Wireless Networks Test-beds: when heterogeneity plays with us

Alessio Botta[1], Antonio Pescapé[1], and Roger Karrer[2]

[1] Universitá di Napoli Federico II, Napoli, Italy {`a.botta,pescape`}`@unina.it`
[2] Deutsche Telekom Laboratories, Berlin, Germany `roger.karrer@telekom.de`

1 Introduction

The Internet has been designed for heterogeneity. In particular, Clark formulates in the design goals of the DARPA Internet that the network must support (i) multiple types of services and (ii) accommodate a variety of physical networks. These design goals are the most important goals besides the interconnection of existing networks and survivability [1]. Moreover, they have led to two design principles: the end-to-end argument and layering. These principles have coined the Internet architecture and were among the key enablers of the stunning success of the Internet. In particular, they have shaped the architecture of the Internet into the well-known hourglass (see Figure 1).

However, is the heterogeneity envisioned four decades ago still the same heterogeneity we experience today? We argue that the notion and the challenges of heterogeneity have significantly changed over time. In particular, the heterogeneity targeted in the early days focused on *co-existence*, i.e. the ability to seamlessly connect different network technologies and shield the upper layer protocols and end systems from the details of the underlying technologies and protocols. In contrast, today, we are challenged to make the heterogeneous technology *concurrently collaborate*. In particular, in the wake of the fixed-mobile convergence, networks suddenly face the challenge to either dynamically choose one of the available technologies or even to concurrently use multiple technologies. For example, modern cities typically provide multiple wireless access technologies, such as GSM, 3G, WLAN oder even WiMax. All these technologies are concurrently available and modern devices are even equipped with multiple radios to take advantage of the concurrent availability of the heterogeneous technology.

The concurrent availability of heterogeneous resources puts forward a set of unprecedented challenges. A first challenge is to decide who controls the resources. Given a modern device with multiple radios that can be used in parallel, some instance has to decide which and how many resources should be used. Should the end system control the resources? Technically, an end-system

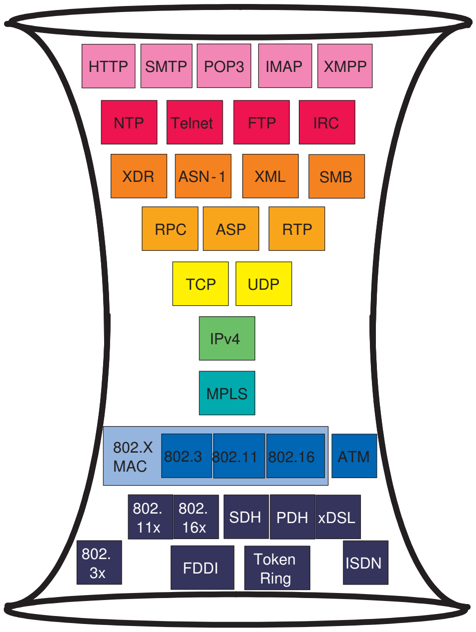

Fig. 1. Internet "hourglass"

approach is able to take the entire end-to-end path into account, including the non-wireless access network as well as potentially multiple providers, whereas a network provider only has information about the technology he deployed. Moreover, an end system may also take end system resources into account, such as battery life. Economically, the end user ultimately needs to be informed about the costs of using multiple technologies. On the other hand, to make such decisions, an end system needs information about the network, including the availability of the different technologies as well as the actual resource usage.

The second challenge is to maximize and manage the usage of the different resources. The availability of multiple technologies in the new heterogeneity allows to exploit the features of the different technologies and consider and arrange them according to some specific metrics: short-delay paths may be used for delay-sensitive applications such as VoIP (Voice over IP), high bandwidth paths for bulk traffic. Again, the question of who controls the resources influences the ultimate outcome.

Finally, the third challenge is to implement the cooperation. In particular, the end-to-end argument, the layering principle and the economical separation between ISPs and end users prevent an easy information exchange among the different layers that hide the technological diversity as well as among networks and end systems. Thus, from a design perspective, it is far from obvious how such an implementation could be done with the Internet protocol stack. Should layering, cross-layer implementation and intermediate layers be considered harmful, or should they be considered as necessary steps towards efficiency for a future development of the Internet?

