Digital Advanced Rural Testbed for Next Generation Access

Paris Skoutaridis and Graham Peters

Avanti Communications,
74 Rivington Street, EC2A 3AY, London, UK
{paris.skoutaridis,graham.peters}@avantiplc.com

Abstract. This paper provides a first look at the Digital Advanced Rural Testbed (DART) project, which is funded by the Technology Strategy Board under the Network Services Demonstrators Programme. The objective of the project is the design and implementation of a network demonstrator incorporating advanced network infrastructure elements and service enablers. This is done with a view to allowing third parties to experiment with novel business models, applications and services, compliant with the vision of next generation access services.

Keywords: Bandwidth Variation, Content Caching, Demonstrator, Multicast, Next Generation Access (NGA), Quality of Service Variation, Satellite, Technology Strategy Board (TSB), Testbed.

1 Introduction

Over the past couple of decades the internet has played an increasingly integral role to our lives both at a personal and a business level. The resulting increased demand for internet access is also accompanied by a shift to increasingly feature rich content and applications with video content currently being the most bandwidth demanding. The trend for more feature rich, bandwidth demanding content access over the internet has thus far evolved with increasing pace, and is expected to continue to do so. Consequently, Next Generation Access (NGA) services are receiving significant interest in an effort to develop solutions capable of catering for the ever increasing access requirements of both businesses and the public.

The DART project is funded by the UK Technology Strategy Board (TSB) [1] under the Network Services Demonstrator Programme [2]. The project aims to facilitate the experimentation and trialling of business models and capabilities deemed pivotal to the provision of NGA services to remote and rural communities. To that end, DART will offer a test-bed platform incorporating advanced infrastructure and technologies using a mix of satellite, terrestrial wireless and community fibre schemes. The test-bed will be built around a specific set of enabling technologies: real time (or near real time) bandwidth management, quality of service (QoS) management, multicast and content caching. While bandwidth and QoS management are mostly relevant to direct-to-premises end users, the DART network will enable the

G. Giambene and C. Sacchi (Eds.): PSATS 2011, LNICST 71, pp. 131–142, 2011.

trialling of multicasting and content caching to both direct-to-premises end users and remote community central network points (central caches) from where the content can then be redistributed to individual end user's premises.

The DART team comprises five organisations well placed for the design and implementation of the demonstrator, as well as the operation of the demonstrator and provision of support during the experimentation phase of the project. The team is led by Avanti Communications who will make bandwidth available over the recently launched HYLAS1 satellite, and be responsible for the design and implementation of the enablers over the satellite network. Metabroadcast, a London based design and technology company that develops content (video and audio) discovery products using content and social metadata, will utilize their existing software to provide a stream of recommendations for specific content to be cached. The University of Aberdeen will support the recruitment and experimentation with direct-to-premises end users and lead the dissemination activities. The University of Lancaster will provide access to the Wray community network thus enabling experimentation with localised content caching and multicasting. Finally, 21Media, a company specialising in the development of intelligent, platform independent software systems for IPTV, will focus on the exploitation of IPTV usage data over the Wray network and the development of a localized content distribution system.

The remainder of this paper discusses the motivation behind the design and implementation of the DART demonstrator, as well as presenting an overview of the demonstrator itself. The paper is organized as follows. In section 2 we discuss the concept of next generation access and the process that is required for its eventual realization. Section 3 discusses the Network Services Demonstrations Programme of the TSB and how that program aims to address the eventual roll out of NGA services in the UK through stimulating the development and testing of relevant business models and applications. Section 4 presents the DART project in more detail, providing an overview of the key network components and architecture, the specific demonstrator enablers and the way in which it is expected experimentation will be set up and undertaken by third parties. Finally, section 5 concludes the paper with the timeline of DART and the contact details for the project's contact points.

2 Next Generation Access

The ever increasing reliance of modern society on the internet and broadband access is currently putting a strain on existing networks. This problem is compounded by the shift towards more feature-rich, bandwidth demanding applications and services, as well as a shift in consumption patterns and the continuously increasing amount and variety of content available on the internet.

The increased reliance of modern society on broadband internet connections has also led to another phenomenon; that of the digital divide. The fibre networks predominantly used for delivery of broadband connections are largely unavailable in remote and rural areas, meaning broadband access to such locations is extremely limited if not completely unavailable. Consequently, there is a clear distinction between those with and those without access to high speed broadband connections.

