

and managed services are provided in parallel over a single broadband access connection. The quality of the managed services is guaranteed by, among other things, the use of bandwidth reservations in the access network. For instance, many Digital Subscriber Line (DSL) and fiber network operators use Ethernet Virtual LANs (VLANs) to partition and assign bandwidth to services. As a result, the IP traffic flows belonging to the managed services are separated from one another, and from the flows to and from the public Internet.

An important observation here is that in the triple play model described above, the guaranteed-quality connections that are created through the bandwidth reservations are typically only used for the network operator's own telephony and TV services. But guaranteed-quality connections could add value to important applications in other sectors as well, such as health, energy and security [11]. There is clearly an issue when this value is not realized, as there are large expectations for the contribution of broadband networks and services to solutions addressing the previously mentioned large societal challenges. Thus, the question arises as to whether it is useful to extend the use of guaranteed-quality connections to, for example, health and education services, which are currently delivered via either dedicated service-specific infrastructures or the best-effort public Internet. It is assumed that the broadband connection offers enough bandwidth to accommodate the additional guaranteed-quality connections without substantially reducing the bandwidth available for public Internet access. Otherwise, the deterioration of the public Internet access is bound to raise net-neutrality-related concerns [12].

Demand-side requirements from non-triple play services

We have interviewed five Dutch SPs from non-telecom sectors about the services they plan to provide in the future using broadband networks. The goal was to gain better understanding of the demand for guaranteed-quality connections for non-triple play services. Through open questions, the providers were asked about existing services that are provided over the Public Switched Telephone Network today, such as a medical alert service, burglar and fire alarms, as well as services that they plan to provide in the future, such as remote examinations in education. They then assessed whether or not the current broadband networks meet the requirements for the delivery of these services. Table 1 shows their responses mapped to four requirement areas: bandwidth, quality, availability and security. The SPs were not obliged to cover all requirement areas in their answers.

Table 1 shows that both the downstream and upstream bandwidth available in Dutch broadband networks is not perceived as a limiting factor. This should be seen in the context of the Netherlands being one of the countries with the largest internet penetrations and average access speeds in the world, and may not be applicable to other countries. Other requirement areas show clear limitations, though:

- **Quality:** the quality of two-way voice and video is seen as insufficient for services in which close and uninterrupted communication is required.
- **Availability:** the down-time of connections because of network problems is seen as too large for services that are important in emergency situations.

Sector	Service	Bandwidth	Quality	Availability	Security
Health	Medical alert service: alert emergency response center by pushing button	✓	✓	✗	
	Two-way video-based smart living services: extension of medical alert service with richer communication	✓	✗	✗	
	Medical monitoring: remote monitoring of weight, blood pressure, body temperature	✓	✓	✓	?
Energy	Smart meters: remote measuring of energy usage	✓		✓	✗
Security	Burglar and fire alarms	✓		✗	
Education	Virtual class: remote participation in classes through two-way video	✓	✗		
	Remote examinations: take official exams through two-way video	✓	✗		

Table 1. How service providers assess the suitability of the current best-effort Internet for their services.

- Security: there are doubts as to whether public Internet connections are secure enough for services where protection and integrity of data are important.

The interviews thus point at other factors than bandwidth as inhibitors for new services. The security issues mentioned seem to be related to the application layer rather than the broadband infrastructure, and are therefore not further discussed in this paper. Quality and availability, on the other hand, can be directly related to properties of the network infrastructure. It is therefore worthwhile to investigate how broadband infrastructures can meet these requirements.

Broker/VPN concept and operational implementation
 Broker/VPN (Broker / Virtual Private Network) is a broadband network architecture developed and trialed by the Dutch network provider Reggefiber together with local communities and SPs. It is an implementation of the two-lane model, i.e. it provides guaranteed-quality connections for telephony and IPTV but also for other services that require more quality and reliability than offered by the public Internet. The first new service that uses the guaranteed-quality connections is a high-quality Neighborhood TV service in the Dutch town of Zeewolde (see Figure 1). Through this service, live broadcasts of local events such as town council meetings, sport matches and theatrical performances are distributed to local community members.

