetnet devices will connect to both centrally managed networks and to more ad hoc or peer-to-peer ones, requiring far more flexibility than is currently available to users' devices.

For all the reasons outlined above, we think that it is necessary to consider a new Internet architecture that makes a split between the edge and the core, and between the user and services. It should also be pointed out that many services, while not integrated into an ISP's business are in fact network server-centric (e.g. data centers, such as search engines, transaction services, community web sites and compute services as offered by Amazon and others). This represents a high-level split which should be included in a new architecture – though not mandated but modeled as a common choice. The Y-Comm architecture outlined below addresses these issues by proposing the specification of two frameworks: The Peripheral Framework aimed at the mobile user and the Core Framework which looks at a framework for networking infrastructure and services in the core of the network.

3 The Peripheral Framework

Figure 1 shows the Peripheral Framework developed for heterogeneous networking. A more detailed explanation of the architecture is found in [MCS+06].

- Layer 1: **The hardware platform layer**: this layer is used to define the hardware components and technologies required to support a particular wireless network, including electromagnetic spectrum, modulation techniques, Media Access Control (MAC) algorithms, etc. We can therefore represent each different networking technology as a vertical slice of the hardware platform layer. A key function of this layer is to determine which technologies are compatible and hence can be operated simultaneously.

- Layer 2: **The network abstraction layer**: this layer specifies a common networking interface which all networks employing this architecture must support. This interface is used to maintain and control the network on the mobile node. Different wireless device drivers must be written to map onto this layer. A major feature of this layer is to determine which networks are available and which networks will shortly be out-of-range of the mobile node.

- Layer 3: **The vertical handover layer**: this layer is concerned with the specification of mechanisms including state engines and triggers for vertical handover. There are two kinds of vertical handovers. The first is network-controlled and is managed and maintained in the core network. The second is client-controlled in which the client controls handover. All commercial systems that support vertical handover, such as BT Fusion[1], use network-controlled handover. The authors however believe that client-based handover is a more elegant solution for a number of reasons. Firstly, the network abstraction layer running on the mobile nodes allows the mobile node to have up-to-date information about its interfaces and the state of wireless networks on the mobile node hence allowing a better environment to support vertical handover. In addition, the client can take into consideration other factors such as the state of its TCP connections in its decision about when to execute handover. The client-based solution is also more scalable as it can easily access information on several networks. Network based handover will involve an entity in the network knowing about the status of all the networks which would require network operators being willing to hand over sensitive information to third parties.

- Layer 4: **The policy management layer** : this layer is used to evaluate all the circumstances when handover should occur. The layer can be implemented by defining certain rules with regard to all the relevant parameters and their values which are evaluated with respect to handover. Policies can be essentially divided into two categories: reactive and proactive policies. Reactive policies are triggered by changes in the condition of the networks to which the mobile node is connected. Such triggers are conveyed by the network abstraction layer. Pro-active policies attempt to know the condition of the various networks at a specific location before the mobile node reaches that location. Pro-active policies allow mobiles node to calculate the Time Before Vertical Handover (TBVH) which will allow the mobile node to minimize the effects of vertical handover.

- Layer 5: **The End Transport System**: this layer looks at moving data to and from the mobile node. Since most peripheral networks will be wireless, it is therefore important to ensure that network and transport systems operate efficiently so that applications running on the mobile node can receive sustainable qualities-of-service. TCP/IP which is used throughout the Internet has been shown not to perform well in wireless networks [Mey99] [XPPS01]. In addition, recent work has shown that TCP adapts very slowly to network conditions after vertical handover [CV05]. Several attempts have been made to modify TCP as well as efforts to make TCP more responsive to temporary outages without changing the protocol engine [SM03]. However, no solution has been generally adopted. Alternatives to TCP/IP, such as SCTP [SXM$^+$00], are being developed to obtain better performance when transporting data in heterogeneous wireless networks.

- Layer 6: **The Quality-of-Service (QoS) Layer**: this layer helps to ensure that quality-of-service required by applications can be maintained as the quality-of-service being offered by the networks is dynamically changing as the mobile node moves around. This framework defines two types of QoS support. The first is called Downward QoS which allows applications to specify the QoS they require and

[1] `http://www.bt.com/btfusion/`

leaves the system to support such requirements over available network channels. The second type is called Upward QoS in which applications themselves attempt to adapt to changes in network conditions. Current applications which cannot adapt to changing conditions will employ Downward QoS while newer applications such as multimedia and networked games would need to use Upward QoS.

