A Position Paper: Towards an Utility Computing and Communications Infrastructure

Peter Komisarczuk, Kris Bubendorfer

School of Mathematics, Statistics and Computer Science, Victoria University of Wellington,
PO Box 600, Wellington, New Zealand, Tel: +64 4 463 5661, Fax: +64 4 463 5045

***Abstract*—The future of the Internet may be a conjunction of Grid computing and service overlays hosted over Next Generation Internet (NGI) technologies, consisting of both wireline and wireless networks. The Grid and overlay networks can be seen as a key service layer for the future where the combination of computing and communication resources is dynamically allocated to virtual organizations on demand to enable optimal service and business deployment. To enable this synergy the efficient allocation of resources at minimum cost needs to be enabled. Today our resource description, resource matching and pricing in Grid/NGI needs to be extended/addressed and our viewpoint on these aspects are introduced in this paper. Initially this paper looks at some of the drivers and business models that could be envisaged with some background on developing grid networks and a brief review of some recent work on QoS description.**

***Keywords*: business model virtual organization NGI**

1 INTRODUCTION

The Next Generation Internet (NGI) could be formed of the product of an underlying advanced communications infrastructure and Grid/utility like computing resources over which are deployed overlay networks. The product of these platforms could effectively form the basis for an extensible and dynamically allocated services network.

The network overlay abstraction provides a flexible, extensible, application level paradigm to deploy services easily and incrementally despite heterogeneous underlying network and service technology. However if we combine overlay networks with utility/Grid computing infrastructure and advanced networks using NGI technology we could experiment in the provision of the key infrastructure for future networked services.

It has been postulated that utility computer service providers [1] could be the basis for cost effective, easily extensible solutions for enterprise and service providers, including network service providers. However today "On demand Computing" is typically bounded in terms of the type and number of resources under control, in fact these are usually limited just to the data centre. Typical these current solutions provide aspects of the following:

- Resource Management
- Load Balancing
- Asset Assignment
- Metered payment for additional resources – resources on demand

Current resources are processor blades, virtual machines or various storage assets (DAS, NAS, and SAN) that are assigned on demand and charged on some metered basis [2]. Within the datacenter the Enterprise viewpoint is that of investment proofing their IT infrastructure, whilst the service provider viewpoint is obtaining maximum revenue for their deployed assets through efficient use of the provisioned IT infrastructure – by maximizing sharing/contention. However we are seeing that tomorrows "on demand computing" concepts are beginning to borrow from "Grid Computing" in order to provide processing power and storage "on tap" across a wider perspective.

Extending utility computing concepts to include resource allocation of the enabling communications infrastructure could enable ubiquitous overlay service provisioning, where the utility computing communications providers [UC^{2}P], leases resources to Virtual Organisations [VO] that are created to provide services to users or the network as a whole. In this paper the VO concept is bigger, more encompassing than the Grid computing concept of a virtual organisation [3], where the VO is created in order for a group of distributed computing resources to collaborate to achieve some task, such as a simulation. The Grid computing resources, data and code may be part of different commercial entities who collaborate to create the VO. In this paper and the UC^{2}P we want to view the VO in a larger business context, where it is "*applicable in social as well as in information systems*" [4]. A unified field theory equivalent for the description of VOs in all its myriad of combinations and facets: computer, communications, operations, data/value (knowledge) and workflows, management, business relationships, right through to human personalities is being explored [4]. Although a unified VO descriptor is one of the ultimate goals in this proposed work the concepts expressed within this paper look at the computer, communications and business aspects of a VO in the context of the potential future overlay network services in the future Internet.

The Grid computing model may be considered a basis for exploring utility computing. In fact the European Union GRASP [5][6] project is exploring the Grid paradigm (based on .NET technology) for developing a commercial Application Service Provider model. Our work has taken a similar course and in particular our work has initially focused in the area of enabling virtual organizations in cellular applications [7].

The UC^{2}P infrastructure envisaged as the future service model within this paper demands the "Next Generation Grid" (NGG), encompassing more resource types and providing greater flexibility in terms of mobility and resource allocation based on an economic model than current Grid implementations. Probably the predominant Grid toolkit today is Globus [8], which has come some way to providing an NGG through the Open Grid Service Architecture (OGSA) and Open Grid Service Infrastructure (OGSI) and enabling web services though the Web Service Resource Framework (WSRF). Work on the Globus 4 is progressing, however this is moving to a more standards based approach on Tomcat and SOAP [9] in common with more commercial utility computing solutions (not discussed further here, but include LSF etc). Within these the current Grid communication services model may not encompass all the potential services that could be required in a general purpose UC^{2}P scenario and resource allocation does not provide for the speediest allocation mechanism nor does it currently provide an optimal charging mechanism. These should be required in order to enable effective provisioning and cost optimize the operations of a VO.