To sum up, we notice that the challenges raised here are fundamental problems that touch to the very core principles of the Internet design, such as the end-to-end argument [1]. Solutions should therefore be considered only in this entire context. Moreover, it is likely that the road towards solutions will have to consider the tussles raised by the competing and conflicting demands, preferences and needs of the different stakeholders [2].

Our work focuses on shedding light on these questions by planning, deploying and experimentally evaluating test-beds. We argue that many of these questions will ultimately be decided by convincing arguments that are supported by hard facts from real data. Measurements provide insight into the real benefits an operator or an end system may gain. The deployment yields detailed numbers on the deployment costs and ultimately on the incentives for an operator to invest into enhancing collaboration. Ultimately, test-beds contribute to the debate on the fundamental principles of systems and networks design.

This chapter first gives an overview of the heterogeneity in the Internet today. Section 2 thereby emphasizes the tremendous *heterogeneization* of the Internet along various dimensions. The discussion highlights the challenges and the need for a clear structure in the control plane to monitor and manage the heterogeneous devices. Then, Section 3 discusses the impact of the heterogeneity on the Internet architecture and the protocols. The section digs into the fundamentals of the Internet architecture and shows that addressing the heterogeneity requires a fundamental consideration, potentially even a rethinking, of the design principles of the current Internet. We emphasize the need for test-beds to address these challenges, in particular to verify that novel approaches are feasible and comply with the requirements, such as scalability. Section 4 then discusses two heterogeneous test-beds at work: the Magnets test-bed in Berlin and the small scale test-bed at the University of Napoli. We describe such test-beds and show how they allow us to tackle the above challenges. Finally, Section 5 concludes this article.

2 Dimensions of Heterogeneity

One of the greatest achievements of the Internet has been to expand and evolve in spite of the increasing heterogeneity. Heterogeneity has steadily increased over the past decades for a number of reasons and in different dimensions.

A key technical driver for heterogeneity has been the miniaturization of the integrated circuits. In the first years of the Internet, only a few devices were built to send and receive packets over the Internet. With the increasing miniaturization, entire operating systems with complete TCP/IP stacks fit on pocket devices and even tiny sensor nodes. Therefore, the Internet today consists of a plethora of devices with a wide variety of physical capabilities, ranging from low-speed battery-conserving sensor nodes to Gigabit routers.

The technical advances have been joined by a rapid decrease in costs. The Commodore 64, the first computer sold for the mass market, contained only 64 kB of memory - a size that seems ridiculously small when comparing it to today's portable devices. Similarly, CPU speed, disk space and network interfaces have rapidly increased their capabilities while the prices constantly dropped. In 2000, Vint Cerf wrote "By 2020, so many appliances, vehicles, and buildings will be online that it is likely there will be more Internet devices than people online at any given moment." [3]. If any part of this statement is wrong, it will be the time frame by when this vision is achieved. This means that now, and ever more in the future, a plethora of devices will exist that are able to communicate.

Similar to the devices, but at a much lower speed, did heterogeneity increase at the physical and data-link layer. Quite interestingly, though, heterogeneity is limited in the wired Internet where Ethernet has largely triumphed over competing technologies such as ATM. In the wireless world, in contrast, we are still in the infancy of the technological deployment. Significant improvements at both devices and antenna technology will therefore continue to change. Today, WLAN has established itself as the dominant technology for wireless communication, but it is unclear how WLAN will compete with WiMAX. Finally, WLAN is currently also challenged by the advances and the integration of cellular technologies, such as UTMS and HSDPA.

Third, the heterogeneity of the end-host devices has an immediate impact also on the operating systems. Different appliances have often different operating systems. The OS are typically tailored to the capabilities of the end systems and therefore differ for PCs and handheld devices. It is this optimization that causes the heterogeneity: a wireless device will e.g. also send data via TCP-IP, yet the sent traffic pattern may differ from the traffic sent from a PC because it optimizes its resources differently. For example, the traffic pattern of a battery-powered device may significantly differ from a wired-powered PC.