Next Generation Access is the term used to collectively describe the services and applications that will be available over the internet in the future, as well as the networks and technologies that will be employed in their provision. As the current trends in services and applications indicate, next generation services and applications will be characterised by significantly increased bandwidth requirements as well as, in many cases, stringent QoS requirements. Given this, next generation networks will be required to cater for the increase in bandwidth requirements as well as to guarantee the provision of suitable QoS for those applications that require it.

At the time of writing next generation access in urban areas of the UK where broadband internet is already attainable is envisaged through the extension of the fibre-optic backbone connections with fibre-to-the-home, instead of the currently employed copper wire or cable connections. Such an upgrade in the last mile access method would significantly increase not only the bandwidth attainable over a broadband connection but also the reliability and quality of the resulting service.

The case is not the same for remote and rural areas where broadband connections are predominantly unavailable. The high cost associated with the deployment of fibre-optic connections removes the economic viability of providing such connections to rural and remote areas of a low population density. Consequently, there is a clear danger that remote and rural areas may be bypassed by future broadband access developments, thus falling even further behind in the broadband revolution compared to high population centres. This would further impact life in such areas both at a private and a business level, and would likely put increasing pressure on the already low populations. Alternative access methods as well as content and application provision paradigms are therefore required in order to ensure that next generation services do not bypass the rural UK.

DART proposes that next generation access to rural areas can only be realised through the combined use of a host of access technologies, in conjunction with intelligent methods of content distribution and application design. In accordance with this proposition, DART proposes the implementation and trialling of specific enabling technologies and network architecture characteristics that have the potential to significantly improve the available offering to the rural community both in terms of service level and quality, as well as in terms of the variety and quantity of the available content and applications.

3 The Network Services Demonstrators Programme

As part of the Digital Testbed Programme, the Network Services Demonstrators program of the UK Technology Strategy Board (TSB) aims at establishing test-bed platforms for the demonstration of network services.

It is the intention of the TSB that demonstrators developed under this programme become national hotspots for experimentation and trials. This should in-turn stimulate development and innovation in fields such as business models, applications and services that rely on advanced network infrastructure and service enablers. Thus, the

Network Services Demonstrators program aims to address the high, unmet demand for advanced platforms over which experimentation can be undertaken in an open and live manner.

The Network Services Demonstrators program comprises two phases that all funded projects are required to follow. Phase 1 involves the definition, design, and implementation of the demonstrator over a period of four to six months. This phase culminates with the generation of a proof of principle business model, application, or service that is enabled by the demonstrator's features.

Phase 2 of the program is the experimentation phase. During this phase the demonstrators will "go live" with access to third parties being provided. Depending on each demonstrator's features, third parties will then be able to experiment on relevant novel business models, applications, and services. The second phase of the Network Demonstrators Programme has a duration of at least one year.

4 The DART Demonstrator

4.1 Objectives and Motivation

The key requirement set by the Network Services Demonstrators Programme on all funded projects is the development of a demonstrator system that enables third parties to trial and experiment with novel and innovative business models, applications and services centred around advanced network infrastructure capabilities and service enablers.

To that end, DART will design, implement and make available for experimentation a demonstrator system that will allow third parties to experiment with specific network infrastructure capabilities and service enablers that have traditionally been either prohibitively expensive to access, or even unattainable altogether.

Through the DART project, experimenters will be provided with access to an advanced Ka band satellite network, thus bypassing the relatively high cost traditionally associated with bandwidth leasing. In addition, the enhanced nature of the design and capabilities of the HYLAS1 network and end user equipment will present experimenters with the opportunity to trial advanced capabilities and service enablers that have thus far been unavailable for testing over satellite networks. These are the bandwidth variation capability, the link QoS variation as well as the ability to experiment with multicasting and caching of content either locally at the end user premises or centrally at a community cache facility.

Through the supported service enablers, DART will be one of the first test beds to allow in depth experimentation with innovative services such as telehealth, telemedicine and precision farming techniques. Experimenters will also have the opportunity to test and develop new business models based on the observation of end user behaviour in rural areas as well as other demographics, or by developing innovative content distribution methods through exploitation of DART's multicast and caching capabilities. In all cases monitoring and feedback of activity over the DART network will be undertaken to allow both experimenters and the DART consortium to analyse the offered services.

