Next Generation Networks and the Internet

By Glenn Gilmour and Joshua Saunders
Research, Statistics and Technology Branch of DCITA

A paper for the Communications Policy & Research Forum

25-26 September 2006

Acknowledgements: Dr Geoff Peak (RS&T) for helpful hints and comments to the authors while researching and writing this paper.

Disclaimer: *The views expressed in this paper are those of the authors and do not necessarily represent those of the Department of Communications, Information Technology and the Arts or of the Commonwealth Government.*

INTRODUCTION

This paper describes some technical and economic aspects of ‘next generation’ networks and the relationship of a next generation network to the Internet. The paper is meant to give an understanding of the technology and protocols that underpin next generation networks and a review of pricing mechanisms that may be utilised for next generation networks. The paper is divided into two sections. Part 1 covers the technical aspects of next generation networks, while Part 2 covers the economic aspects of pricing.

The technical aspects discuss the Internet as a ‘network of networks’ and why a next generation network by virtue of using common Internet Protocols, is a constituent network of the Internet. Next generation networks are viewed from the perspective of the contractual relationships embodied in peering and transit relationships, and the interconnection protocols that make connection to the ‘global’ Internet possible. A brief description of the Internet as a set of underlying Autonomous sub-systems is provided.

The network topology of a next generation network is then described by analysing the network’s constituent ‘layers based on the Open Systems Interconnect model, and how current telecommunications networks, embodying the PSTN, are likely to transform on a layer-by-layer basis to an IP packet switched network supporting the ‘triple play’ of voice, high speed Internet, and video streaming, and new forms of distributed network services, such as file sharing.

A rudimentary analysis of basic Internet Protocols (IP) and IP packet headers is undertaken to show how service classes and congestion fields within packet headers could be utilised for different levels of Quality of Service (QoS) for network products.

The second part of the paper analyses from an allocative efficiency viewpoint how congestion and priority are currently handled by the Internet, and what forms of future pricing models might be used to prioritise Internet traffic in congestion situations. As ‘next generation’ networks appear utilising IP protocols to handle both ‘on-net’ proprietary traffic, and ‘off-net’ traffic destined for the remainder of the Internet, different levels of bandwidth resources and priority may be required by different network services, and this may need to be allocated by a combination of static and dynamic pricing systems.

Part 1: Technical Aspects of Next Generation Networks

The Internet is a 'Network of Networks'

The Internet is not one single network but rather a "network of networks" consisting of 'public' and 'private' networks that interoperate to form the 'global Internet. The 'Public' networks are Government, and quasi Government institutions (.gov extension) while the 'private' networks encompass organisations (.org extension) and companies made up of large enterprises, telecommunication carriers, and other service providers (.com extension). Whether the economic definition of a 'public good', a good which is non-excludable, and non-rival or non-depleteable pertains to the Internet, is dealt with below in the context of congestion and priority pricing.

The constituent networks interoperate technically by using common transmission protocols, such as IP, and commercially by contractual arrangements between the network owners. These contractual arrangements are usually of two types:

- peering relationships which agree to swap equivalent traffic volume across reasonably equivalent routes, sometimes called 'bill and keep'; or 'sender keeps all' and
- transit contracts, essentially paid 'non-equivalent' traffic volume exchange where the larger network is able to provide routes to the Internet for the smaller network.

It is these contractual relationships that provide 'any-to-any' connectivity across the global Internet, by advertising the routes which are reachable from their individual network.

The transformation of the PSTN to a next generation network for communication is being driven by 'edge intelligent' user devices: fixed, mobile, or portable, which are identified by a unique IP address. The IPV6 standard allows a large enough address space to give every IP device a unique global address and hence the promise of global connectivity wherever the device is, but a challenge to the development of global interoperability standards for connectivity and QoS across the 'network of networks'

The figure below, which is illustrative only, shows two devices A and B which connect and communicate across the 'network of networks'. The first part of the figure shows the two devices exchanging packet data which collectively form an 'IPStream' in each direction. It is assumed that they have previously established a connection for this exchange. Looking 'deeper into the cloud' the second part of the figure shows that device's A home Internet Service Provider (ISP) is network 2, while B's is network 4. A's 'IPstream' reaches B by paying for transit through network 1. B, however, reaches A by having a peering arrangement with network 3, who in turn has a peering arrangement with A's network 1. B's network 4 also has a transit relationship with network 1.

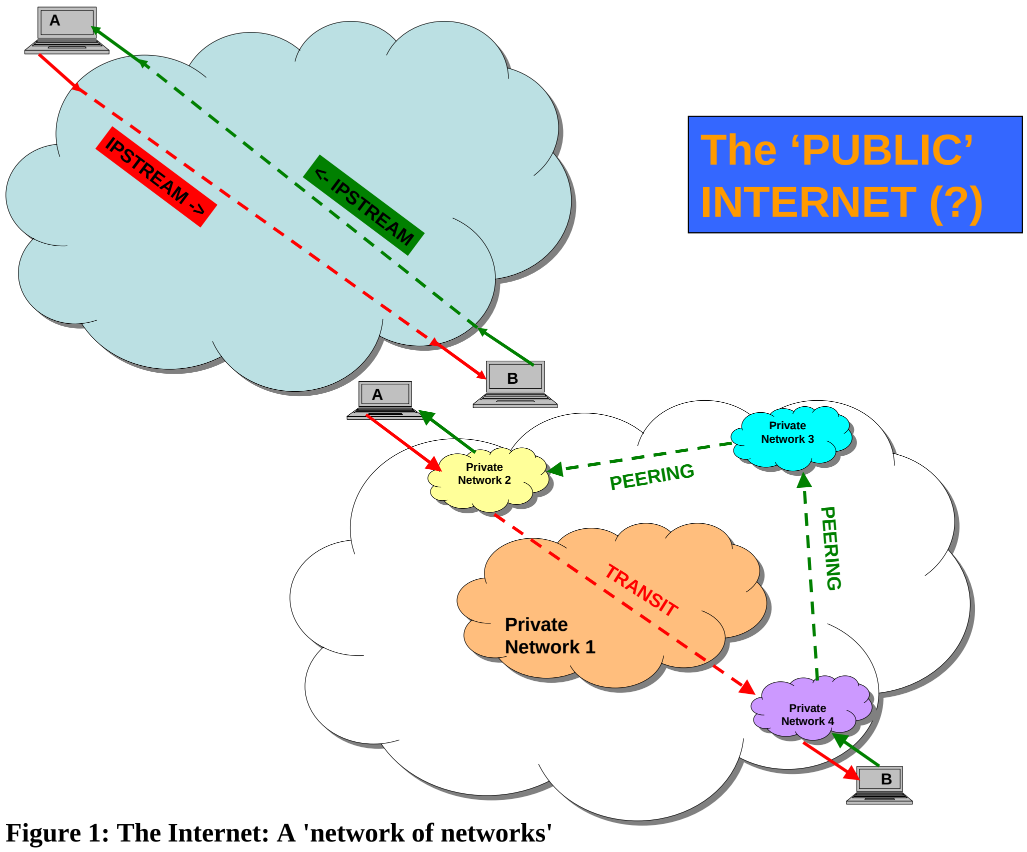

Figure 1: The Internet: A 'network of networks'