- Layer 7: **The Application Environments Layer** specifies mechanisms and routines that allow applications to be built which can use all the layers of the framework. There have been a few attempts to build complete application layers for mobile environments, such as Ambient Networks [NSH$^+$04]. An interesting approach is adopt a toolkit approach which allows different types of application environments to be built. This approach is similar to that used in the deployment of X Window System [SG86]. Hence various objects in the toolkit can be used build a particular application environment. So for example, using the toolkit, we could specify a context-aware, location-aware, application layer for high mobility environments using Upward QoS mechanisms. Such an application environment will use the lower layers of the Peripheral Framework via the objects in the toolkit to build relevant applications.

4 Implications for the Core Network based on the Peripheral Framework

Some key observations about the Peripheral Framework: firstly, the Peripheral Framework is designed to be implemented on hetnet devices. However, in order for client-based handovers to work, the mobile node would need access to network resources in order to facilitate vertical handover. This must therefore be a key requirement of the Core Framework.

Secondly the Policy Management System is meant to support proactive policies. This can be done using a knowledge-based system in which the parameters at that location have been previously measured. The other approach is the use of a simple mathematical model based on the distance of the mobile node from nearby based stations as well as the direction and velocity of the mobile device. In order to calculate Time Before Vertical Handover (TBVH), the policy management layer would need to know something of the network topology and the performance of the relevant base-stations. A major new nomenclature is the concept of a Boundary Base Station (BBS) which is a base station on the boundary of the network. A BBS is the last base station that must be traversed just before vertical handover. So it is important for different networks to be able to describe their network topology in a way that facilitates the calculation of TBVH.

Finally, though there will be a weakening of the end-to-end transport mechanisms, successful communication does demand end-to-end QoS support. It will therefore be necessary to map the QoS available in the peripheral networks to the QoS in the core network and vice-versa.

4.1 Networking Issues for the Core Network

There are additional issues which need to be considered. It would be beneficial if the Core Framework attempted to address key issues in the management of large telecommunication systems. Firstly, network operators presently dominate telecommunication systems resulting in a highly vertically integrated architecture. The new framework should attempt to define an architecture that allows a more horizontal approach. In this regard, it is felt that a tightly layered approach in which the functionality of the layers are clearly defined will allow entrepreneurs to specialize in providing particular services resulting in a more horizontal business model.

Secondly, there is also a requirement to be able to define and manage non-overlapping networks on a single hardware platform. Such a design would allow the development of city-wide or regional wireless networks which can be better tailored for relevant users. This would also allow the deployment of new technologies in a limited geographical environment making a more viable business by the gradual deployment of services. Presently, large scale national networks such as 3G must be deployed at a national level. This requires a lot of expense and thus can only be done by companies with deep pockets. A major feature of the framework is the ability to support network virtualization and partitioning which can be used to define a virtual network which is managed by a network operator and can be viewed by a subset of an extensive hardware

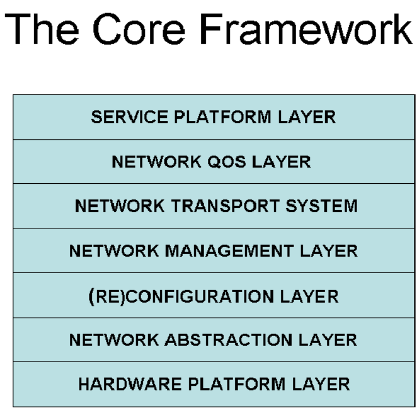

Figure 2: The Core Framework

platform.

Finally, we believe that in order to enhance the new architecture, it is necessary to provide support global service platforms. These platforms would allow services to be implemented and managed independently. Currently, this is not possible as the deployment of a given service must be done sequentially on each individual network.

5 The Core Framework

The Core Framework is shown in Figure 2. The first two levels of the Core Framework are similar in purpose to the first two layers of the Peripheral Framework, but while the Peripheral Framework specifies software such as device-drivers in order to support a given network on a mobile node, in the Core Framework these layers represent the specification and software needed to run in the base station of a given technology. The relationship of these layers in the two frameworks is analogous to the specification of Data Communication Equipment (DCE) such as a modem and Data Terminal Equipment (DTE) definitions found in wired data communications. However in this case we look at specification from a wireless and not a wired context. So the base-station specification corresponds to the DCE while the mobile node is analogous to the DTE end of the interface.