Forth, heterogeneity exists due to the wide variety of applications. While the Internet has originally been dominated by file transfer data, applications and protocols such as the World Wide Web [4] and email have tremendously increased the popularity of the Internet. Today, virtually any imaginable application has been ported to the packet-based Internet, even phone calls are increasingly replaced by VoIP, e.g. via Skype. Finally, two of the greatest surprises are the network games and the file sharing applications. In the first Internet years, network games were predicted to have a very small share of users and traffic[4]. On the contrary, the last games released for network-equipped console games have caused an observable increment in overall Internet traffic[5]. As for the file sharing applications, it is almost known that nowadays they generate a large percentage of all the Internet traffic [6]. Similar evolutions have been noticeable with respect to content: while the Internet was originally text-based, we see an increasing *mediazation* of the content, from text to pictures to multimedia content. The latest push has come from the Web 2.0 that simplifies the exchange of content.

Finally, heterogeneity comes from the different protocols that implement the various functionality. We hereby distinguish three types of heterogeneity. First, different protocols have been specified to implement different functionality. The prominent example here are TCP and UDP, which present a stream and a packet-based transport-layer interface to the application. Moreover, the transport-layer interface shows an interesting evolution. On the one hand, TCP is becoming the dominant transport-layer protocol, and it is even used for e.g. real-time streaming of multimedia content, even though the retransmissions and the strong reaction to congestion make it difficult for TCP to maintain the required streaming rate. UDP, in contrast, is increasingly blocked by firewalls to prevent security exploitations such as DDoS attacks. On the other hand, we see a push towards diversification of TCP, away from the point-to-point protocol towards multipoint-to-point communication. This diversification is motivated by the increasing server-side replication of data (e.g. in Content Distribution Networks and even peer-to-peer networks) as well as client-side multi-homing. Protocols such as Stream Transmission Control Protocol (SCTP[8]) or Structured Streams [9] emphasize the need for enhanced communication support at the transport layer. Thus, in the future we will see an increasing heterogeneity of TCP-friendly protocols that open multiple streams in parallel. The second level of heterogeneity comes from the different flavors of a protocol, e.g. TCP. Over the past years, many TCP variants have evolved, such as TCP Reno, Tahoe, NewReno, WestWood, FAST, BIG, etc. These variants were the response to the increasing heterogeneity of the lower layers: some TCP variants target high-speed wired networks, others target wireless networks. While all flavors have their pros and cons, we largely ignore today how the different variants inter-operate, e.g. in the case that they are concurrently deployed in wired-cum-wireless networks. Finally, we notice an increasing "heterogeneity" in protocols, especially at the application layer, due to security constraints. In particular, today's firewalls increasingly block potentially suspicious ports. As a result, applications "hijack" ports and protocols to tunnel content through the firewalls. HTTP is one of the most (ab)used protocols for this purpose today because port 80 is most frequently open. Similarly, Skype is known to actively search for holes in the firewall. In the future, as long as the binding between port and service identification prevails, we expect that the raising security concerns lead to more heterogeneity in (ab)used application-layer traffic.

3 Impact on architecture and protocols

The heterogeneity at both the application layer and the lower layers raises questions how the Internet architecture should evolve in the future. At the