For the purposes of exposition network 1, the largest, is vertically integrated in terms of being a carrier and a large scale ISP, and is at the highest tier of the Internet for its operational region which theoretically means it can see all routes on the Internet in its region. These providers are also known as Internet Backbone Providers (IBP). This has been referred to in Australia as the ‘Gang of four’, consisting of Telstra, Telecom New Zealand, Singtel-Optus, and MCI, the latter when they had ownership of Ozemail.

The routes encapsulated by the transit contracts that network 1 has means that it can offer not only its own end-users routes to its peering partners, but also to its customer’s customers, for example user A whose ISP (network 2) buys transit from network 1.

Depending on the relative cost of peering compared to transit for network 2, which is a function of the volumes exchanged, and the extra routes exposed by network 1, and to a lesser extent network 3, it may benefit network 2 to peer directly with network 4. Alternatively, it could re-route the traffic bound for network 4, via network 3 which it has a peering relationship with, and then rely in turn on the peering relationship between network 3 and 4, to forward packets to device B.

However, network 1 may provide transit to network 2 for many more Internet routes than that needed to reach networks 3 and 4, and as network 1 is vertically integrated it can offer a bundle of routes (transit services) as well as underlying carriage or transport. Network 1, as discussed above, peers with larger IBP networks (not shown in the figure) and may rely on the transit relationships it has with networks 2, 3, and 4, in order to peer at the higher tier. In this way market forces can play out between the tiers as to the relative price of transit via a higher tier, to the cost of peering at a lower tier between the smaller networks 2, 3 and 4.

There are many such transit and peering relationships supporting the 'public' or 'global' Internet, which are mostly enforced between the boundaries of the constituent networks. The majority of these boundaries are between what the Internet defines, as an autonomous systems (AS). Each AS is a collection of IP networks and routers under the control of one entity that presents a common routing policy to the Internet. The AS announces to the rest of the Internet the routes (groups of IP addresses) that are reachable from it, which as discussed above includes:

- the AS entity's own customers;
- customers of ISPs that use the AS for transit; and
- customers of the AS entity's peering partners.

Unique Numbers are assigned to Autonomous systems which are registered and maintained by the relevant regional Internet registry in the same way that domain names are registered and maintained.

A common protocol known as the Border Gateway Protocol (BGP) implements the common routing policy presented by each AS, allowing the AS's to communicate. Unlike other routing protocols which have some form of 'hello' protocol for automatic neighbour discovery BGP routers at the edge of an AS must be configured manually so that the neighbours 'discovered' are those that the AS has peering and transit contractual relationships with, hence the term 'policy based routing'. [1]

Policy based routing extends the Internet to not just advertise and forward packets based on IP address routes, but to also combine IP addresses with:

- port numbers (see below for explanation);
- IP precedence fields to determine classes of service and hence priority (see below for more on priority),
- protocols and
- packet size.

In this way packet filtering or 'traffic shaping' using combinations of these attributes develops route maps and access lists to permit / prioritise, or deny entry of packets to the AS.

Layer Topology of Next Generation Networks

The PSTN is a set of circuit switches that in many countries are now some 10 to 20 years old. With the rise of IP based routers, to support the new edge intelligent devices taken-up by the population, the economic depreciation profiles[1] of the PSTN legacy switches may fall even faster. However, a next generation network is likely to be based on an incumbent's underlying PSTN infrastructure and will re-use many of the lower level PSTN infrastructure components such as:

- copper in the access network based on the xDSL standard which allows previously unused capacity in the form of higher frequency bands to be utilised;
- fibre in the core network, by utilising a process known as Wave Division Multiplexing (WDM) each fibre strand can be "lit up" to provide a number of light waves each light-wave carrying a separate traffic stream, providing large amount of bandwidth capacity;
- existing ducts, trenches, conduits, and right-of-ways; and
- existing exchange buildings and network operation centres.