The VO or service overlay is required to dynamically and expressively construct descriptions of network, computing and ad hoc resources, locate and reserve these in a timely and ordered manner, in order to provide the services required by the VO. Such an infrastructure could for example enable:

- A VO infrastructure to expand or shrink dynamically as required
- Pervasive computer and network services, allowing the VO to orchestrate services/ connectivity
- Dynamically obtain the best price for a service.

This paper specifically provides a discussion of business models, communication models and QoS specification in the context of moving the Grid from e-Science toward the UC^{2}P infrastructure. It is proposed that these areas need to be addressed as we look into the extension of Grid like computing and overlay networks towards the future UC^{2}P infrastructure.

The paper firstly looks at the business oriented computational model diversity that could be required within the UC^{2}P infrastructure. The second aspect explored is that of the Grid communications model and QoS specification, where we discuss enhancements that could be envisaged.

2 Business Model Diversity Requirements

There are many business models that might be enabled in the future UC^{2}P, two examples of which are highlighted in this paper. TINA (Telecommunication Information Networking Architecture) business model and computational model represents the height of the wireline telecommunications industry model, shown in Figure 1 [10]. Figure 3.shows the second model of the Ambient Networks for Beyond 3G (B3G)[11] wireless networking, providing on demand network connectivity.

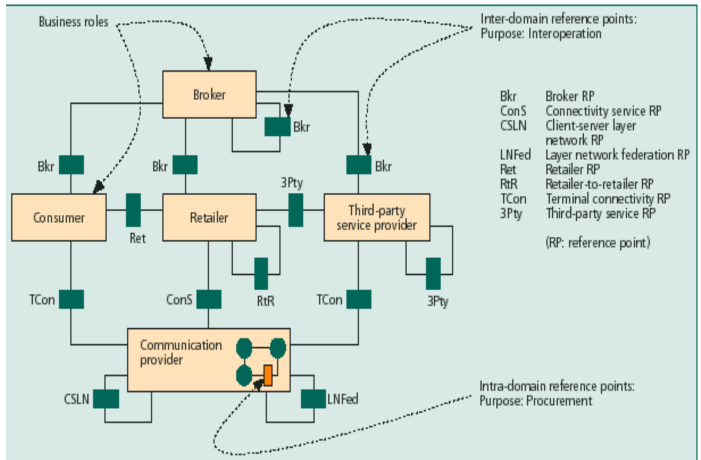

Figure 1 TINA Telecoms Business Model, from [10]

The TINA model provides a number of business role definitions and inter-domain reference points interconnecting the roles. Firstly the broker acts as a means for locating services and service providers. The Retailer sells services to Consumers by interacting with 3rd party Service Providers, e.g. content providers, and providing QoS enabled connectivity through the Communications Providers [10]. The reference points in the model are defined in ODL (Object Definition Language) comprising multiple interfaces and methods covering broadband service provisioning.