The **(Re)configurable Network Layer**: this layer provides a control plane for (re)configuring networking resources in the core network. This would include various network switching elements such as mobile switching centres, gateway GPRS support nodes and routers [HSBH03]. This interface will also be used by the Vertical Handover layer in the Peripheral Framework to obtain network resources for a vertical handover before it occurs. Network events as well as the configuration of new resources to satisfy the QoS requirements may generate new reconfiguration needs in order to guarantee the stability of the whole system. Reconfigurable systems benefit greatly from the virtualization of hardware components such that it is possible to have a small number of virtual units, for example, switchlets [dML97] or routelets. Though it is clear that a lot of research has been done in this area and some of it has found its way into commercial products, what is missing is the opening up of these interfaces [BLH+98] to hetnet devices. The drive to open up these systems has not gone far enough and without this it will not be possible to build networks that are different in scope and functionality using the same hardware. This ability is necessary for the next stage in network evolution [Laz97]. Of course such an effort must be accompanied by the required security framework to prevent hetnet devices attempting to abuse of core networking resources.

The **Network Management Layer**: This layer is highly significant as it acts as a management plane that uses the programmable network layer to bring together various hardware and software components to build enterprise class networks. Each network will have an operator that controls it. To do this, the layer must also provide Authentication, Access Control, Accounting and Charging (AAAC) systems [RHKS01]. It must also support the use of policy mechanisms that would allow operators to dictate which hardware components may be used on their networks. The Policy Management Layer in the Peripheral Framework can interact with the Network Management Layer in the Core Framework to help inform mobile nodes about network resources to which it could have access on specific networks. The Policy Management layer uses this information to tell the Vertical Handover layer on the mobile device about which network resources can be obtained for a vertical handover. Since both the Network Management Layer and the Policy Management layer have their own policies to follow, a conflict resolution process should be carried out between them.

The **Network Transport System**: This layer is about network addressing and transport mechanisms in the core network. Currently TCP/IP is used in the core network and we are of the opinion that it should continue to be used, though a move to IPv6 is necessary to add enhanced network capabilities and integrate the various value-added technologies into one core protocol.

The **Network QoS Layer**: This layer is responsible for QoS issues within the core network. It looks at how QoS may be defined and the mechanisms used to establish and maintain QoS at different points in the system [DS03]. With the failure of IntServ [BCS94] and the slow deployment of DiffServ [Gro02], a new model for handling QoS issues is required. A lot of motivation for the development of IntServ and DiffServ was the belief that the Internet would soon be unable to deal with the huge increase in traffic that would be spawned from its high growth rate [HK00]. However what has happened is that the core network has become faster, minimising the threat of congestion in the core network. In addition, the development of heterogeneous wireless networks means that there are more severe QoS issues in the peripheral network than in the core network. We therefore believe that QoS issues in the core network should be approached from the network level rather than from the application or device level. Hence, a novel approach is to develop a QoS architecture based on the ability of peripheral networks, rather than individual machines, to calculate and specify their QoS requirements. These networks will then negotiate with the core network to obtain the required resources to meet their QoS needs.

The **Service Platform**: The Service Platform allows different agents to install and operate various services in a secure and controlled fashion. The service platform will provide the ability to install services in component form on several networks simultaneously, or on a single network. This will therefore allow the provision of both national and regional services to be easily constructed, e.g. traffic information in London would be a local service accessible to networks operating in London. There is enormous scope for such location-based services.

6 The Y-Comm Framework

In this section, we attempt to put the Peripheral Framework and the Core Framework together to represent a future telecommunications environment which supports heterogeneous devices, disparate networking technologies, network operators and service providers. This is shown in Figure 3.