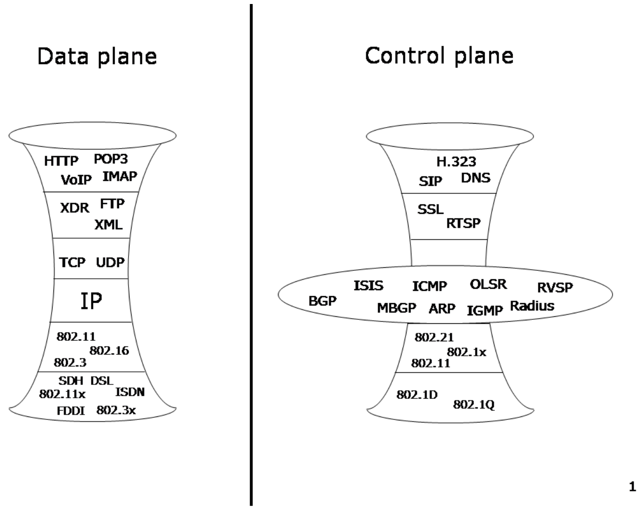

Fig. 2. IP Waist

data plane, the hourglass model has shown to have many advantages. However, at the control plane, the picture looks quite different, as Figure 2 shows. The constant demand for new services has led to an almost inverted hourglass shape, with the thick waist around the network layer. Moreover, cross- and inter-layer protocols have started to blur the layer boundaries. While some of the problems can only be solved with a clean slate approach, e.g. splitting locators and ids as well as services and ports, other issues concern the layering structure of the control plane. We identify two challenges.

The first challenge is to address the question whether layers are necessary, and how strict they need to be. Or, in other words: is the layering principles a necessary precondition to ensure the future evolution of the Internet? Or are they even preventing evolution today where the Internet challenges are no longer just technical, but also economical and social? Just consider security problems: how many problems arise because the information exchange among the layers does not exist? Similarly, to which degree will it be possible to organize and optimize multiple heterogeneous (wireless) access technologies without an integration at the control plane?

To give a brief example, consider routing in wireless mesh networks. Routing is typically addressed at the network layer with IP. However, in wireless mesh networks, the discussion is ongoing if routing in the mesh should be performed at layer 3 or, as the 802.11s standard prescribes, at layer 2? The advantages of a layer 2 approach are that the entire mesh cloud is visible as a single "node". It integrates well with the broadcast properties of the mesh and promises high performance. In contrast, a layer 3 approach reveals the mesh topology, allows a re-usage of IP-layer mechanisms such as IP multicast and mobile IP. Finally, layer 2.5 approaches promise a combination of the advantages, but also require reprogramming. Thus, this example shows that the traditional distribution of functionality onto the different layers may be subject to change in the future. Similar considerations are possible for network coding approaches [10], multipath [11] and security.

The second challenge is the distribution of control. As outlined above, it is unclear how the control of network resources should be divided in the future: end systems or the network (i.e. ISPs). Consider the case of multi-homing. Should multi-homing be implemented in the end device or inside the network? Both possibilities have pros and cons. Or is there even a third, compromising option, e.g. in the form of middleboxes that provide a limited support of customization for end systems but are still under the control of ISPs?

All these questions and the subsequent decisions must be supported by test-bed implementations and evaluations. Only test-beds provide the necessary power to assess the performance gain and the implementation overhead. By deploying and experimenting with these parameters, vital insight can also be gained on the incentives for either solution. For example, by deploying a wireless test-bed in a city, both ISPs and users can be integrated into the test-bed and their needs and their willingness to cooperate can be investigated under real conditions.

4 Heterogeneous Wireless Test-beds at work

This section describes two test-beds and some experiences we had with them. We aim at providing some information useful to set up heterogeneous test-beds in reality, to perform measurements on them, and to interpret the outcomes of the experiments in order to shed light on the potential and limitations of current heterogeneous networks.

The first test-bed we describe is located in Berlin, Germany. This project is called Magnets and has a lot of interesting features which span from being a joint research-operational network to mixing different access network technologies, from being located in the center of a very big city to having a multi-hop Wireless Wide Area Network (WWAN) as a backbone, from having such a wireless backbone made with off-the-shelf components to being able to reach more than 60 Mbps of throughput. We describe the complete plan and then provide some details regarding the different components: the wireless backbone, the wireless mesh networks that will be interconnected, and the points of integration with other technologies (i.e. GPRS, UMTS, and WiMAX). Also, we present the experimentations we have performed providing some interesting results.

The second test-bed has a smaller scale with respect to the first. However, it comprises a large mix of different devices, operating systems, and access networks. The smaller scale improves the ability to control the environment, the behavior, the measurements and thus, in turn, the predictability of the experiments and the interpretation of the results. Therefore, while Magnets allows to assess what a real user would experiment, the small scale test-bed allows to go in deep into the root causes of the observed behaviors. We present the architecture of this network, our measurement methodology, and some obtained results.