Crucially, in addition to the ability to experiment with the above features, DART also presents the unique advantage that these capabilities and service enablers will remain available after the completion of the project for both further experimentation and development as well as for potential commercial exploitation.

4.2 The DART Enablers

Bandwidth Control. Dynamic occasional variation of the link bandwidth is a particularly useful feature that is not currently widely supported by broadband internet networks. Particularly in the case of broadband over satellite, where bandwidth is a premium resource, the ability to dynamically vary the link bandwidth on an occasional basis would not only be extremely useful to both private and business end users but also to network operators.

Dynamic bandwidth control would enable end users to reduce their satellite broadband cost by opting for a lower bandwidth connection for their every day use, while having the option to augment their bandwidth allocation at times that a higher speed connection is required. Examples of such cases include instances when an end user would need to download or upload a large file, attend a videoconference or access other bandwidth demanding content and applications. User-initiated bandwidth variation would require the end user to issue a request to the network, which would respond by increasing the non-contended bandwidth allocated to that link for a predetermined amount of time. In this model there can be several “service levels” defined with respect to the amount of bandwidth increase, each with a predefined cost per unit time.

Dynamic bandwidth variation could also prove attractive to content providers that require high bandwidth, high quality connections for their service. In this scenario a content provider could initiate the bandwidth augmentation of a specific link while their content is being forwarded to the end user. The cost for the temporary bandwidth increase could either be born by the content provider or passed on to the end user via the fee that the content provider charges for their services.

Further to the above, dynamic bandwidth variation would also be a particularly attractive proposition for network operators. Context aware bandwidth variation could be provided by the network operator through the inclusion of appropriate software components in the network management system. Irrespective of the initiation point, dynamic bandwidth variation can potentially increase revenue and reduce network congestion and contention ratios. It also affords the network operator with increased flexibility and adaptability in their bandwidth management processes. Novel and innovative service level agreements and charging methods could also be developed as a result of the ability to dynamically vary the bandwidth allocated to individual links.

Variable Quality of Service. QoS variation is another capability that is currently unavailable in today’s broadband internet services. Similarly to bandwidth variation, QoS variation is a capability that has the potential to not only significantly improve the service experience of end users but also to decongest satellite communications networks and improve the overall quality of the provided service. QoS variation is becoming increasingly important in internet networks as a result of the increasingly varied types of services and applications available for consumption, and the significant differences in the QoS requirements. Similarly to bandwidth variation, QoS variation may be provided in a context aware, user initiated or content provider initiated manner.

It is envisaged that users, content providers, and network operators alike would benefit from the dynamic use of QoS variation on individual links. Broadband connections could be allocated default QoS attributes for the majority of the time, guaranteeing acceptable performance during "typical" connection use, with on-demand variation of the connection's QoS attributes taking place when more demanding application or services are accessed by the user. As an example, a lower latency and lower jitter traffic class could temporarily be allocated to a user's connection that engage in a videoconferencing session, or a better protected (e.g. through FEC) traffic class could be allocated to a user that is downloading a large file in order to reduce the download time.

Multicast and Caching. Multicast data distribution is particularly well tailored for satellite communication systems as it takes advantage of the wide area of coverage typical to satellites. Specifically in the case of broadband over satellite multicast can have a significant impact on the bandwidth requirements associated with distribution of common data to a large number of receiver units and, by extension on the number of data distribution sessions that the satellite network can support simultaneously, as it eliminates repetitive transmissions for multiple receivers. Through moderating the bandwidth requirements for data distribution over satellite multicast also improves the scalability of the satellite network to both the receiver population size and the amount of content distributed.

Multicast is equally suitable to both streaming and background traffic classes. Particularly in the case of live content streaming, content providers would benefit significantly from the use of multicast as opposed to distribution of the content over several unicast links. Relevant examples of content would include academic lectures, medical operations, telemedicine where several non collocated experts are required to observe or live sporting events. Similarly, multicast can reduce the bandwidth requirements and overall costs associated with background traffic distribution to several receivers by utilising a single link instead of several unicast links, with the benefit attained being proportional to the content volume and rising sharply with the number of receiving terminals.