In the Broker/VPN, Ethernet VLANs are used to separate the traffic flows of the best-effort public Internet access and the various managed services. Each VLAN is terminated on a separate Ethernet port on the Network Termination (NT) in the Home Gateway (HG). Depending on the customer's subscription, the

total available bandwidth is typically 100 or 200 Mbit/s, with 10 or 30 Mbit/s assigned to the Neighborhood TV service in VLAN no. 4. The Neighborhood TV service registers video signals and distributes them live to neighborhood viewers. The Broker/VPN connection is used for feeding the live video signal upstream, from the camera over the Fiber-to-the-Home network towards the IPTV platform. In the IPTV platform, the video feed is incorporated in the commercial set of IPTV channels and distributed downstream to the viewers using standard IPTV delivery protocols.

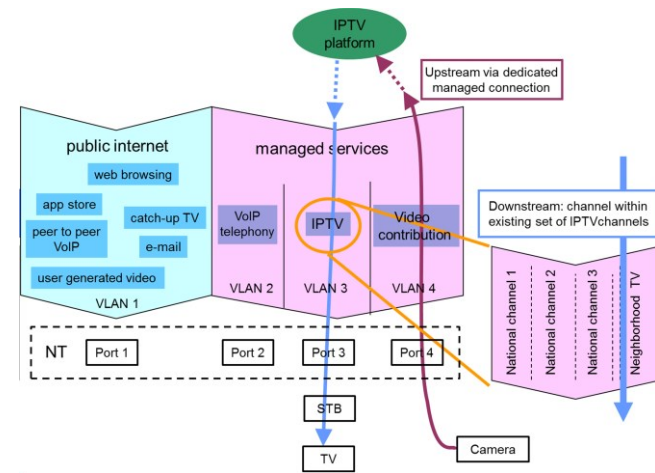


Figure 1: The delivery of managed IPTV services in Zeewolde, including a 3rd-party Neighborhood TV channel, using a Broker/VPN managed connection.

Viewers of the Broker/VPN-supported Neighborhood TV service perceive the video quality as better, and in particular more stable, than in earlier implementations in which the upstream contribution occurred unmanaged over the public Internet. Future research at the network and video layers should confirm this result,

Abbreviations

DSL	Digital Subscriber Line
DPI	Deep Packet Inspection
E2e	End-to-end
HG	Home Gateway
IP(TV)	Internet Protocol (TeleVision)
MSL	Managed Services Lanes
NT	Network Termination
OTT	Over The Top
QoS	Quality of Service
SP	Service Provider
SSID	Service Set Identifier
STB	Set-Top Box
UPnP	Universal Plug and Play
V(LAN)	(Virtual) Local Area

and relate it to the relevant technical performance indicators. Also much appreciated by viewers is the ease of use that the Neighborhood TV service offers. Live video streams can be contributed to the IPTV platform from every fiber connection equipped with Broker/VPN, just by hooking up the camera to the specific Ethernet port on the NT. This also works in public buildings, such as the theater and the town hall, where upstreaming the video over the public Internet was often problematic because of various configuration issues in the private networks. These configuration problems are avoided in the Broker/VPN approach: if connected to the correct port, the camera does not have to be (auto-) configured locally by the HG first.

The VLAN used to implement the guaranteed-quality connection extends from the NT to a Point-of-Presence within the same city. This local VLAN is independent from other VLANs defined in the larger, national fiber network. Other local services may therefore share the VLAN for upstream video contribution, as long as they are offered by the same local SP. This SP has knowledge of the local services, and is therefore in a position to perform connection admission control to preserve the quality of these shared connections.

The administration of local VLANs is a substantial management effort for the operator. An alternative implementation therefore defines the VLANs in the national network, rather than locally within the smart cities. Such VLANs are especially attractive for SPs with a national customer base. At the same time, national VLANs require more coordination among the relevant SPs regarding connection admission control. This may lead to undesired dependencies between SPs.