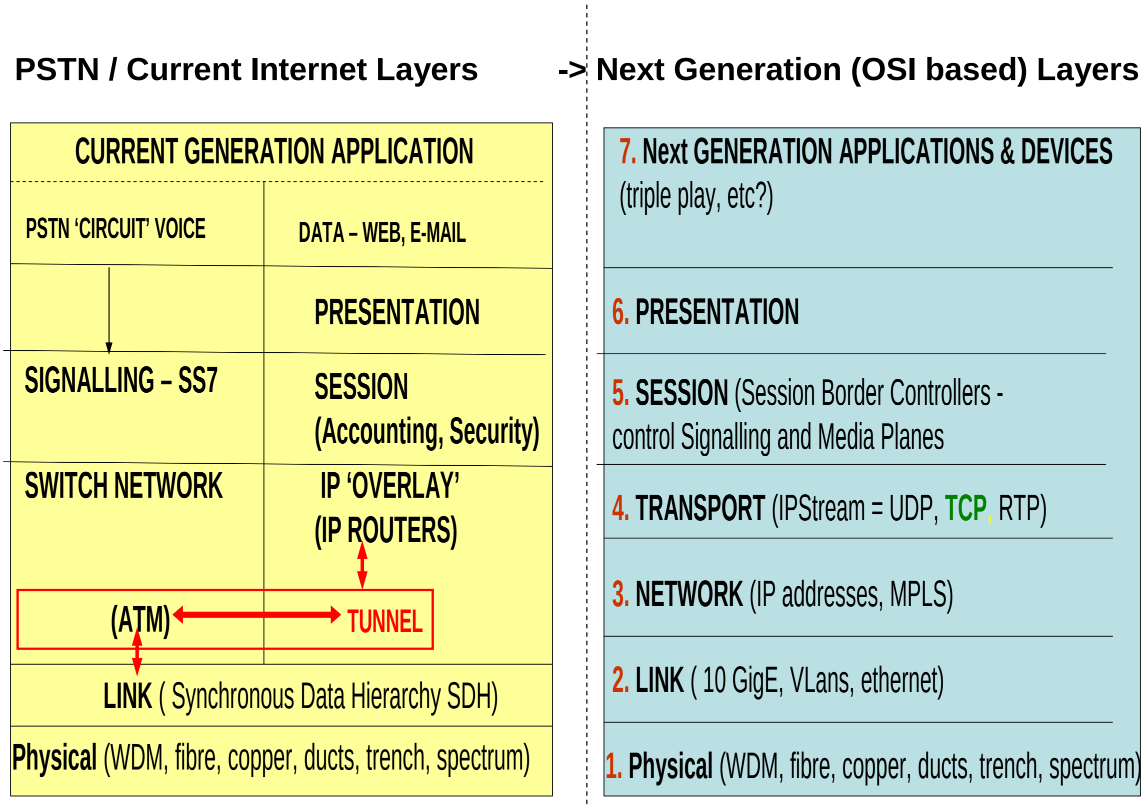

Figure 2: Network Layers

---

[1] The asset still has an economic life if it generates an excess of revenues over costs. Economic depreciation is the annual change in economic value defined as the discounted present value of expected future revenues from the output produced by the asset, less the present value of associated future operating costs of the asset.

The transformation of the PSTN to a ‘next generation’ network can be analysed from the point of view of considering the PSTN and the next generation network as consisting of a set of layers[2], as shown in the figure above.

The left hand side of the figure above shows the PSTN and an approximate representation of how the current Internet is structured, sharing as it does the lower layers (1 and 2) with the PSTN circuit switches for voice traffic. The IP layer can be regarded as an ‘overlay’ in that it ‘tunnels’ through the Asynchronous Transfer Mode (ATM)[3].

Application Layers 6 and 7

The right hand side is the representation of the ‘next generation’ network in which applications / devices originate and accept IP packets whose payload could be voice, data or media (pictures and sound) - layers 6 and 7.

Session Layer 5

The session layer (layer 5) carries out signalling between devices (often called the ‘signalling plane’ in a ‘next generation network) so that connections for two-way or one-way packet exchange between devices can be established. This layer is responsible for Authorisation, Authentication, and Accounting (‘AAA’), and must also determine that end-user devices can exchange packets, by determining if the devices have compatible codecs[4]. If this is not the case the session layer may undertake to arrange for trans-coding of the ‘media’ exchange between incompatible codecs.

The Session Layer may also need to perform directory look-up services if VoIP is to replace PSTN voice calling services. ENUM is an example of such a service that enables telephone numbers to be converted into an IP address that can locate Session Initiation Protocol (SIP) or H.323[5] enabled devices by translating a telephone number to a Uniform Resource Identifier (URI) using the Internet’s Domain Name Services (DNS).

The DNS already performs conversion from web site addresses (http//:www.xxx.com.au) to IP addresses on a global scale and the idea behind ENUM is to use it as a global numbering plan organised into a similar hierarchical structure

---

[2] Seven layers are chosen here based on the Open Systems Interconnect (OSI) model, but other layer representations may be equally or even more valid.

[3] ATM was developed under the auspices of telecommunications carriers and the ITU to provide a networking standard that could support both synchronous channel networking (SDH) and packet-based networking (IP, Frame relay, etc), whilst supporting multiple levels of QoS for packet traffic.

[4] A codec (coder / decoder) resident inside a device is normally based around a Digital Signal Processor (DSP) which encode (transmitter end) and decode (receiver end) sound or video into digital bit frames utilising various compression techniques which are traded off against voice or picture quality.

[5] H.323was originated by the ITU, but is looked upon as a more complex standard being based on telephony engineering, in comparison to SIP which was originated by the IETF, and based on IP standards.

as for websites' Uniform Resource Locators (URLs) into the DNS structure of Top Level (TLD), second level (2LD), and third level (3LD) domains, such that a 'telephone number' can be translated into a URI.

The calling device uses the E.164[6] number to look up the appropriate DNS server to translate the 'telephone number' to the URI of the called device. The E.164 number can resolve to multiple URIs which leads to the possibility of 'find me follow me' applications, or preferences on behalf of the calling party as to where the call is to be received or alternatively preferences on behalf of the called party on where it wants the call to be received. Thus, networks 2 and 4, in figure 1 above, would rely on the Session layer to inter-connect devices A and B to provide a VoIP service across the Internet.

While it is necessary for the session initiation and session termination to be centralised, for 'AAA' and numbering translation services, this is not necessarily the case when IP devices begin to exchange data, voice or media. In fact the network soft-switches involved in the 'media' exchange can determine the shortest / least cost path for the 'media plane' transport and so optimise the chosen router path through layer 3 (see below) through the network, subject to 'cooperating' transit and peering relationships as outlined above. As the most efficient path, the media plane does not necessarily have to take the same routes that the signalling plane took.

However, there are also legitimate institutional and Government concerns involving fundamental applications such as: emergency calling, and legal interception in the face of encryption, that may need to operate across the 'islands' of 'next generation' networks. In order to allow monitoring and centralisation of the traffic moving in the 'media plane' Session Border Controllers (SBC) will increasingly interpose themselves between the soft-switches and the devices in 'next generation' networks as SBCs transit both signal plane and media plane exchange between devices.

Transport Layer 4

The transport Layer (layer 4) is responsible for ensuring packets which are routed through the Internet arrive at the destination device. If the communicating devices use the Transaction Control Protocol (TCP) then the destination device will notify the packet originating device to re-transmit lost packets. There is a cost to this reliability in the form of delay and this is discussed below in the context of congestion. User Datagram Protocol (UDP) means the devices are responsible for arranging any re-transmission protocol(s) between each other. This makes it fast but unreliable, as lost packets will not be re-transmitted if devices use UDP as is.

Real-time Transport Protocol (RTP) is used for video streaming which demands strict latency and jitter tolerances by time stamping packets so that they can be re-ordered as they flow along the network routing path.

Network Layer (layer 3),

---

[6] E.164 is the numbering standard being developed by the ITU.

The network layer (layer 3), can be thought of as the 'global' Internet's router layer that routes the packets across the Internet, according to 'advertised' routes (essentially IP address groupings), where routing 'policy' reflects the peering and transit agreements between the 'private' and 'public' networks that constitute the Internet, and controlled at AS borders by BGP as discussed above.