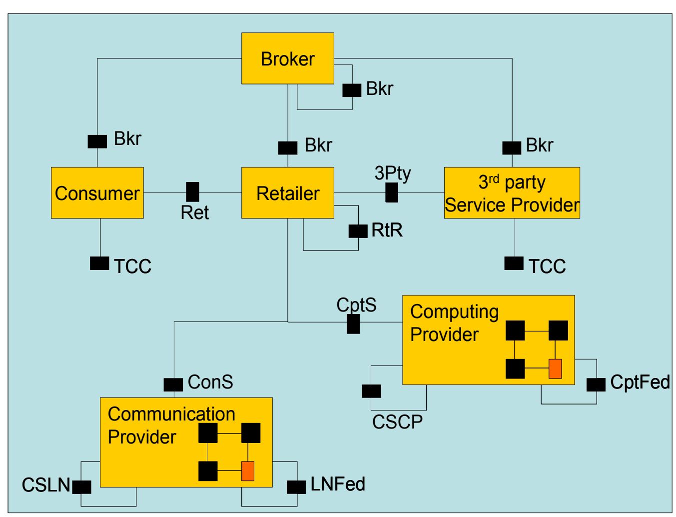

Figure 2 Extended TINA Business Model

An extension of the TINA Business Model to encompass the UC^{2}P model is shown in Figure 2. Terminal connectivity is replaced by a "TCC" interface for Terminal Computing and Communications control, but the main difference is in the interface from the retailer to the communications and the newly defined computing providers. The TCC allows the network and utility computing providers to interact with the terminal environment, e.g. to instantiate service objects and stream connections. The new business role added is the (Utility) Computing Provider which is interfaced through the CptS (computing service provider) inter-domain reference point. The CptS reference point allows the retailer, to provision computing resources (processors, storage and I/O), similarly to ConS which enables the Retailer to provision communication services with communications providers. Similarly Computing Providers can create federations through the CPFed (Computing Provider Federation) reference point and create client-server relationships through the CSCP (Client-Server Computing Provider) reference point, e.g. for sub contracting computing resources. The Computing Provider also has an intra-domain reference point for procurement of computing resources, similarly to the communication provider model.

This rather compartmentalized scenario based on TINA architecture is only one of many potential business models that could be employed in the future.

The model from Ambient Networks [11] is a "beyond 3G" network scenario from the Wireless World Initiative. The model is of flexible relationships that can be transient such as a vehicle or personal area network interaction between passing individuals or devices, which allows flexible wireless infrastructure decisions to be adopted.

In the Ambient Network flexible terminals based on software defined radio can connect to networks in an ad hoc manner through the ANI (Ambient Network Interface) and access services through the ASI (Ambient Service Interface), shown in Figure 3. The ANI and ASI are standard interfaces to the Ambient Control Space (ACS) which provides mobility, security, access, service delivery (media and content) and management. The aim of the Ambient Network architecture is to provide scalable and affordable wireless networking geared towards increasing competition *and* cooperation. This architecture also aims to facilitate the incremental market introduction of new services [11].

In the Ambient network the individual can act as a service provider and as a consumer. The concept is a development of the I-centric communications model of the Wireless World Research Forum (WWRF). The Ambient network enables ad hoc mesh networking potential creating hop-by-hop structured mobile systems. Relationships can be ad hoc network (e.g. on a train) and long term (e.g. with your workplace network (VPN etc.), through to your wireless telephony provider.