Smart City services support in the home

Heterogeneity of home networks

In the Zeewolde pilot, the Broker/VPN VLAN is terminated on a LAN port of the HG, and the camera must be connected to that port using a dedicated Ethernet cable. Only then, the bandwidth is reserved right down to the client device, and ease of use and end-to-end QoS can be guaranteed. This could also be done for other services using dedicated devices (e.g. burglar alarm systems, energy monitors, medical monitoring devices), using techniques like segmenting by IP address or using a different WiFi SSID. However, this configuration prevents the shared use of such devices by other services and devices in the home, especially if it concerns devices with more generic functionality, such as tablets. One can even argue that every device can in principle always be used for other services than the intentional one. Devices can only be shared if the home network is an integrated network. Today, however, most home networks are unmanaged and lack QoS support. We do not expect this to change in the near future, as home networks are becoming increasingly heterogeneous. In such a context, the service's reliability and performance depend strongly on the characteristics of the home network [7]. They include static properties such as the topology, device characteristics, and the network technologies used, but also more dynamic properties such as individual session times and the utilization of the various links. The latter in particular is an issue for providers of managed high-quality services such as IPTV and telephony [4], but also the services of Table 1.

The increasing heterogeneity has led to significant device- and service-management complexity [6]. As home networks are owned and controlled by the

Measurement setup

In every household the local home gateway was replaced with a Linksys WRT54 GL v.1.0 broadband router and a monitoring laptop computer acting as a traffic recorder. The traffic recorder runs PRTG (www.peasler.com), which measures the rate of all incoming and outgoing traffic on every single port of the router at 10 s intervals. We observed incoming IP flows from the public Internet into the Local Area Network (LAN); outgoing traffic from LAN to the public Internet; and LAN-to-LAN traffic (which stays within the home). The router has five 100 Mbit/s Ethernet ports and an IEEE 802.11g access point acting as a single port. One Ethernet port was connected to the WAN, and another to our IP traffic recorder. On the other four ports, we measured and recorded the incoming traffic rate (in bit/s) for a full week.

tenant, who is usually not an experienced network manager, it is often the helpdesk of the broadband access provider who gets called when a service malfunctions, even though he is not responsible for the home network. Such help desk calls are very expensive, because effective tools for remotely home network management hardly exist. Besides, home networks have also become increasingly personalized, meaning that (remote) troubleshooting and problem solving becomes more and more tailor-made work.

Home network monitoring and modeling

Forcing QoS in the home network with standard Diffserv-like techniques has proven to be impractical. These solutions need to be supported by every device in the end-to-end (e2e) path to be effective, which is a requirement that is hardly ever met in home networks. In the research project FIGARO (www.ict-figaro.eu), we take a different approach. We simply acknowledge the heterogeneity of the smart city network, especially it containing best-effort private networks such as home networks. Instead of trying to control the network across domains, we learn about all relevant network characteristics in real time and require the services and content to adapt accordingly. Service and content adaptation can be done in many ways, for instance by means of adaptive bitrate streaming, session admission control, or local caching. Although we realize that this is not an ideal solution, as it does not match the provided quality of the MSLs for 100%, it is still better than best-effort, and for many services that will be enough.

The first step in this process is the development of intelligent home network monitoring tools. Available bandwidth monitoring tools [5] typically provide snapshots of the current state of the home network. To

make decisions about content and service adaptation, also educated guesses should be made about the state of the home network in the near future. For that, models need to be developed describing the (statistical) home network behavior.

Standard Internet models cannot be adopted readily, because home networks have different characteristics. For example, one often assumes that the stochastic properties of traffic in the home network are the same as in the public Internet. However, there are many applications that run at home and congest the home network, but do not communicate with the public network. Examples are network printing, file transfer, content synchronization, internal control traffic (e.g. UPnP), and streaming media from a local storage device. To model these networks, we conducted a traffic measurement campaign in Amersfoort.