Link layer (layer 2)

The Link layer (layer 2) in the PSTN, known as Synchronous Data Hierarchy (SDH) is a carrier grade 'point to point' link which transfers the packets after switching and / or routing has been performed in the network layer. Overtime SDH is expected to transition to 10 Gigabit Ethernet (10GigE), although 10GigE is designed to tunnel through SDH, just as ATM can tunnel IP packets. Thus it is expected that telecommunication carriers will still be using and upgrading SDH for many years yet.

Of interest is that the Ethernet standard, which started as a 'within building' local area network (LAN), then evolved to support metropolitan-wide area networks (MANs) for larger enterprises and now with the 10GigE standard is being positioning to be used as a carrier grade transport at layer2. The GigE Ethernet standard also provides support for layer 2 switching by partitioning user(s) traffic into Virtual LANS (VLANs), and providing traffic class priority similar to IP protocol[7].

Physical (layer 1)

The physical (layer 1) consists mainly of civil infrastructure such as ducts, trenches, fibre, copper and spectrum rights, and it is here the bulk of the asset value in terms of replacement costs is likely to lie. Note WDM greatly expands the capacity of fibre[8] with each 'wavelength' transporting a SDH or 10GiGE traffic stream.

What is apparent is that in the conversion of the PSTN to a 'next generation' network there is very little transformation in the physical layer of the network, where the economies of scope for a telecommunications network are predominantly present. Rather the transformation takes place in:

- layer 2 (depending on the rate of replacement of SDH by 10GigE);
- layer 3, IP and Multiprotocol Label Switching (MPLS)[9] replacing ATM and Frame relay, as core switches are replaced by routers;
- layer 4 as devices move to all IP protocols such as TCP/IP and RTP for reliable transport;
- layer 5 where soft-switches replace class 5 switches providing support for PSTN call-setup for legacy POTS services, and 'AAA' for the 'next generation' of devices; and
- finally, the devices and applications themselves which rapidly proliferate. (layer 6 and 7).

---

[7] Traffic Class Priority is discussed further below.

[8] For example the NextGen AARNET optic fibre cable can on each fibre strand transmit at least 32 separate wavelengths each wavelength capable of 10 Gbps. A typical fibre cable has 48 strands.

[9] MPLS is a protocol which can expedite the router path, without the need to look up destination IP addresses in routing tables, by creating a 'switched soft circuit path' based on the MPLS label(s)which MPLS enabled routers place between the layer 2 and layer 3 (IP) header.

Congestion and Priority on the Internet

Currently, IP protocols such as TCP/IP and UDP/IP allow for "best efforts" packet delivery but not necessarily at priority or in order of transmission. If demand for Internet use is stochastic, utilisation may be such that not all users access the network at the same time, and hence congestion is avoided, and the Marginal Cost (MC) of packet transmission is essentially zero.

.

Devices and their associated Applications that use TCP guarantee delivery of the packet stream by imposing a sliding transmission window at the source device and only advancing the window over each packet or packet group when it is notified that packet has been received at the receiving device. A non-acknowledgment from the receiving device means that the packet must be re-transmitted and TCP will slow down the rate of transmission only returning to a faster rate when failure rates return to an acceptable level. In this way any congestion or broken routing paths are handled on the Internet. Thus congestion cost, and any associated externality such as social cost, is a time and not a monetary penalty to avoid congestion.

However, reliability based on packet re-transmission will generate delays which can not be tolerated by VoIP or video streaming which requires strict tolerances for latency and jitter. Thus, the question can be asked:

At what point do users or classes of users prefer to pay to avoid delay?

This depends on users' "willingness to pay" for the network services. This raises issues of traffic priority, and whether some classes of users are willing to pay for priority, in other words do not wish to wait. Simple analogies are the postal service where yellow express delivery post boxes sit side-by-side with red normal delivery post boxes – some users will pay the extra for express post, and public highways where transit lanes are provided if vehicles utilise more seating capacity.

Debate has now surfaced that a "two-tier" Internet is emerging, and the principle of "net neutrality" (all traffic treated with equal priority) will be no longer supportable if the Internet is to perform: PSTN equivalent voice calling using VoIP, streaming video, and product distribution of network goods, the latter supporting file sharing using Peer-to-Peer (P2P) protocols as exemplified by applications such as Bit-torrent.

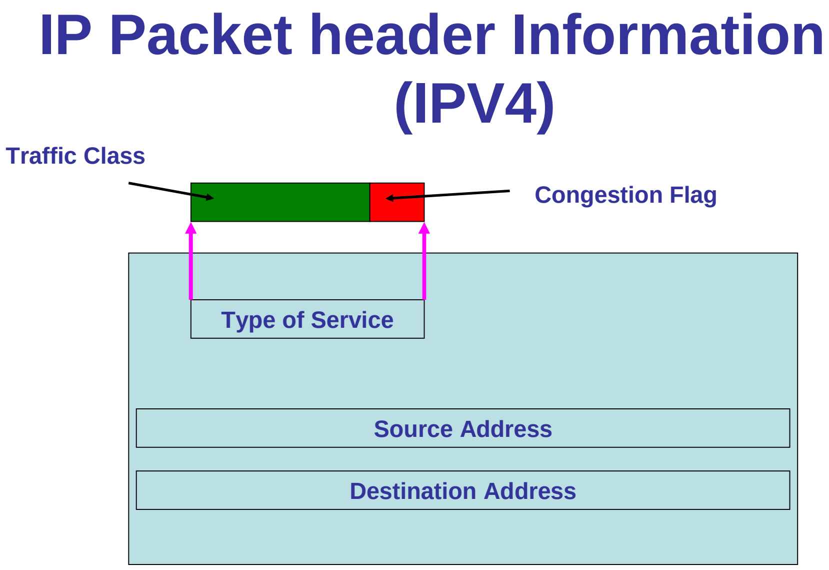

Figure 3: Selected IP header fields

The figure above shows that each IP packet[10] contains a source IP address and a destination IP address which identifies the device where the packet originated and the device of where it is going to. Note there is a Type of Service Byte (TOS) which allows QoS schemes to be developed, a typical scheme being:

- Six bits given over to class of service, an example of which is the Differential Service (Diffserv) protocol[11]; and
- the remaining 2 bits to for a congestion flag known as Explicit Congestion Notification (ECN)[12].

The diffserv protocol allows each router to put packets in separately served queues. For example, voice packets (VoIP) and video streaming would go into higher priority queues, web-browsing may go in a medium priority queue, and e-mail in the lowest priority queue.