When coupled with content caching multicast proves even more beneficial with regards to bandwidth utilisation efficiency. Bandwidth demanding content such as video can be multicast to several caching servers simultaneously thereby significantly reducing the network resources utilisation. The cache servers can then provide end users with access to the cached content for as long as the content is stored. This method of data distribution provides the added advantage that the cached content can be made available to end users for a prolonged time period without the need for additional or prolonged multicast sessions. Furthermore, this is a particularly cost effective and efficient method for making bandwidth demanding content available to communities equipped with a single central cache facility.

4.3 Demonstrator Architecture

The architecture of the DART demonstrator test bed is depicted in figure 1. The DART test bed is composed of the Avanti core network, the recently launched (November 2010) HYLAS1 satellite and two test site categories. The first test site category comprises direct-to-premises users equipped with enhanced two way satellite communication equipment and located in Scotland. The second test site category is a hybrid satellite and cable/wireless mesh network in and around the village of Wray, located in northern England.

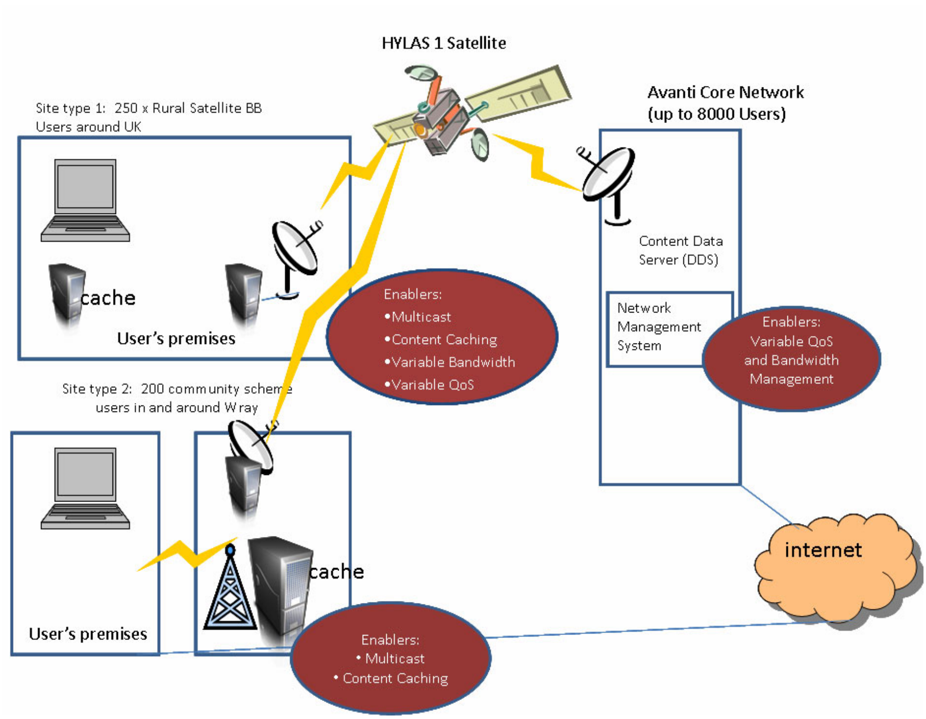

Fig. 1. DART Test Bed Architecture

The end users comprising the first test site category will be primarily recruited from the current Avanti pool of users as well as from the recruitment efforts of the University of Aberdeen. It is estimated that at least 250 existing Avanti end users will participate in testing using DART. The direct-to-premises test site will support trialling and testing of all DART enablers. Multicast and caching services will be tested using an adapted version of the software and hardware developed for an ongoing European Space Agency (ESA) funded project called NXY, whose development schedule is ahead of that for DART. In this scenario, content will be multicast and cached locally on end user equipment and will be accessible by the end users transparently through their browser. A possible realisation of the network components and architecture for multicast and caching is illustrated in figure 2.