Results

So far, we have collected data in 15 selected households. We intend to triple this amount in the coming year, but our first results already look promising. The characteristics of the households are summarized in Table 2. Most participants can be classified as "early adopters", and we assume that prediction models resulting from these networks will be applicable to a much larger section of the total population in the near future. We also carried out a questionnaire-based survey to determine the Internet use profiles in our sample. The results are similar to what is known about the average Western population [3]. There are some minor differences, mainly in the use of audio/video services and terminal services (e.g. telnet), which indicated that our population is indeed significantly skewed towards early adopters.

	Type	#
House	Terrace house	9
	Apartments	5
	Suburb house	1
Network type	Access: Cable	2
	Access: DSL	13
	LAN: Ethernet+WiFi	14
	LAN: Ethernet + Wifi + Power line comm.	1
Av. # of networked devices per household	Laptop or PC	4.6
	Multimedia Servers	1.5
	Personal handheld devices	1.2
	Game Console	0.7
	Others	1.0
User	Families with children	12
	Families without children	3
	Average # of users per house	3
	#hours daily spent on networked applications	4

Table 2. Profile of the households and home networks that participated in our study.

Per household i and per unit of time (10 s) the incoming traffic is summed over the ports, yielding the aggregated home networking traffic $x(n)$ at time n . Figure 2a depicts the obtained probability mass function $pmf(x)$. Traffic rates higher than 40 kbit/s, up to 92 Mbit/s, have also been observed, but with a relative frequency lower than 0.002. Despite the large statistical noise, the home-network traffic-rate distribution differs significantly from that of the Internet, which generally follows a Gaussian function [13]. Our results are better approximated by a Generalized Pareto distribution, with a root mean square error of 0.0005. This can be explained by the fact that the Internet is always active, and statistical multiplexing can be applied, while home networks tend to have long periods of relatively low activity and support fewer applications. This result is remarkable, and important for smart cities. We expect that smart city networks will contain a disproportionately large amount of small private sub-networks like home networks, and must therefore most likely be modeled differently from the mainstream Internet.

We also studied the individual traces $x(t)$ for every i , and looked for sudden steps in $|x(t)|$ larger than 250 kbit/s. The time between two such steps we call the application session time d . $pmf(d)$ is depicted in Figure 2b. The resulting curve can again be approximated with a Generalized Pareto distribution function, and an average application session time of about 500 s. This value is surprisingly consistent with the average user session time reported by [1]. Apparently, content and service optimization systems in smart city networks can learn specifics about use patterns of best-effort private networks by monitoring

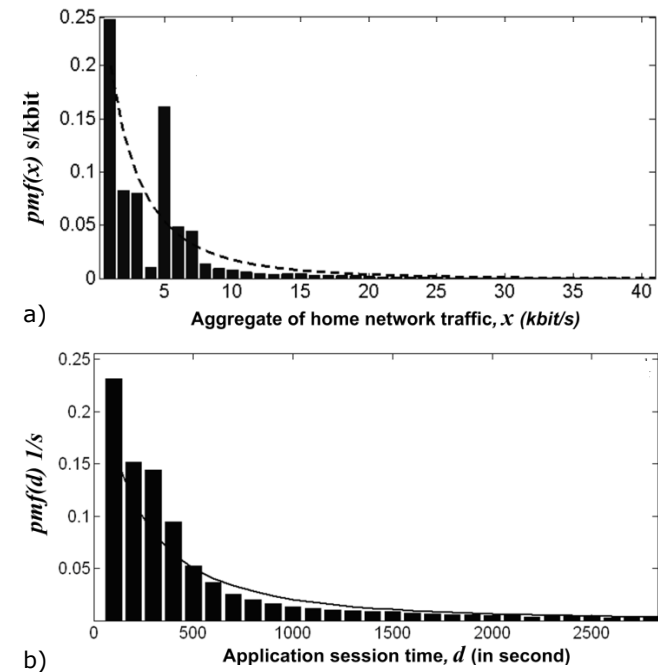


Figure 2: a) The probability mass function of our observed aggregated home-network traffic rates x , b) The probability mass function of the approximated application session time d .

traffic at the IP level, without first performing privacy-sensitive research by means of questionnaires or DPI.