Note that the source and destination port numbers referred to above with regard to packet filtering or traffic shaping, actually reside in the TCP or UDP header and not the IP header. While the IP addresses in the IP header represent the source and destination devices the port number (a 16 bit field) allows the device to find the requisite application software since one device can run many applications. Thus the Internet uses well known port numbers for systems and applications purposes, and by filtering on IP address / port pairs packet 'types', for example P2P traffic can be monitored, if an IP address / port pair is suspected of using inordinate amounts of bandwidth and violating acceptable or 'fair use' policies.

---

[10] There are other fields in the IP header but they are not discussed here.

[11] See IETF RFC 3246 which specifies expedited forwarding which guarantees equivalent service to a dedicated circuit switch, and RFC 2597 which specifies assured forwarding, which relies on priority levels.

[12] See IETF RFC 3168.

A priority queuing mechanism such as Diffserv can only guarantees relative priority within a router, and hence can only determine the order in which packets proceed to the next hop. To guarantee QoS from end-to-end across an AS, or AS group called an AS path, then a protocol such as MPLS can be used to create a path along the route[13]. Note that an MPLS router chain maintains its own list of label switched paths but honours the traffic priority using the service class protocol such as Diffserv in the IP header. In this way emergency calling can be put in a prioritised queue as it transits through each router in a reserved path. The MPLS path could be used for lesser priority packets but packets containing an emergency call must go straight to the head of the queue when they enter the router.

The ECN bits, supported in many but not all manufacturers' routers, allow for further robustness. In the event of extreme congestion, packets marked for ECN will be put through while those that are not are discarded by the router. Of course across the AS Path the constituent AS's must support MPLS enabled routers for end-to-end QoS to be maintained which will need to be reflected in the routing policy at the AS border gateway routers.

---

[13] In some ways this is a return to circuit being created although it is a 'soft' circuit and MPLS paths can be re-configured if circumstances demand.

Part 2: Economic Aspects of next generation networks

Why Price for Congestion in a Next Generation Network?

In a PSTN network the issue of congestion is ‘black or white’. If the PSTN signalling system can not set-up an end-to-end connection across the PSTN switch path then the call will fail before it starts. If the call is set-up then a dedicated circuit exists and congestion will not be an issue for the call duration. This contrasts with packet forwarding and switching where during the course of the packet flow congestion could occur at any point, and any time in the layered connection chain.

Of course, it could be argued that next generation networks should be designed with excess capacity so that congestion is not an issue, but this would be at the cost of overall efficiency.

Economic analysis can attempt to analyse whether capacity must be distributed by a more dynamic pricing mechanism than the current ‘best efforts’ flat rate pricing. It is of interest to ask if dynamic pricing reflects a socially optimal congestion cost, and / or price discrimination (to capture consumer surplus) by service providers. The goal of congestion pricing is to achieve the social optimum of network usage, so that a user only uses the network to the point where the marginal benefit derived from that usage is equal to the social marginal cost.

If consumer behaviour is such that users take into account the cost to them only - Marginal Private Cost (MPC), then this may create congestion through lack of acknowledgement, or arguably knowledge, of the burden their use creates on other network users - Marginal Social Cost (MSC). This problem is illustrated in the figure below by the distance between the MPC and MSC curves. This leads to sub-optimal levels of use of the network and is the origin of congestion issues.

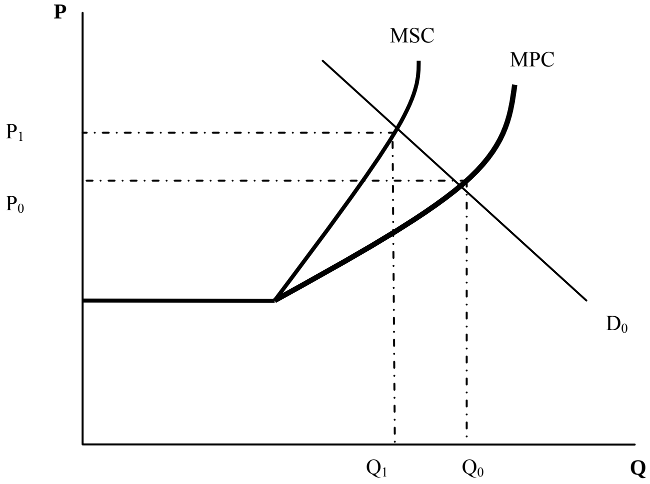

Figure 4: Negative Externality of Network Congestion

Given the point where the network is congested, the demand curve D_0 shows the difference between price and quantity levels between the MPC and the MSC at the same level of demand, and ideally should be set to equal MSC, not where users maximise individual utility at MPC.

As discussed above the Internet copes with traffic congestion by re-transmitting lost packets, and / or slowing the rate of re-transmission. However, it is unlikely that a high proportion of users are, as yet, accessing high bandwidth demand services such as film or television streaming [2]. It is also unknown as to the distribution of demand over time, whether there will be concentrated bursts, or peak demand periods or a more stochastic distribution.

However, the "early signs" of a growth in the area of streaming media and media downloads are there, as the "busy hour" for the Internet is not necessarily during the day as determined by business activity, but occurs early to late evening as volumes grow for P2P applications which exploit file sharing across a community of users, to distribute video and audio content[14]. This is likely to be a key change in the nature of demand from the current internet to the 'next generation' as the availability of this type of service increases along with the technology to adequately facilitate it.

However, as next generation networks are 'bulked-up' to offer 'on-net' proprietary video streaming content, then the majority of new video serving and router capacity may be for private distribution only, and will stay on-net, while access to non-proprietary web-sites, e-mail and VoIP will need to go 'off-net' to the remainder of the Internet[15].

---

[14] Sourced from a discussion with a telecommunication carrier representative.

[15] One can only speculate as to the mix of off-net and on-net content, for example new release movies and rights to events are likely to be proprietary on-net while off-net content may proliferate with the popularity of 'self-generated content' websites such as http://www.youtube.com/.

If congestion issues do arise then should the pricing of network services of a 'next generation' network reflect congestion levels in real time and charge users according to their individual demand, and its affect on other's demand? In answer to this question, forcing users to reveal their true willingness to pay could alleviate congestion and lead to a more socially optimal outcome. While the idea of any element of price discrimination in an Industry based heavily on high levels of infrastructure is often seen as monopolistic behaviour, often characterised by high levels of profit, it may be fair to argue that a degree of Price discrimination may not only be efficient but also Socially Optimal when correctly applied to the market situation.

The figure below shows the case of first degree price discrimination where the provider can identify the 'willingness to pay' of each user and charge each user a different price. This earns higher profits than charging the monopoly profit maximising price at P_4 marginal cost = marginal revenue because users who are willing to pay above P_4 do not retain their consumer surplus. However, users paying P_5 and P_6 who were previously excluded from the market are now able to enter.