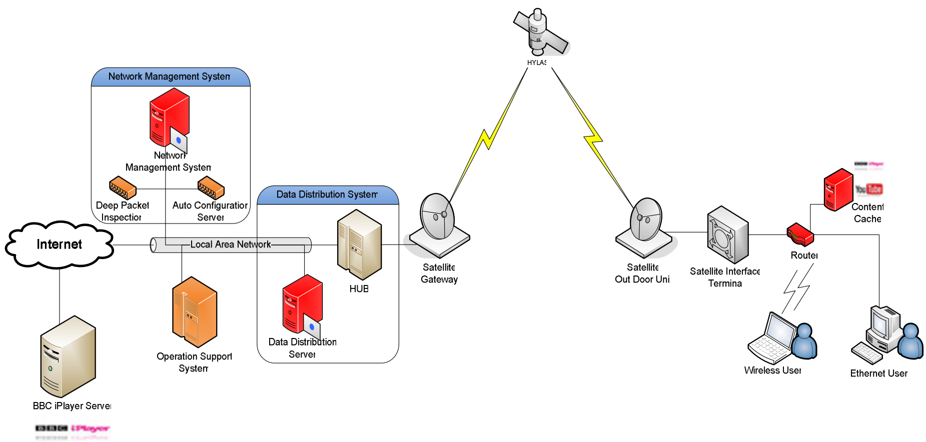

Fig. 2. Multicast and Caching Implementation Architecture

Variable QoS and variable bandwidth testing will also be supported by the registered direct-to-premises end user equipment. Testing of these enablers will be facilitated by the development and distribution to end users of software that enables them to issue bandwidth augmentation requests or changes to the QoS attributes of their links on a short or long term basis. Suitable software for the reception and processing of such requests, as well as the tracking and (possible) billing for the resulting service on the network management side will also be developed and integrated with the existing Avanti network management system.

In addition to the direct-to-premises portion of the network, DART will also provide access to a second site category. This will be the mesh wireless network installed at Wray in 2004 through the National Rural Support Programme (NRSP) [3] and operated since by the University of Lancaster. The Wray network [4] is comprised of a number of mesh wireless routers that connect the village's houses to a central mesh router that provides internet access. Figure 3 illustrates the current architecture of the Wray network. For the purposes of DART, a community satellite caching facility will be added to the network and interfaced with the central mesh router. This facility will allow the caching of content that the community members will then be able to access over the mesh network. The Wray network, augmented with the community satellite caching facility, therefore will provide an excellent opportunity for experimentation and trialling of central caching and redistribution of bandwidth demanding content and the effects that this content delivery method would have on both the backhaul and the last mile networks.

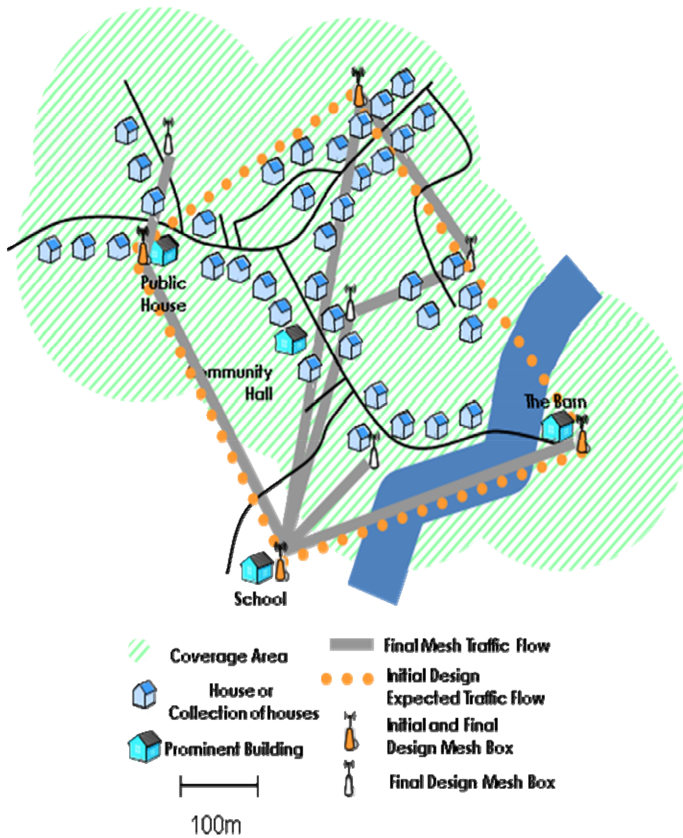

Fig. 3. Wray Network Architecture

4.4 Accessing the TestBed

The aim of both the TSB and the DART consortium is that accessing the DART demonstrator should be a simple and straightforward process. Given that DART is only one of several demonstrators being developed under the Network Services Demonstrator Programme it is also important that there should be a common process established between the various projects towards registering interest, accessing and using the demonstrators. To that end, a common interface will be developed that will provide experimenters and end users with a means of access to all TSB funded demonstrators.