4.1 Large Scale Test-bed: Magnets, a next-generation access network

Magnets is designed as a next-generation wireless infrastructure. It consists of two main parts: a wireless mesh with 100 nodes and a high-speed wireless backbone with a raw end-to-end throughput of 108 Mb/sec. Besides the size of the network in terms of nodes and link speed, a key distinguishing characteristic of Magnets is its heterogeneity along several dimensions: it features multiple wireless interfaces with diverse link characteristics, nodes with varying degrees of processing and storage capabilities, and interconnection of multiple mesh networks with disparate routing protocols.

To further exploit this uniqueness, Magnets is designed with a three-fold goal. First, Magnets is designed as a semi-productive network. That is, the network is used as a testbed, e.g. to experimentally evaluate protocols, but at the same time the network is integrated into the productive campus network of the TU Berlin. Therefore, Magnets will extend the Internet coverage of the students. This combination eventually allows us to also perform measurements of real user traffic and to evaluate protocols under realistic conditions. Second, with Magnets we will systematically assess ways to build wireless mesh networks. For example, the WiFi backbone is designed as a fully planned network, optimized for throughput. In the mesh, we encounter several constraining factors, such as the buildings, the density of already deployed access points. We will compare the capabilities of this mesh with other mesh networks in Berlin that are driven by communities and their structure is therefore unplanned.

The Magnets architecture consists of 3 parts: a high-speed wireless 802.11 backbone, an 802.11-based wireless mesh network and integration points to alternative technologies (GPRS, UMTS and WiMax). Next, we describe the three parts in more detail.

Backbone

The Magnets backbone is designed to interconnect 2 facilities in Berlin with a high-speed connection that is purely wireless. After a careful planning that involved network-specific parameters, such as finding buildings to provide line-of-sight, but also economical parameters such as the deployment costs or rent for space, we decided on a layout that consists of 5 nodes, as depicted in Figure 3. The total distance between the two end points at the T-Labs and T-Systems is 2.3km. All nodes reside on top of high-rise buildings and have

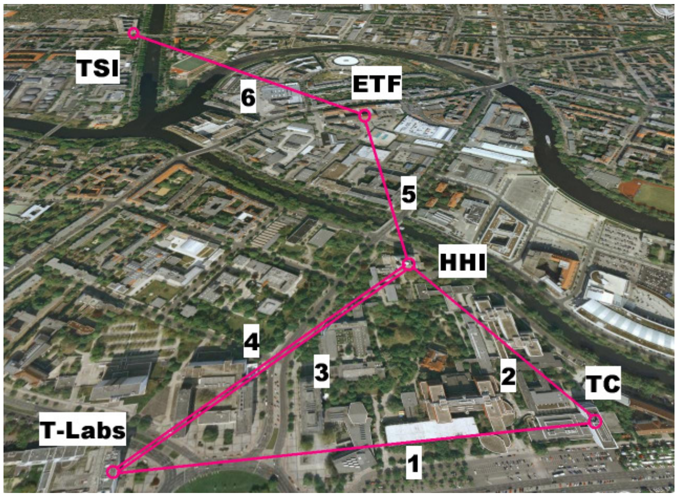

Fig. 3. Magnets WiFi backbone in the heart of Berlin.

unobstructed line of sight. All transmissions are in the unlicensed spectrum (2.4 GHz and 5 GHz) range.

The nodes are designed to ensure an efficient multi-hop communication. In particular, to avoid well-known performance and unfairness problems in multi-hop communication [16], we decided that each wireless link should be operated by an individual access point. Thus, a total of 12 WiFi access points (APs), suitable for outdoor usage, are mounted along the antennas to shorten the cable length between the antenna and the AP. While the APs all support 802.11a/g modes at 54Mb/sec, with the option to improve to 108Mb/sec via Super-A/G, the use of the frequency band is defined by the antennas. We decided to operate 8 APs at in the 2.4 GHz band and the rest in the 5 GHz to also have heterogeneity in the transmission frequencies, e.g. to observe the impact of interference in a dense urban area.

Since most nodes on the buildings consist of multiple access points, we decided to inter-connect them via a workstation. In addition to the pure connecting of the APs, these workstations can additionally be used to inject traffic and to monitor the forwarded traffic. To ensure that the workstation is not becoming the bottleneck, we equipped them with a fast 3GHz processor and 1 GB of RAM.