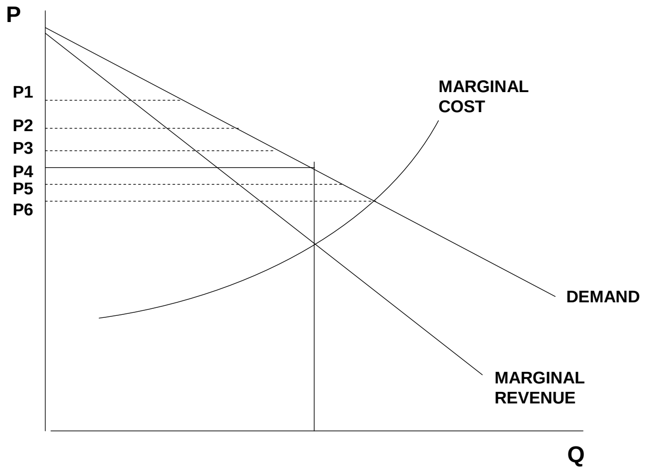

Figure 5: First degree price discrimination

In anticipation of high speed 'next generation' networks with the potential of integrated services such as streaming video and audio, a pricing system which responds to user demand at any point in time may more efficiently price bandwidth if the network is congested. With the Internet as a distribution platform for service provision broadening in scope, it could become very difficult to price a uniform[16]' flat-rate encompassing the true costs of the average consumer's usage and the cost of supplying the service without 'hidden' cross-subsidies between users.

---

[16] Pricing an internet service whereby the scope for service delivery ranges from simple low bandwidth applications such as email to high bandwidth applications such as streaming video would be difficult. Given the nature of the costs involved, the supplier in the absence of a dynamic pricing system could only charge a flat rate. Therefore the average price would not necessarily reflect a consumer's heaviness of usage.

The limits of flat rate pricing mechanisms can be seen initially in its inefficiencies in providing the market with a true level of demand, and as a new user establishes a network connection, they impose a burden on other users of the network. As the network approaches a congestion state, the marginal social cost that the presence of each subsequent user imposes on the network increases. It is at this point that the inherent problems with the current flat-rate pricing structure and best efforts network services may begin to manifest. In addition, the marginal social cost is not signalled back to the user via prices distorting the relationship between supply and demand.

Service Level Agreements

In terms of 'best efforts' service provision, the negative externalities arising from the presence of a user impose a burden on the QoS and transit delays of other user's packets [3]. Of course, in designing and implementing a radical change in pricing in the face of congestion for users the price elasticity and the dependant diffusion of the next generation of network services must be considered. If a congestion pricing scheme deters users from consumption, its costs may far outweigh its benefits, and in this case flat-rate pricing may be more 'consumer friendly'.

Given the origins of congestion, it probably can be assumed that users do not perceive their use of internet effects others or other network services inferring a perceived marginal usage cost of zero. As such installing usage based pricing from the consumer perspective, creates disequilibrium by identifying when Marginal costs are non-zero. If the user is aware of these extra usage costs will this effect the diffusion of next generation higher capacity services?

Historically, the movement from dial-up to broadband in Australia has been associated with lower price points, and thus can this history be used as a guide to the diffusion of higher speed broadband services as provided by ADSL2/2+ technology? The answer to this invariably lies in the price elasticity of demand, and given that broadband demand in the past has been price elastic[17], then a usage-sensitive pricing system may discourage the uptake by users. However, suppliers unable to apply first degree price discrimination can apply second degree price discrimination by offering pricing packages with varying prices dependant on purchased bandwidth volumes. If next generation networks exhibit economies of scale or at least average fixed costs that decrease more than average variable costs are rising so that average and marginal costs are declining over a range of bandwidth supply, then by expanding output second degree price discrimination can take the form of lower prices representing discounts for increased volume increasing both consumer welfare and supplier profits.

In cases such as this Service Level Agreements (SLAs) may be employed to provide a degree of flexibility to partially address the issue of congestion. The benefits of SLAs lay in the basis of a user self identifying their type – based on their own knowledge of degree of usage and types of usage. In this way, the problems of unknown user types and levels of demand are partially solved, at least for larger users such as wholesale

---

[17] For example the $29.95 'flat price', albeit with limited usage quota, introduced in early 2004 for an introductory broadband service, with low bandwidth usage quota, greatly accelerated broadband demand in Australia.

bandwidth buyers and corporations. It also allows both the supplier and user greater awareness of cost. On the side of the user, knowledge of the price of the service almost in total is of high value and is more likely to be well received. On the side of the supplier ex-ante investment certainty is provided for in the form of long term bandwidth supply contracts.

Unlike other forms of telecommunications, demand for bandwidth fluctuates periodically and therefore so does the spare capacity of the network. This stochastic demand trend often called "Bursty" demand [4], could be assumed of the new network in order to assist in the development of suitable pricing structures.

During periods of low demand or no congestion, the Marginal cost of sending a packet through the network is near zero. In this state there is no need for a dynamic pricing system and a flat rate or fixed scheme achieves efficiency in this situation. In situations where the Market approaches congestion through demand exceeding supply a dynamic or reformed pricing structure is required. This is the case for a number of reasons. Firstly, the peaks in demand currently present as a result of non-discriminate pricing will be somewhat smoothed out according to a users willingness to pay or priority of packet transit. Secondly it will eliminate the *'tragedy of the commons'* [5] of over-use of a common good or resource(s).' And thirdly it will provide market signals to the network suppliers to improve infrastructure and services as a solution to the problem in the longer term.

Although it is unknown as to what pricing mechanism to utilise in order to achieve efficiency of network use, there have been many suggestions, allowing for the absence of technological restrictions and implementation costs, to achieve network efficiency and social optimality. Any such pricing scheme whether static in nature as SLAs or dynamic as per 'Smart-markets', or dynamically smoothed such as Edge pricing (see below) should fulfil the following:

- Demand peaks are "smoothed" as those with lower willingness to pay are forced to shift their usage to periods of lower marginal cost, or no congestion;
- Signals, or incentives to invest in providing greater capacity should be provided to the network owner or service provider; and
- Low-priority applications, such as email should still be available at low cost to users with low valuation of these services.

Providing the model fits these guidelines, allocative efficiency within a given time period at least will be satisfied. In addition to a dynamic priority usage price for packets, fixed and recurring costs must still be recovered as a component of network access price. Although a situation of market power may lead to a sub-optimal outcome, it is assumed in the analysis here that this is not a monopoly situation, although it should also probably be assumed that it is not a perfect market situation either. As capital costs are known ahead of time, such a fee could remain a fixed amount per user and have no impact on the efficiency of any usage pricing structure.