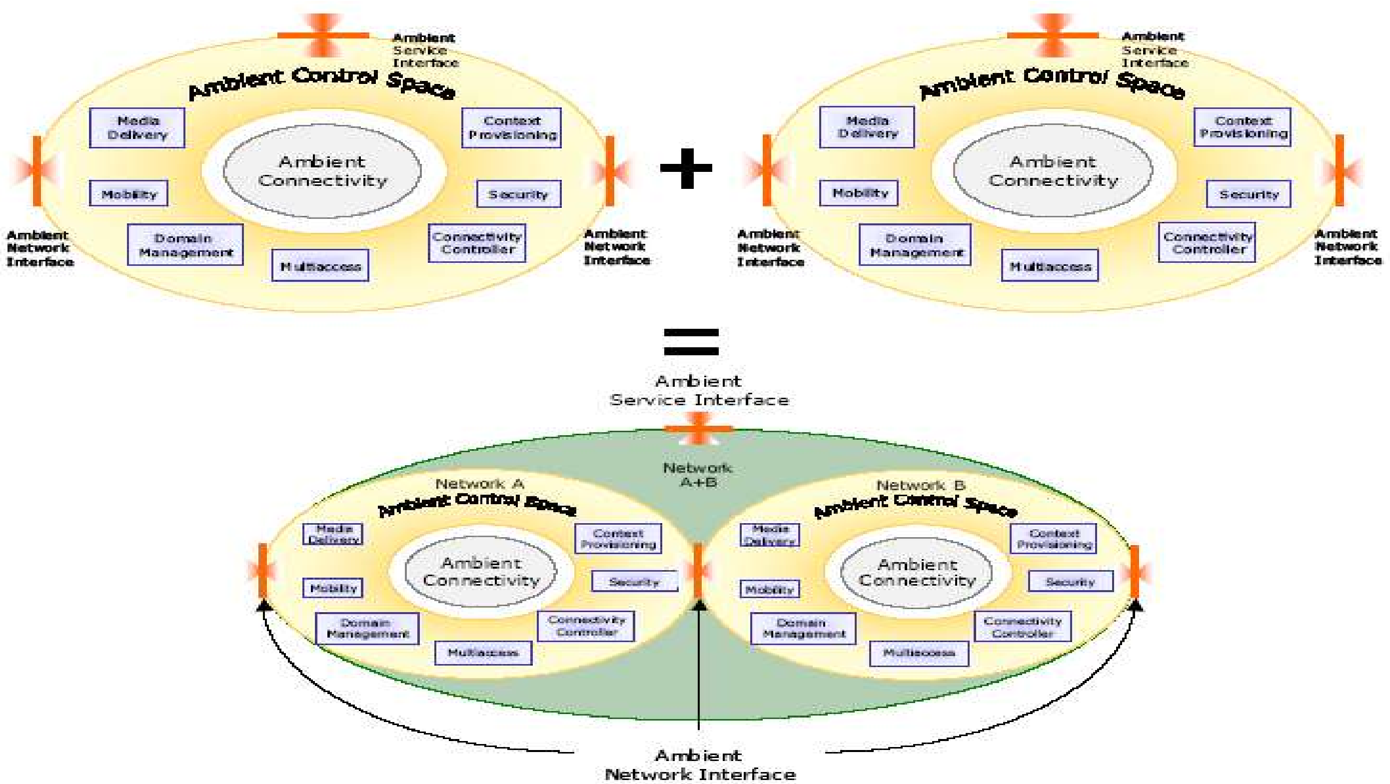

Figure 3 Ambient Networks composition from [34]

The ad hoc networking capability of the Ambient Network allows entities to negotiate connectivity and access to services through the ANI and ASI which on the fly dynamically compose networks into larger entities. The example in Figure 3 shows the composition of two networks to create a new composed network entity “A+B”, which could be for example a VPN grouping at the workplace.

Within the Ambient environment we need to capture:

- How do we describe what is going on?
- How do we describe who we communicate with?
- What computing/communication resources is required
- When do we want the resources?
- What service level agreement do we require?
- What is the service going to cost us and what is the best price I can get that service for?

Furthermore we can say that the UC^{2}P should enable or provide:

- Flexible service provisioning
- Flexible service delivery (networks/computing)
- Optimised object/service locality
- Location based service delivery
- Dynamic resource allocation
- Optimised lifecycle management
- Minimal price for resources to be deployment

The baseline scenario for our initial VO models is based on TINA like business model as these are well known. An M2M over cellular service provider scenario was explored initially [8], whereas in Figure 5 a content distribution model is envisaged. In each case the classic value added service provider value chain in Figure 4 can be envisaged. Here the value added service provider resells communications services by aggregating user/enterprise demand over a set of resources where contention/sharing provides business value. Secondly the value added service provider seeks to optimize the processing and storage infrastructure required to serve its customer base.

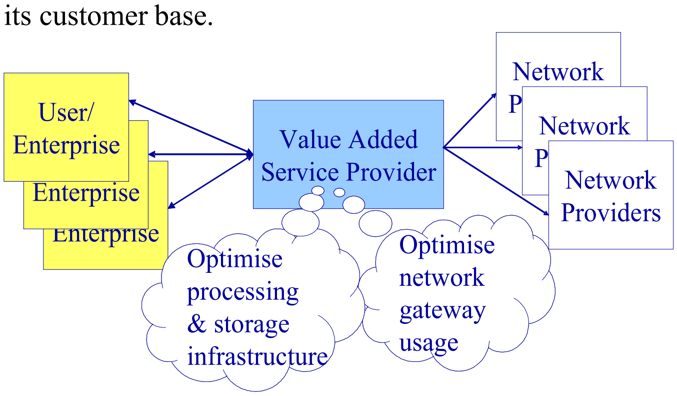

Figure 4 Initial Value Added Business Model Explored

Figure 5 illustrates an example of a content distribution overlay network scenario which provides a cluster of servers for head-end content storage, a cluster of virtual machines for management (AAA, billing, etc.) and video on demand (VoD) servers connected to content and management through backbone and access networks. An alternative scenario based around mobile value added service provisioning is explored in [8]. In such example applications the resource specification needs to provide versatile and extensible specification of required resources and any alternatives; in our current mobile agent middleware implementation called NOMAD [12][13], this is accomplished through the use of the Resource Description Graph (RDG) [14][15] which could describe both the resources available within an Ambient Network as well as a Virtual Organizations application requirements.

NOMAD is just one of several middleware options for dynamic resource allocation, utilizing an economic based resource allocation mechanism. NOMAD uses an auction based mechanism to provide distributed economic resource allocation to provide lowest cost provisioning of resources and the fair allocation of resources [16].