Access to the DART test bed will therefore require potential experimenters and possibly end users to access the common interface from where they can gather information about the test bed characteristics and capabilities, the trials and experimenters that can be supported over the test bed and any contractual or requirements associated with engaging with DART. The common interface will also allow experimenters to define the key aspects of the trials they intend to perform which will then be notified to the contact person responsible for the setup and management of DART trials.

The next step will involve the exchange of more detailed information between the experimenters and the DART consortium in an effort to scope and define the trials in more detail. The trial would then be set up and "go live" over the DART network. A series of API's will also be created so that experimenters can provide DART with content and associated information such as distribution requirements, charging requirements and so on.

Once a trial is ready to go live over the DART network, the DART consortium would also advertise the upcoming trial to end users registered with the test bed. Advertisement of upcoming or current trials may be done on a suitable part of the common interface, on a suitable web page generated by DART or by direct notification of the DART end users. End users would then be able to view the available content by means of a separate API or through a network portal.

The process of accessing the DART testbed, setting up a trial, and undertaking the trial is summarized in table 1 below.

Table 1. Experiment/trial setup and operation process

Advertisement and Interest Attraction	1. The testbed activity is generally advertised to all end users so that they are primed for upcoming "offers". End users register general interest	2. The testbed is promoted to potential experimenters who e.g. content providers, rural service and application providers, network operators and community networks.
	3. An experimenter finds out about the test bed and becomes interested in potentially using it.	4. The experimenter approaches TSB or a DART partner for further information.
First Contact and Trial Scoping	5. Information exchange between experimenter and TSB the experimenter takes place (advice, questions answered etc).	6. Information exchange between the experimenter and Avanti. (more detailed information about the testbed, information about experimenter's goals).
Trial set up	7. The experimenter decides to go ahead with the trial	8. Experimenter agrees to terms and conditions to participate in the trial (probably on-line)
	9. Experimenter provides more details of the test / trial to be performed. This could be automated (through an API.	10. Details of the availability of the trial service are advertised to end users. Those interested sign up to T&Cs.

Table 1. (*continued*)

	11. Content provided to the Avanti network through an API along with key information about when to be delivered, charging details, price level etc.	12. End Users see the content as part of a "walled garden". They are free to choose this or free content such as iPlayer.
Trial operation	13. An end User chooses the content they wish to view and agree to pay for the content, in this case either in SD or HD resolution (they could pay more for HD). For non-cached content they would pay using a different charging tariff perhaps.	14. The end user views the material. The service usage is monitored and accounted according to a predefined billing mechanism. The user is charged a low tariff to view videos.
	15. The service tariff is increased after some time to allow understanding of the effect of price elasticity. The tariff change is advised by email to the end user at both the start of the trial and the time of the changeover.	16. Data is collected from the trial, anonymised, and discussions are held with the users on how they value the service and used to assess the commercial exploitation opportunities.

5 Conclusions

Under the TSB Network Services Demonstrators programme, the DART project will address the future provision of NGA services to rural and remote areas through facilitating testing and trialling centred around the key service enablers and advanced network features discussed in section 4. The project will comprise a design and implementation phase and an experimentation phase.

In terms of timeline, phase 1 of the project kicks off on 1st February 2011 with the design activity, followed by implementation, integration and testing and deployment expected to be completed in late July 2011. Phase 2 of the project, the experimentation phase, will kick off in August 2011 with the demonstrator being made available to third party experimenters and will last for 12 months. The project will complete in August 2012 with the final project review.

For further information readers are invited to contact the project manager, Trevor Barker, at Trevor.Barker@avantiplc.com.

References

1. UK Technology Strategy Board (TSB) home page, http://www.innovateuk.org/
2. TSB Digital Testbed: Network Services Demonstrators programme description, http://www.innovateuk.org/_assets/pdf/competition-documents/briefs/digitaltestbednetworkservicescomp.pdf
3. National Rural Support Programme (NSRP) home page, http://www.nrsp.org.pk/
4. Ishmael, J., Bury, S., Pezaros, D., Race, N.: Deploying Rural Community Wireless Mesh Networks. IEEE Internet Computing 12(4), 22–29 (2008)