In the short term, during periods of network congestion, users with a higher willingness to pay will remain in the market for bandwidth, whereas those with values relatively lower will drop out and wait until a period where the network has spare capacity and hence price is within their range. In summary, a user should remain using the system until the point in time where their marginal benefit is equal to the marginal cost of their presence on the network.

With a flat rate pricing regime, smaller users may actually be subsidising larger users through average price, not actual usage price. This occurs mainly due to the inability to effectively 'choose' or affect the price small users pay for packet transit via altering their behaviour according to their personal valuation of that service at a point in time. By directly targeting those who have a high willingness to pay for high priority network services, the cost is being disaggregated and redistributed according to the actual demand preferences and priority for each consumer, or consumer '*type*'.

An interesting characteristic of an individual consumer's demand is that it can vary over time and also vary for different uses of the network. This leads to further problems in designing either a SLA scheme, or any other static pricing scheme for that matter. In this sense, many authors [4, 7, 8] have discussed the ideas of using auction systems to allocate bandwidth, whereby individual demand can vary and therefore the Sum of the Marginal Social Cost can be calculated in "real time"[18].

Smart Markets

Provisions for real-time demand management using an auctioning system such as the "Smart Market [6]," seem in theory, to be an efficient method of pricing congestion to bring Demand and Supply into equilibrium based on adjusting QoS for network services. In the absence of technology restraints, the smart market achieves economic efficiency by both providing the supplier with a Marginal Social Cost figure and allowing users to express their true willingness to pay for use of the network and inadvertently the avoidance of congestion based packet delay.

In the real world however, the viability of such a model is decreased due to the lack of technical and software support, high cost billing and accounting overheads and the implementation of a suitable time scale for dynamic pricing which may not truly reflect dynamic demand due to long durations between price updates [7]. In this way the authors consider it a precursor to further advancements in congestion pricing through auction, and an attainable social optimum.

The Smart Market, developed by Mackie-Mason and Varian [6] revolves around an auction system in order to address congestion. From an economic perspective looking at congestible pricing systems, the Smart Market arrangement would not only alleviate congestion at times when Marginal Cost is a non-zero value, but also provide benefits to both the incumbent and the consumer to varying degrees. The proposed model revolves around auctions being held at congested routers. The system works

---

[18] "Real Time" refers to a system whereby the market is cleared and price re-estimated at a given time-scale, not truly "real time." Billing and Accounting costs and system and software support will determine the time scale in which the market is cleared.

whereby users submit a bid for transit of their packets through that router. Basically, the idea consists of packets being prioritised according to the price the users are willing to pay to have that packet sent, where the users who submit a bid above the marginal users bid are admitted access and pay the social opportunity cost. Those whose values are lower than the market clearing price are either dropped or buffered. As supply of capacity in the short run can be said to be near constant, any shift in demand will be reflected in the price. At the level that demand is equal to supply and equilibrium is observed, the price will be that of the Marginal Users bid, not of the bids submitted.

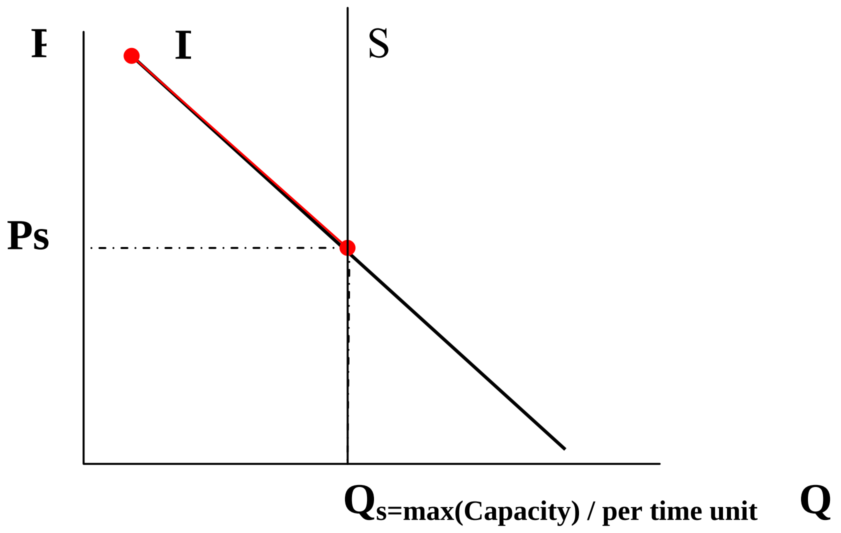

Figure 6: The market for Network Bandwidth

Assuming that bandwidth supply in the immediate future is constant due to technology rollout restraints, the figure above illustrates, albeit simply, the mechanics of the Smart Market and the way in which it allocates bandwidth. The price P_s is the price that all admitted packets are charged at, where capacity is exhausted at Q_s.

As opposed to other pricing options where the tendency of a user is to submit a bid lower than their true value, the model shows that the incentive for the user is to submit a true valuation of their willingness to pay because they do not necessarily pay the price they bid. Rather the mechanism ensures that the market clears at the Marginal User's bid price – that is the price of the bid of the last admitted user for a particular packet. Aside from being economically efficient by restoring market equilibrium via prices that bring demand and supply into equilibrium, the system also maintains benefits for users; with all users above the Marginal User reaping a consumer surplus of their bid minus the Marginal User's bid or the market clearing price.

On the supply side, the Smart Market system provides the supplier with better information which can lead to a more informed and less risky investment decision. Ultimately, this may lead to more being spent on increasing the total capacity of the network which in the long run will decrease the price a majority of users will face on average. A decision by supplier(s) to invest, *ceteris parabis,* will be made based on the total bid value of the packets not accepted, or those below the Marginal Users Packet Bid value. The service provider will then compare this value to the cost of increasing total capacity to accommodate those packets. Simplistically, if the revealed total value of those packets exceeds the costs, then it would be worthwhile for the incumbent to provide that capacity. Of course, this is an ideal outcome. In the case of a monopoly firm or a firm with significant market power, for reasons of profit maximisation, a firm may choose to constrain the supply at that level in order to force higher packet prices. In a secondary market or competitive market however where access may need to be granted to essential facilities[19], this is unlikely to be the case, as it would mitigate such market power.