In NOMAD, loosely coupled cooperating machines called Depots provide resources such as CPU, memory, disk and additionally specific hardware and software components, e.g. certain communication bearers, resources such as codecs, java classes etc. A Depot can be made up of a number of computing and communications resources, plus security components such as firewall and Intrusion Detection Systems like a conventional Grid computational node [17]. In recent work the NOMAD resources that can be allocated have been extended from computer resources to a more extensive set of resources such as network based services [8][15].

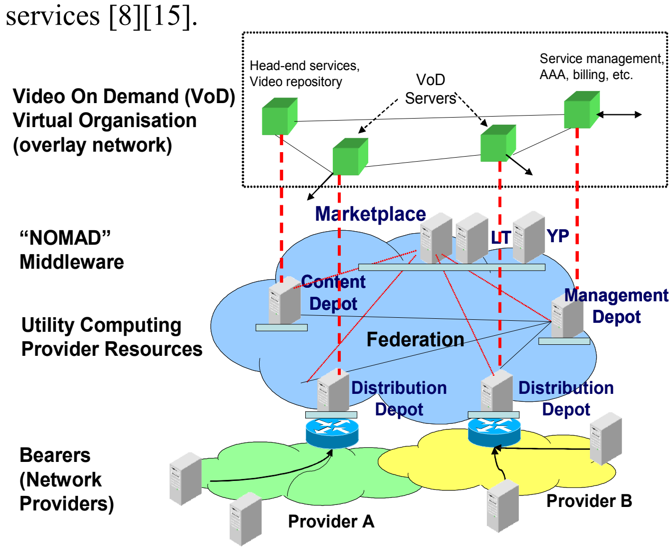

Figure 5 Scenario Overview

From a network perspective the resource description must include the properties outlined in section 3 for QoS provisioning of network resources. Computational resources are described similarly [12][15] and should include the information provided in the Grid RSL language [18].

Effective mechanisms for resource matching in real-time have been initially studied and shown to be achievable in a few milliseconds [15] even for brute force matching of VO resources to agent resource allocation. These measurements were undertaken on a small scale laboratory setup, however moving to a WAN will add transmission, propagation and queuing delays to the performance of the system. Specifically for Utility Computing/Grid based overlays [6] we need to support dynamic resource description, resource matching and resource allocation of both computing and network platforms to enable dynamic service deployment. For example a dynamic overlay could be envisaged for content distribution, such as the deployment of multicast streaming video of live events or video on demand services, e.g. for e-Learning, where the number of users and the amount of content varies throughout the day, week, month and year [19][20].

The resource description, matching and economic valuation components are a key requirement of the middleware framework in such a system. However there are many other properties of frameworks that need to be provided to effectively enable this vision of flexible UC^2P based service overlays. These would include the following (which is not an exhaustive list):

- Flexibility – from large complex virtual organizations though to single users and their services
- Scalability – worldwide numbers of users, VOs, network elements, networks and applications/services
- Deployment/Lifecycle Management – cross domain with commercial constraints (trust management, reputation, contracts, allocation, pricing etc.)
- Resiliency – from network related resiliency such as overcoming network failures to service resiliency
- Security – integrated from network to the service layer
- FCAPS (Fault Configuration Accounting, performance and Security) management – both Intra- and Inter-domain management
- Flexible business model support – as discussed above
- To these can be added Grid and overlay specific framework aspects, such as:
  - Cross domain co-ordination of resource allocation from services to network
  - Overlay network and workflow support

In section 3 we briefly overview one key aspect – the current networking model in Grid Computing and current work in the description of QoS in the networks. This is the baseline for network resource description and the network models currently in Grid computing systems from which we can build future communication models.

3 Communications and QoS in Resource Description

Quality of service (QoS) and class of service (CoS) in the Internet can be enabled through Integrated Service (IS) and Differentiated Service (DS) models, however today most customers still only have a best effort service. To enable QoS and/or CoS in networks and the potential new services provided through overlay networks, significant control and management information is required for provisioning, maintaining, monitoring, accounting and billing. Heterogeneous devices from supercomputers and Grids, through to hand held devices require appropriate connectivity, which can include wireline and wireless connectivity, plus user, session and device mobility.