The two frameworks share a common base sub-

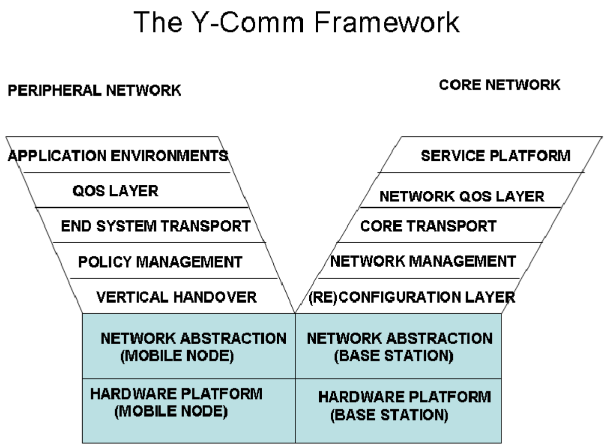

Figure 3: The Y-Comm Framework

system consisting of the hardware platform and network abstraction layers. Both frameworks diverge in terms of functionality but the corresponding layers interact to provide support for heterogeneous environments.

6.1 The Security Model for the Y-Comm Architecture

It should be noted that security is not a layer in the Y-Comm Architecture. This is because security is needed on a number of levels in the architecture. For the Peripheral and Core subsystems, we describe three levels of security which need to be present.

The first of these levels is called **Network Architecture Security** or NAS. This concurs with the security characteristics of a given network technology and the security threats that ensue from the deployment of such a technology. A good example is the fact that wireless systems are inherently broadcast (over the air) so that security measures must be taken with respect to the transmission of packets compared to a wired systems. In Y-Comm, NAS algorithms must be implemented in the Policy Management System within the Peripheral Framework and in the Network Management Layer of the Core Network Subsystem. It is necessary in the Policy Management Layer because this layer decides when and to which network a vertical handover should be directed, hence it must be aware of the security risks posed by each network. The Network Management Layer must also implement security based the type of technology being used in the Core network to implement reconfigurable systems, e.g., MPLS or

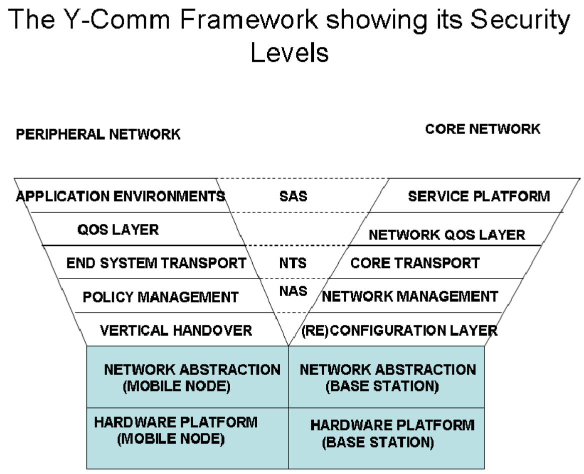

Figure 4: The Y-Comm Security Model

ATM as well as the algorithms to obtain access to programmable hardware.

The second level security is applied at the Transport Level and is concerned with the reachability of, access to and communication with networks and hetnet devices. This is defined as **Network Transport Security** or NTS. In the Peripheral Framework, NTS takes the form of NAT, firewalls, intrusion detection systems (IDS), etc. while in the Core Framework, NTS takes the form of algorithms that ensure secure transport through the core network such as IPSec [DH03].

The third and final security level occurs at the highest layer and is called the **Service and Application Security** or SAS. For the Peripheral Framework, SAS determines what access is given to applications to use resources on host machines or LAN services. These will involve the use of authentication algorithms. For the Core network, SAS defines which resources a service provider will be allowed to access or control and on which networks the service provider can install various services [39]. The authors believe that in the Y-Comm architecture it is necessary to have these three levels specified for each subsystem to obtain complete security. The unified architecture with its security levels is shown in Figure 4.

7 Previous and Current Work

With regard to the Peripheral Framework, much recent research in mobile networks has looked at vertical handovers. This was also explored by the development of the Cambridge Wireless Testbed [CVS+04] which was built by the Computer Laboratory, University of Cambridge. The testbed was unique as it explored vertical handover in LAN, WLAN and 3G networks. It pioneered the use of client-based handover techniques [PM03]. In addition, the Cambridge testbed was used to look at reactive mechanisms for policy management called PROTON [VCP04]. This was implemented as a three-layer subsystem. The lowest level was the hardware execution layer, which performed the actual handover. The second layer was the policy layer which allowed policies to be specified as rules which were used to decide whether handover should be initiated. The final layer was an input/output layer which fed events and triggers into the policy layer.