To avoid inefficiencies in dealing with the problems of a congestion externality, a flat-rate fee or reserve price during times of no congestion is optimal. The proviso here is that during times of congestion, the prices reflect this. In essence a flat rate when Marginal Cost≈0, combined with a pareto-efficient and incentive compatible dynamic system when Marginal Cost >0 would be efficient.

Edge Pricing and the Dynamic Capacity Contracting (DCC) model.

Instead of pricing congestion per router or only at a border gateway router[20] when a packet needs to go off-net, as a fully dynamic auction system suggests, edge pricing estimates the cost determined by expected congestion of a packet's expected and best path to reach its destination. Expected congestion could be calculated, depending on variables such as congestion history, time of day [5] or some other type of 'peak load' identification.

Related to the idea of an average or estimated congestion price across the expected path is the idea of Dynamic Capacity Contracting (DCC) [21][8]. Edge Pricing contains essentially only one variable, the state of the network. Of course, this does display some flexibility of price given the variables that determine price. However, the above problem of 'bursty' demand could render the congestion cost inaccurate and lead again to higher than desired levels of congestion.

The DCC model takes the idea of Edge Pricing and expected congestion and introduces a time-scale variable to allow for more flexibility of price in order to enable it to better reflect the Marginal Social Cost of congestion. In essence, the DCC model allows allocation to be more efficient. As discussed with the Smart Market, the pricing "updates" or when prices are reset and determined again through whatever mechanism, are vital to deciding the actual state of congestion and in turn the efficient market clearing price. The introduction of Short Term contracts - possibly separated into Class of Service, taking the form of essentially two parameters: volume and

---

[19] Broadly defined as facilities that are uneconomic to duplicate.

[20] It would not be unreasonable to assume that an auction process at a single router could be extended to an Autonomous System where border gateway protocols negotiate the transaction. However, for time critical streaming bandwidth such auctions would need to be held at session set-up time, and so bandwidth reservation would need to be made over the life of the packet stream transit, and hence edge pricing would be more suited than the smart market approach.

[21] The analysis within [8] discusses a Point to Point scenario only.

contract length (time), is the mechanism by which it achieves greater efficiency over a basic Edge Pricing or expected capacity model. Logically, the closer to zero the time scale gets, the more a dynamic pricing model will need to trade-off between allocative efficiency, and monitoring and transaction costs. After the contract ends, the agreed time limit and the congestion cost will again be calculated to reflect the current state of the network. It is through this process that the DCC could be considered more or less a dynamic pricing system. To put it simply, the DCC is a middle solution to an Expected Capacity Model, encapsulated somewhat in SLAs, and an Auction system such as the Smart Market.

Conclusion

If efficiency of allocation of network bandwidth is to prevail, the system which determines the price a user pays should contain both of the following elements:

- A fixed component that covers network access and fixed costs incurred by the incumbent for supplying the service; and
- A dynamic pricing scheme for usage of the service in the presence of congestion.

Although a dynamic system is the “utopian” pricing scheme from the viewpoint of an economist with the possibility of maximising efficiency, there are other factors to take into account which may hinder the application of a truly dynamic system. Although many of them are technical, such as the lack of software and appropriate systems support for implementation of such pricing schemes, many of the problems still originate from the historical and cultural elements of network usage by users. From this perspective, fully dynamic congestion pricing via auctions may not be able to be successfully implemented.

Thus, it could be that concessions need to be made on Economic efficiency in order to accommodate user preference and technical support. It is probably evident that most dynamic pricing systems are suited to a wholesale and medium to large business market such as the National Electricity market, rather than a household or small business market. In this case larger wholesale providers may contract for variable priced bandwidth capacity but undertake to provide fixed prices to the end-user base. If such providers do not rely totally on ‘second best’ bandwidth volume quotas to their end-user base, they may need to price risk using hedging strategies in derivative or futures markets for bandwidth. However what markets for bandwidth will emerge is unknown until the intricacies of ‘next generation’ competitive access regimes are ironed out, and congestion pricing remains just a theory hoping to ensure efficiency in allocation by changing the way pricing works in the future.

Finally, if these decisions are made at the borders of the AS whether at session set-up time or during packet transit when packets need to go ‘off-net’, unifying standards for QoS will need to be defined requiring different pricing mechanisms than what currently suffices for SLAs and transit and peering agreements.

References

[1] White R., Slice D., and Retana A, *Optimal Routing Design,* Cisco Press, USA< 2005.

[2] Dolan., T., *Internet Pricing Is the end of the World Wide Wait in view?.,* Communications and Strategies., No 37, 1st Quarter 2000., p15.

[3] Nguyen., T., Armitage., G., *Pricing the Internet – A visual 3-Dimensional Evaluation Model.,* Swinburne University of Technology., Melbourne., AU., 2003

[4] Mackie-Mason, J., Varian, H., *Some Economics of the Internet*., University of Michigan USA, 1992

[5] Henderson., T., Crowcroft., J., Bhatti., S., *Congestion Pricing Paying Your Way in Communication Networks.,* University College., London., October 2001

[6] Mackie-Mason, J., Varian, H., *Pricing Congestible Network Resources*., University of Michigan, USA, 1994

[7] Singh., R., Yuksel., M., Kaltanaraman., S., Ravichandran., T., *A comparative Evaluation of Internet Pricing Models: Smart Markets and Dynamic Capacity Contracting.,* Renselaer Polytechnic Institute, Troy, NY

[8] Sinha., R., Kalyanaraman., S., Ravichandran., T., *Dynamic Capacity Contraction: a framework for Congestion Sensitive Pricing.,* Renselaer Polytechnic Institute, Troy, NY

Acronyms

10GiGE: Switched Ethernet protocol with a capacity of 10 Gigabits per second

AS: Autonomous System

ATM: Asynchronous Transfer Mode

Diffserv: Differential Service Protocol

ECN: Explicit Congestion Notification

ENUM: tElephone NUmber Mapping

E.164: an ITU-T recommendation which defines the international public telecommunication numbering plan used in the PSTN. It also defines the format of telephone numbers. E.164 numbers can have a maximum of 15 digits and are usually written with a + prefix.

MPLS: Multi-Protocol Label Switching

P2P: Peer to Peer

QoS: quality of Service

RTP: Real-time Transfer protocol

SDH: Synchronous Data Hierarchy

SS7: Signal System 7

TCP: transaction Control Protocol

UDP: User Datagram Protocol

URI: Uniform Resource Identifier

URL: Uniform Resource Locator

VLAN: Virtual Local Area Network

VoIP: Voice over IP

WDM: Wavelength Division Multiplexing