Quality of service encapsulates a wide variety of non functional properties such as reliability, performance, security, and timing/workflow. Although there are many QoS services available, applications cannot easily take advantage of QoS, one possible reason is the lack of a universal description language [21] which should be rectified in the future. As with resource allocation the negotiation of QoS parameters requires a universal language so that all hosts and network elements can understand, although currently there are several QoS languages available including: XQoS [22], QuO [23] and SLAng [24] which cover different network aspects, which are briefly introduced in this section.

Management based on the client-server SNMP model to provide FCAPS functionality is not expected to scale to enable the NGI/Grid/overlay network infrastructure [21]. Using decentralized management frameworks we obtain advantages of scaling and operational speed, however maintaining state and service information becomes more complex. In fact with overlay networks defining virtual Internets consisting of virtual routers and hosts plus tunneled links, e.g. for virtual private network services [20] the complexity of service and network management increases as the number of client-server relationships and management information increases. This increase in management complexity needs to be overcome to enable effective deployment of overlays and new services.

The QoS description need to incorporate a variety of different levels as discussed below, from a Service Level Agreement (SLA), which can span multiple domains, through to fine grained QoS descriptions. SLAng [24] provides a reference model for inter-organisation service provisioning over multiple levels: storage, network, middleware, and application. SLAng provides a format for the negotiation of QoS and uses a language suitable for automated reasoning. The focus of SLAng is the provision of contracts between entities, which are represented by Service Level Agreements (SLA).

SLAs describe an agreement between a customer and a provider that include technical and non-technical characteristics of the service. These include QoS requirements and a related set of metrics with which the provision of these requirements is measured. The SLAng syntax is XML based, defined using an XML schema. The justification for this choice is the favorable integration with existing service description languages (especially in an e-business environment). The SLA contains descriptions of the contractors, contractual statements (e.g. start time), service level specifications (SLS). Resource matching is not explicitly mentioned, the use of XML allows matching via generic tools.

QML[25][29] is a specification language developed by Hewlett Packard, its aim is to support specification in an object orientated context derived from work in TINA [10]. QML has three major abstraction types for QoS: contract type, contract, and profile. QML makes use of a set of quantifiable properties, where QoS is characterised along dimensions, which are grouped into categories. For example dimensions latency and throughput may be in a category called performance.

Each QML service specification includes an interface and a QoS profile. The interface describes the operations and attributes exported by the server, the profile describes the QoS and specifies the attributes and operations exported by the server. Additionally an extension to QML has been developed to allow QoS profiles to be associated with relationships, thus aiding the design of more complex systems. As yet there has been no work on ensuring QoS requirements are satisfied, and quantification of QoS server performance.

BBN technologies developed Quality Objects (QuO) [23][30], a framework for providing QoS in distributed applications. QuO is based on CORBA (Common Object Request Broker Architecture) and supports runtime performance tuning and configuration by specifying operating regions, behavior, alternatives, and reconfiguration strategies thus enabling resiliency and other adaptation features. These specifications allow QuO to adaptively trigger runtime reconfiguration as system conditions change, for example an operation transitions between operating regions.

QuO is designed as a higher level programming language for specifying application QoS requirements. This allows programmers to specify higher level aspects of real-time requirements, such as the relative priority of events, and the trade off between real-time and other QoS requirements. QoS levels (expected and measured) are characterised by regions which specify measurable quantities such as latency which is established when the client and server agree upon a specific latency region. QoS levels are monitored to ensure they remain within the expected region.

QuO is described using a Quality Description Languages (QDL) made up of three aspect languages: a Contract Description Language (CDL), a Structure Description Language (SDL), and a Resource Description Language (RDL). CDL describes the contract between a client and an object, including the QoS required by the client, the QoS the object expects to provide, and a range of possible levels. SDL describes the structural aspects of the application, including adaptive behaviour and interfaces.

Finally XQoS [22] is a QoS specification language, which makes a mapping between the QoS needs of an application and the underlying communication services. The XQoS design is based on a formal model called the Stream Petri Network Model (TSPN), which offers a formal context for the modeling of synchronisation constraints within weakly synchronous systems, e.g. for specification of multimedia synchronisation scenarios in terms of time, sync, order, and reliability. The system allows multimedia applications to negotiate their QoS requirements, which are turned into QoS specification. The results [19] show the positive impact on multimedia performance using a universal description language.

In Grid computing support for networking is through GRAM (Grid Resource Allocation and Management protocol) and GARA (General purpose Architecture for Reservation and Allocation). GRAM deals with resource discovery through service registration to some centralized management entities. GARA provides a framework for various underlying resource schedulers to provision applications with resources, including storage, computing and networking resources described through the RSL (Resource Specification Language) [18].