The efforts detailed above concentrated on the layers 3 and 4 of the Peripheral Framework. Work is also looking at defining the lower layers of this framework. Recently, the IEEE convened the the **802.21** [2] Working Group to examine the possibility of standardising the interface to different wireless MACs. In our view, this work can be used as a prototype of the mobile-node side of the network abstraction layer.

Recent work has been looking at developing proactive management policies. At the University of Cambridge algorithms for constructing coverage maps of wireless networks such as WLAN, GPRS and 3G throughout the city are being developed [Cot07]. This will allow hetnet devices to ascertain coverage at a particular location and predict handovers accordingly. This is being carried out as part of our wider work on wireless network provision for vehicles [CD07]. Meanwhile, proactive systems-based mathematical modelling is being pursued by the Networking Research Group at Middlesex University [SLM06]. The aim is to find a simple and efficient way of calculating the Time Before Vertical Handover (TBVH). Analytical models have been developed and are being verified using simulations in OPNET. A prototype model for providing support for QoS is also being developed [SLM05]. Finally, work has also begun to look at Network and Transport protocols for the End Transport System using the Plutarch Model [CHM+03].

In terms of the Core Framework, work has begun to look at the (Re)configuration Layer. There has been a careful review of the Programmable and Active Network Research that was done in the late

[2] http://www.ieee802.org/21/

1990s [Kou03]. The idea is to use this work as a starting point to define a layer that can allow mobile nodes to acquire the necessary resources to aid client-based handovers. We are also beginning to look at extending the work being done by the IEEE 802.21 Working Group to cover base-station functionality.

Finally, the implementation of the Y-Comm architecture involves the participation of many agents and parties. It is therefore extremely important that there are mechanisms to allow entities to acquire the necessary information and resources to operate effectively. For example, to have an efficient vertical handover will involve network operators, user preferences, location and context information as well as economic considerations. We are therefore looking at the use of ontologies [BYBN07] to provide the required functionality for managing these issues.

8 Conclusions and Future Work

This paper has proposed a new architecture called the Y-Comm architecture which we believe can be used to build future telecommunication networks for heterogeneous networking.

With regard to future work, we would like to look at network management, in particular, network virtualization and partitioning techniques. This would allow us to define and manage new networks based on an extended hardware platform. After this, we would like to look at layers 6 and 7 of the Core Framework.

We recognize there are tremendous challenges in trying to prototype the Y-Comm Framework and therefore appeal to the networking research community to engage seriously with this effort.

References

[BCS94] R Braden, D Clark, and S Shenker. RFC 1633: Integrated Services in the Intranet Architecture: an Overview. In *Technical Report*. IETF, 1994.

[BLH+98] J Biswas, A.A Lazar, J.F Huard, K Lim, Mahjoub S, L.F Pau, M Suzuki, S Torstensson, W Wang, and S Weinstein. The IEEE P1250 Standards Initiative for Programmable Network Interfaces. *IEEE Communications Magazine, Special Issue on Programmable Networks*, 1998.

[BYBN07] I.G Ben Yahia, E Bertin, and Crespi N. Ontology-Based Management Systems for the Next Generation Services:State-of-the-Art. In *Proc. ICNS*, 2007.

[CD07] D Cottingham and J Davies. A Vision for Wireless Access on the Road Network. In *3rd International Workshop on Intelligent Transportation*, March 2007.

[CHM+03] J Crowcroft, S Hand, R Mortier, T Roscoe, and A Warfield. Plutarch: An argument for Network Pluralism. In *Proc. ACM SIGCOMM Workshops*, August 2003.

[Cot07] D Cottingham. Wireless coverage mapping in vehicular contexts. In *Submitted to ACM CoNEXT*, 2007.

[CV05] D Cottingham and P Vidales. Is Latency the Real Enemy in Next Generation Networks. In *Proc. First International Workshop on Convergence of Hetergeneous Wireless Networks*, July 2005.

[CVS+04] R Chakravorty, P Vidales, K Subramanian, I Pratt, and J Crowcroft. Performance issues with vertical handovers - experiences from GPRS cellular and WLAN hot-spots integration. In *Proc. IEEE PerCom*, 2004.

[DH03] N Doraswamy and D Harkins. IPSec: The New Security Standard for the Internet, Intranets and Virtual Private Networks. In *Prentice Hall Security Series*. Prentice Hall, 2003.