Conclusions

Network heterogeneity inhibits service quality assurance considerably, especially when considering the high demands for robustness that new services require. Our survey results indicate that over-dimensioning the public network’s bandwidth is not the solution. The focus should rather be on improving the networks in a way that increases the quality and availability of the respective services.

One way to do so is by introducing MSLs. The first results of our experiments with MSLs are promising. Network-neutrality concerns can be avoided as long as the MSLs do not influence the openness and the bandwidth of the public Internet services negatively. The MSLs cannot be extended into the home networks, but are terminated on the home gateway. For quality assurances in the home network we therefore propose the use of advanced network monitoring tools in combination with service and content adaptation. Such tools do not yet exist widely, and require the use of dynamic home network traffic models. From measurements in various homes we conclude that the amount of home network traffic as well as the observed application session times follow Generalized Pareto distributions. As a consequence, smart city networks cannot be understood with standard Internet models.

Acknowledgements

We thank Patrick Beitsma, Jouko Huismans, Jan Wester, Alwin Sixma, Freek Bomhof, and Robin Purchase for valuable input. This work was partly supported by the European Commission under the 7th Framework Program, Grant Agreement 258378.

References

- [1] Balachandran, A., Voelker, G. M., Bahl, G. P., and Rangan, P. Characterizing user behavior and network performance in a public wireless LAN, *SIGMETRICS Perform. Eval. Rev.* 30, 1 (2002), 195–205.
- [2] Caragliu, A., Del Bo, C., and Nijkamp, P. Smart cities in Europe, *J. Urban Tech.* 18, 2 (2011), 65–82.
- [3] Cheeseman Day, J., Janus, A., and Davis, J. *Computer and Internet Use in the United States: 2003*, U.S. Department of Commerce, Economics and Statistics Administration, U.S. CENSUS BUREAU, 2005.

- [4] Delphinanto, A., Hillen, B.A.G., Passchier, I., van Schoonhoven, B.H.A., and den Hartog, F.T.H. Remote discovery and management of end-user devices in heterogeneous private networks, In *Proc. of IEEE CCNC 2009*, IEEE Press (2009).
- [5] Delphinanto, A., Koonen, A. M. J., and den Hartog, F. T. H. Real-time probing of available bandwidth in home networks, *IEEE Commun. Mag.* 49, 6 (2011), 134–140.
- [6] Develder, C. *et al.* Delivering Scalable Video with QoS to the Home, *Telecommun. Syst.* 49, 1 (2010), 129–148.
- [7] DiCioccio, L., Teixeira, R., and Rosenberg, C. Impact of home networks on end-to-end performance: controlled experiments, In *Proc. of HomeNets 2010 (ACM SIGCOMM workshop)*, ACM Press (2010).
- [8] ETP Expert Group, *Strengthening the role of European Technology Platforms in addressing Europe's Grand Societal Challenges*, EUR 24196 EN, Publications Office of the European Union, Luxembourg, 2010.
- [9] Loeb, H.P., Liss, C., Ruckert, U., Sauer, C. UMAC – A Universal MAC architecture for heterogeneous home networks, In *Proc. ICUMT '09*, IEEE Press (2009), 1–6.
- [10] Marcus, J.S., *et al.* Network Neutrality: Challenges and Responses in the EU and in the US, European Parliament, IP/A/IMCO/ST/2011-02, PE457.369, 2010.
- [11] Middleton, C. Delivering services over next generation broadband networks: Exploring devices, applications and networks, *Telecommunications Journal of Australia*, vol. 60, no.4, pp. 59.1 (2010).
- [12] Nooren, P., *et al.* "Net neutrality and the value chain for video," *info* 14, 6 (2012), 45–58.
- [13] Pras, A., van de Nieuwenhuis, L., van de Meent, R., and Mandjes, M. Dimensioning network links: A new look at equivalent bandwidth, *IEEE Network* 23, 2 (2009), 5–10.