This setup allows us to perform a range of measurements. First, we are able to observe the per-link characteristics. Since every node is physically located at a different environment, we are able to monitor the node and link performance as a function of the link distance, the capacity and the interference at the receiver. In terms of distance, the links vary from 330m up to 920m. We are able to assess short-term statistics, e.g. to assess the frequency and the impact of link-layer retransmissions, as well as long-term statistics, e.g. the evolution of the link speed over several days or even months. Moreover, we are able to monitor low-level information, such as link-layer retransmissions as well as end-to-end throughput, e.g. the performance of different TCP versions over multiple wireless hops and even over wired-cum-wireless connections.

WiFi Mesh

The WiFi mesh with 100 mesh nodes is deployed on the campus of the TU Berlin in collaboration with the IT department. The mesh will consist of 100 nodes deployed as a combination of in- and outdoor nodes. The mesh shall cover the entire campus area and thereby provide Internet access to the students.

For the selection of the hardware for the mesh nodes, we opted for two hardware platform: routerboards and Avila Gateworks. The routerboards provide maximum extensibility. The RB500, e.g., has the ability to attach external storage via a compact flash card, which is important to add management and measurement tools and eventually to collect traces. With the help of a daughterboard, up to 6 MiniPCI slows provide ample opportunities to attach WiFi cards or to insert alternative technologies on each node. The drawback of the Routerboards is the limited CPU speed. Therefore, the largest portion of the mesh will be built with Avila Gateworks. These network processor-based boards easily achieve throughputs of 100 Mbps and are therefore well suited for high capacity.

To perform experiments, we set up the nodes with 2 particularities. First, all nodes run openWRT, a Linux-based operating system that provides the flexibility to access kernel information. Moreover, we use Atheros cards with MadWiFi to get access to the MAC and PHY statistics. This software provides ample opportunities and flexibility to deploy and evaluate protocols at any layer. It allows experimental evaluation of benefits and drawbacks of cross-layer optimizations that have been proposed in the research literature [14]. Our main objective here is to shed practical, experimental light on the ongoing discussion. Second, we equip most nodes with 2 boards: a main board for data transmissions (Avila or Routerboard) and a secondary board for monitoring (mostly a cheaper ASUS board). The reason for the monitoring boards is that CPU, memory and network speed may cause limitations to concurrently transfer data and perform monitoring on the nodes. For example, running tcpdump on multiple interfaces may severely slow down the performance of the nodes.

In terms of deployment, we decided to build three different "mesh networks": a "smoke" test-bed, consisting of pairs of boards only, an indoor test-bed of 20 nodes and finally the outdoor test-bed of 100 nodes. The smoke test-bed is used for node configuration and testing. That is, before any software is deployed on a mesh, it must be tested and shown to be runnable at least on two nodes. While the smoke test does not guarantee that software

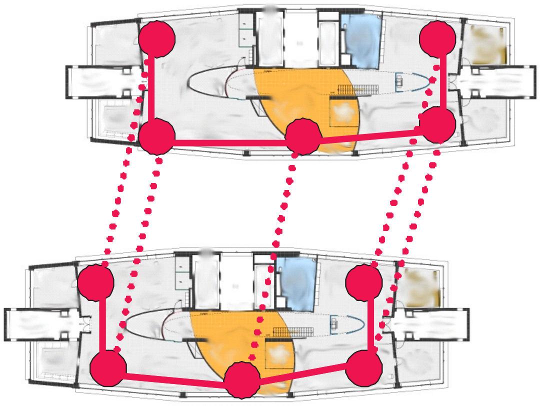

Fig. 4. Deployment of the indoor test-bed.

does not crash deployed nodes, it reduces the risk that somebody has to climb and unscrew the mounted indoor or outdoor boxes.