Grid computing can support optical (wavelength, SDH/SONET) networking, for example the LightPath Services offered by CA*net4. The LightPath [28] network resource architecture provides a TINA like network control hierarchy. The LightPath architecture provides User, Service, Resource (network) and Physical (element) layers with support for end-to-end, inter and intra domain interaction. Optical networking through GLIF (Global Lambda Integrated Facility) [27] which defines interfaces and protocols for the control planes of the contributed Lambda resources for example with the CPL (Common Photonic Layer) from Nortel.

Support for optical networking, as well as Ethernet/LAN services provides support for storage area networking and Enterprise connectivity services. Currently the LightPath network resource architecture supports layer network federation (i.e. TINA LNFed reference point). This could be extended with TINA intra-domain procurement and client-server network composition functionality, and thus can be applied to the communications infrastructure for the UC^2P infrastructure [31].

4 Further Work

This paper has overviewed a potential business and service architecture in future NGI deployments, consisting of Utility and Grid computing resources and overlay networks. The creation of value added services through the combination of computing resources and QoS network resources can enable VOs, potentially scaling from an individuals requirements through to large organizations utilizing current advances in network deployment, such as LightPath [28] and Grid resource management overlays [26] to create dynamic services. This combination may provide a useful building block for tomorrows NGI deployments.

A brief summary of QoS resource descriptions that can be harmonized to enable the deployment of VOs. This input should drive QoS and network description work in the NOMAD middleware, in which the development of VOs is actively being undertaken. Discussion of network design architecture based on TINA architecture and business model considerations should be taken into consideration in the initial development of Grid network infrastructure.

Future work in this area aims to develop the framework for an extended communications model in the context of Grid computing and overlay networks. The work aims to extend resource descriptions and economic models for resource allocation, middleware services and the development of VOs and business models, within the Distributed Systems Research Group at Victoria University of Wellington. This work has been funded by Victoria University of Wellington as the Network Enabled Virtual Organisation (NEVO) project.

5 References

[1] V. Machiraju, J. Rolia, A. van Moorsel, "Quality of Business Driven Service Composition and Utility Computing", Software Technology Laboratory, HP Laboratories Palo Alto, HPL-2002-66, March 15th, 2002.

[2] M. Kallahalla, M. Uysal, R. Swaminathan, D. E. Lowell, M. Wray, T. Christian, N. Edwards, C. I. Dalton and F. Gittler, SoftUDC: A Software-Based Data Centre for Utility Computing, IEEE Computer, November 2004, Volume 37, Number 11, pp38-46.

[3] I. Foster, C. Kesselman, S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations", International Journal of Supercomputer Applications, volume 15(3), 2001.

[4] Abbe Mowshowitz, Virtual Organization, Communications of the ACM, September 1997, Vol. 40 No. 9.

[5] Grid-based Application Service Provision (GRASP) http://eu-grasp.net/english/ last access 24 July 2004

[6] T. Dimitrakos, et al. "An Emerging Architecture Enabling Grid Based Application Service Provision", 7th IEEE International Enterprise Distributed Object Computing Conference EDOC 2003, Brisbane Australia, September 16-19, 2003.

[7] P. Komisarczuk, K. Bubendorfer, K. Chard, Enabling Virtual Organisations in Mobile Worlds, IEE 3G2004 conference, 18-20 October 2004, London, UK.

[8] A. Abbas, GRID COMPUTING: A practical Guide to technology and Applications, Charles River Media, ISBN 1-58450-276-2, 2003.

[9] The Globus Alliance, Globus Toolkit 4.0 Release Manuals, http://www-unix.globus.org/toolkit/docs/development/4.0-drafts/, last accessed 8th March 2005.

[10] Hendrik Berndt, Takeo Hamada, and Peter Graubmann, TINA: Its Achievements and its Future Directions, IEEE Communications Surveys and Tutorials, First Quarter 2000, Vol. 3, No. 1

[11] Norbert Niebert, Hannu Flinck, Robert Hancock, Holger Karl, Christian Prehofer, Ambient Networks - Research for Communication Networks Beyond 3G, 13th IST Mobile and Wireless Communications Summit 2004, Lyon, France, 27-30 June 2004.