[dML97] J.M der Merwe and I.M Leslie. Switchlets and Dynamic Virtual ATM Networks. In *Proc. Integrated Network Management V*, 1997.

[DS03] A Duda and C Sreenan. Challenges for Quality of Service in Next-Generation Mobile Networks. In *Proc. IT&T Annual Conference*, 2003.

[Gro02] D Grossman. RFC 3260: New Terminology and Classification for DiffServ. In *Technical Report*. IETF, 2002.

[HK00] J Harju and P Kivimaki. Co-operation and Comparison of DiffServ and IntServ. In *Proc. 25th Annual IEEE Conference on Local Area Computer Networks*, 2000.

[HSBH03] G Hjlmtsson, H Sverrisson, B Brynjlfsson, and Helgason. Dynamic packet processors - A new abstraction for router extensibility. In *Proc. OPENARCH-2003*, April 2003.

[Kou03] M Kounavis. Programming Network Architectures (PhD Thesis). In *COMET Group*. Columbia University, June 2003.

[Laz97] A.A Lazar. Programming Telecommunications Networks. *IEEE Networks*, 11(5), 1997.

[MCS+06] G Mapp, D Cottingham, F Shaikh, P Vidales, L Patanapongpibul, J Balioisian, and J Crowcroft. An Architectural Framework for Heterogeneous Networking. In *Proc. WINSYS*, pages 5–10, August 2006.

[Mey99] M Meyer. TCP performance over GPRS. In *IEEE WCNC*, pages 1242–1252, 1999.

[MZ04] J McNair and F Zhu. Vertical Handoffs in Fourth-Generation Multinetwork Environments. In *IEEE Wireless Communications*, volume 11, 2004.

[NSH+04] N Niebert, A Schieder, Abramowicz H, Malmgren G, J Sachs, Horn U, Prehofer C, and Karl H. Ambient Networks: An Architecture for Communication Networks Beyond 3G. *IEEE Wireless Communications*, 11(2), 2004.

[PM03] L Patanapongpibul and G Mapp. A Client-Based Handoff Mechanism for Mobile IPv6 Networks. In *Proc. 8th IEEE Symposium on Computers and Communications*, June 2003.

[RHKS01] C Rensing, Hasan, M Karsten, and B Stiller. A survey of AAA mechanisms, protocols and architectures and a policy based approach beyond: Ax. In *Report 111*. Swiss Federal Institute of Technology, 2001.

[SG86] R.W Scheiffler and J Gettys. The X Window System. *ACM Trans. Grap.*, 5(2):79–109, 1986.

[SLM05] F Shaikh, A Lasebae, and G Mapp. SBM layer for optimum management of co-existing telemedicine traffic streams under varying channel conditions in heterogeneous networks. In *Proc. MIT*, 2005.

[SLM06] F Shaikh, A Lasebae, and G Mapp. Client-based SBM layer for predictive management of traffic flows in heterogeneous networks. In *Proc. IEEE ICTTA*, 2006.

[SM03] J Scott and G Mapp. Link layer-based TCP optimisation for disconnecting networks. *Computer Communications Review*, 33(5), 2003.

[SRC84] J. H Saltzer, D.P Reed, and D. D Clark. End-to-end arguments in system design. In *ACM Transactions in Computing Systems*, pages 277–288, 1984.

[SXM+00] R Stewart, Q Xie, K Morneault, C Sharp, H Schwarzbauer, T Taylor, I Rytina, M Kalia, L Zhang, and V Paxson. RFC 2960. In *Stream Control Transmission Protocol*. IETF, October 2000.

[VCP04] P Vidales, R Chakravorty, and C Policroniades. PROTON: A Policy-based Mobile Client System for the Fourth Generation Mobile Environments. In *Proc. IEEE Fifth International Workshop on Policies for Distributed Systems and Networks*, June 2004.

[XPPS01] G Xylomenes, G.C Polyzos, Mahonen P, and Saaranen. TCP performance issues ver wirless links. *IEEE Communications Magazine*, 39(4):52–58, 2001.

[Zim88] H Zimmermann. OSI reference model - the ISO model of architecture for open systems interconnection. *Innovations in Internetworking*, pages 2–9, 1988.

[ZS07] Q Zhao and B Sadler. A survey of dynamic spectrum access. *IEEE Signal Processing Magazine*, pages 79–89, May 2007.