Figure 4 shows the deployment of the indoor test-bed. The figure shows that 5 mesh nodes are deployed on opposite sides of the T-Labs building and one node at the center. Generally, we made sure that two neighboring nodes are within range of each other, whereas two-hop neighbors have a limited or no connectivity due to glass and other obstructing material. The same applies to the connectivity among floors. Here, two nodes at the same physical location but on different floors are within range, but nodes that are two floors apart or nodes at different floor locations have limited connectivity.

Heterogeneous Nodes

Besides the heterogeneity given by environmental factors (interference) at the different stations, the frequency range, the physical distance of the links and the node density, another degree of heterogeneity can be added to the network by augmenting with alternative wireless technologies. Taking advantage of the up to 6 Mini PCI slots of the nodes, alternative technologies can be added to the network, such as GPRS, UMTS, and WiMax. By superimposing multiple technologies within the same area, we are able to address questions on how to operate, manage and optimize future 4G networks. Again, important here is the ability to gain first-hand experience in a semi-productive testbed with

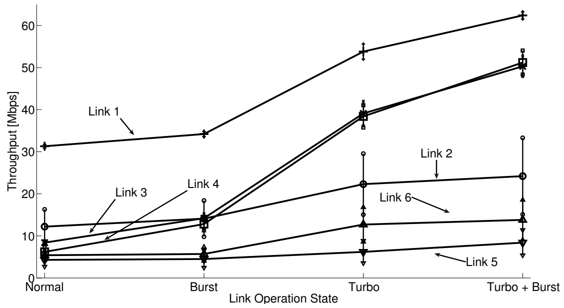

Fig. 5. Throughput of Magnets backbone links.

real user traffic. Issues such as TCP performance during horizontal and vertical handovers between multiple access technologies can be experimentally assessed and evaluated.

Constellation of Mesh Networks

Berlin is the center of the Freifunk community. Driven by the need to provide Internet connectivity in an area where DSL is not available because (ironically) fiber but not copper is available, the community stepped up to build the Freifunk network [3]. Interesting research questions arise when we think about the options to inter-connect wireless mesh networks that are under different administrative authorities. For example, it will be intersting to observe the behavior of two meshes that run different routing protocols or have different routing metrics. Will it be necessary to develop novel protocols to separate the domains (such as BGP in the Internet), or is it possible to weave the two meshes seamlessly together, e.g. to simplify mobility?

Results

To provide initial insight into the heterogeneity we are able to observe with Magnets, we consider the per-link throughput of the backbone links. For this purpose, we measure each link individually. We generate UDP traffic using iperf and vary the mode of each AP between 802.11a/g and the options provided by Super-A/G (Turbo- and Burst mode). For 600 seconds, we generate traffic at 70Mpbs, which lies above the saturation rate of the link. At the receiving workstation, we monitor the incoming packets with tcpdump. Then, to calculate the bandwidth, we sample the traces at 50ms intervals.

Figure 5 shows the resulting throughput. The x-axis denotes the mode, the lines show the average throughput of the different links. Finally, the whiskers show the standard deviation. In normal 802.11 modes, the figure shows that

---

[3] http://www.olsrexperiment.de

link 1 outperforms the others with an average throughput of 31.3 Mbps. Moreover, the low standard deviation of 0.9 Mbps indicates that the link is very stable. Next, links $2-4$ have an average throughput between 6.2 and 12.2 Mbps. These links operate in the 2.4 GHz range and the throughput degradation is attributed to interference. Finally, links 5 and 6 are the weakest links, with an average bandwidth of 4.3 and 5.4 Mbps respectively. Link 5 has strong interference because the ETF building is lower than the others, and link 6 spans a much larger distance with 930m. Thus, we conclude that the link characteristics vary significantly even though they have been measured in the same test-bed.