[12] K. Bubendorfer, "Nomad: Towards an Architecture for Mobility in Large Scale Distributed Systems", PhD Thesis, Victoria University of Wellington, December 2001, http://www.mcs.vuw.ac.nz/research/dsrg/papers/KBthesis.pdf, last access 24 July 2004.

[13] K. Bubendorfer, J. Hine, "NOMAD: Application Participation in a Global Location Service", Lecture Notes in Computer Science, number 2574 in LNCS, pages 294-306, Springer Verlag, January 2003. Proceedings of the 4th International Conference on Mobile Data Management, 21-24 January 2003, Melbourne, Australia.

[14] K. Chard, Efficient Dynamic Resource Description, Victoria University of Wellington, October 2004, Honours project report.

[15] K. Bubendorfer, P. Komisarczuk, K. Chard, Efficient Dynamic Resource Specifications. To appear in 6th International Conference on Mobile Data Management (MDM'05), Ayia Napa, Cyprus, May 9-13, 2005

[16] S. Antlif, "Experimental verification of collective vickery auctions in NOMAD", Victoria University of Wellington, 2003. http://www.mcs.vuw.ac.nz/comp/graduates/archives / honours/2003/antliff-simon/antliff-report.pdf last access 24 July 2004.

[17] C. Chaubal, Grid Computing Engineering, Radoslaw Rafinski, "Building a Global Compute Grid—Two Examples Using the Sun™ ONE Grid Engine and the Globus Toolkit", Sun Microsystems, Poland, Sun BluePrints™ OnLine, April, 2003

[18] GLOBUS. The Globus resource specification language rsl v1.0. http://www.globus.org/gram/rsl spec1.html, 2000. last checked 1/9/04.

[19] Joseph D. Touch, Yu-Shun Wang, Lars Eggert, Gregory G. Finn, A Virtual Internet Architecture, ISI Technical Report. March 2003, http://www.isi.edu/touch/pubs/isi-tr-2003-570/isi-tr-2003-570.pdf

[20] Robert Braden, David Clark, Scott Shenker, John Wroclawski, Developing a Next-Generation Internet Architecture, July 2000, http://www.ana.lcs.mit.edu/papers/PDF/NewArch_wp_v1.pdf

[21] M Eder, S. Nag, Service Management Architecture Issues and Review, RFC 3052, January 2001.

[22] Exposito, E., Gineste, M., Peyrichou, R., Snac, P., And Diaz, M. XQoS: XML Based QOS Specification Language. MMM'03, The 9th International Conference on Multi-Media Modeling, Taiwan, January 7-10, 2003, pp. 114–134.

[23] BBN TECHNOLOGIES, Quality Objects (QuO). http://quo.bbn.com/, © 2002. Last checked 22/11/04.

[24] Lamanna, D., Skene, J., And Emmerich, W., SLAng: A Language for Defining Service Level Agreements, The Ninth IEEE Workshop on Future Trends of Distributed Computing Systems (FTDCS'03), San Juan, Puerto Rico, May 2003.

[25] Frolund, S., And Koistinen, J. QML: A language for Quality of Service Specification. Technical Report HPPL-98-10, Hewlett-Packard Laboratories, February 1998.

[26] Y. Huang, S. N. Bhatti, Decentralised Resilient Grid Resource Management in Overlay Networks, IEEE International Conference on Service Computing, SCC'04, Shanghai, China, September 15 - 18, 2004

[27] GLIF, Global Lambda Integrated Facility, http://www.glif.is/about.html, last accessed November 2004.

[28] Raouf Boutaba, Statement of Works, User Controlled LightPath Project, http://bbcr.uwaterloo.ca/~canarie/SOW.pdf, accessed November 2004.

[29] Frolund, S., Koistinen, J., Quality of Service Aware Distributed Object Systems, Proceedings of the 5th USENIX Conference on Object-Oriented Technologies and Systems (COOTS), San Diego, California, USA, 1999, Hewlett-Packard Laboratories.

[30] Loyall, J. P., Bakken, D. D., Schantz, R. E., Zinky, J. A., Karr, D. A., Vanegas, R., And Anderson, K. R. QoS Aspect Languages and their Runtime Integration, Proceedings of the 4th Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers (LCR), 1998, vol. 1511, Springer-Verlag.

[31] Maniatis, S. I., Nikolouzou, E. G., And Venieris, I. S. End-to-End QoS Specification Issues in the Converged all-IP Wired and Wireless Environment. IEEE Communication Magazine, June 2004, 80–86.