Then, we assess the impact of *Turbo* and *Burst Mode* on the link performance. Even though the reference manual indicates a doubling of the throughput via *Turbo Mode* and an increase of 10 Mbps with *Burst Mode*, it is not obvious how these modes impact the link characteristic of *MagNets*. As we can see in Figure 5, the *Turbo* and *Burst Mode* increase the throughput on link 1 significantly. Compared to the basic mode (31.3 Mbps), the throughput increases with *Burst Mode* to 34.2 Mbps. *Turbo Mode* boosts the throughput to an average of 53.8 Mbps. Finally, with both modes enabled, the average throughput reaches 62.4 Mbps! Thus, we conclude that link 1 matches the original specifications and expectations of *Turbo* and *Burst Mode*. Link 3 also shows throughput gains with *Turbo* and *Burst Mode*. The corresponding rates are 8.4, 14.2, 39.1, and 50.3 Mbps. Note here that the improvement with *Turbo Mode* is more than twice the base rate. All other links obtain results comparable to link 3. With the exception of links 5 and 6 that suffer from the above mentioned problems, we can state that the performance was significantly improved with *Turbo* and *Burst Mode* enabled. Therefore, we argue that the MagNets backbone is able to support a substantial amount of traffic.

4.2 Small Scale Test-bed: a heterogeneous network at the University of Napoli

In this section we present another example of a real test-bed useful to perform experiments aimed to uncover the potential and the problems of heterogeneous networks. Such test-bed is sketched in Figure 6. As shown, it is composed of a number of heterogeneous wireless/wired networks. Over such test-bed a number of different configurations have been produced: we have varied several configuration parameters such as the operating system, end user device, access network, transport protocol, and traffic condition.

All these components constitute our definition of end-to-end path. In details, we define an end-to-end path (e2eP) as:

$$e2eP = (S_{UD}, R_{UD}, S_{OS}, R_{OS}, S_{AN}, R_{AN}, Protocol, Bitrate) \tag{1}$$

where UD stands for the User Devices (S_{UD} at sender side and R_{UD} at receiver side) (e.g. Laptop, Palmtop, Workstation, etc.); OS identifies the

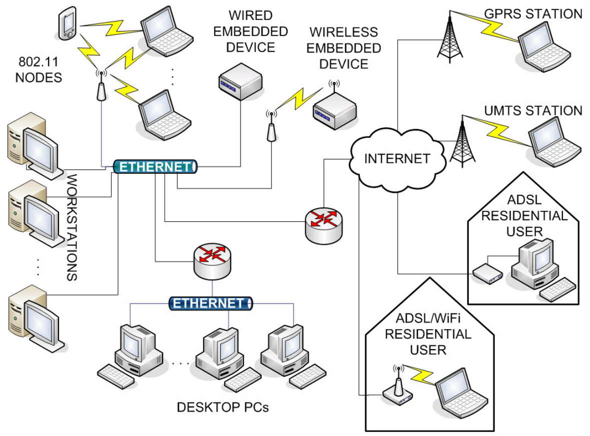

Fig. 6. The small scale heterogeneous test-bed.

Operating Systems of each of the two users (e.g. Windows, Linux, Linux Familiar[4], etc.), S_{OS} at sender side and R_{OS} at receiver side; AN is the Access Networks (LAN, 802.11, ADSL, GPRS, etc.), S_{AN} at sender side and R_{AN} at receiver side; $Protocol$ identifies the protocol the users are communicating through (e.g. TCP, UDP, SCTP, etc.); and, finally, $Bitrate$ is that imposed by the application. By combining all these variables, our test-bed allows to set up about 350 different end-to-end paths.

Measurement methodology

For the measurements we used an active approach and our tool called Distributed Internet Traffic Generator (D-ITG) [17]. D-ITG is able to generate a multitude of traffic patterns by combining pairs of PS (Packet Size) and IDT (Inter Departure Time). In this way it is possible to generate controlled yet realistic traffic. In this chapter, we present UDP Constant Bitrate (CBR) traffic profile obtained with constant PS and constant IDT. This allows to draw a reference curve for successive analysis and to reduce the number of variables. Thanks to its features, D-ITG can be used as an active measurement tool. It can measure and analyze one-way-delay (OWD), round-trip-time (RTT), packet loss rate, jitter, and throughput, using the various components of such platform: (i) sender; (ii) receiver; (iii) decoder; (iv) log server. The experiments have been carried out by producing three traffic conditions named *Low*, *Medium*, and *High Traffic* [18]. The characteristics of such traffic are reported in Table 1.

The measurement stage has been performed between December 2003 and November 2004, in the day hours between 9:00 am and 6:00 pm. Such stage allowd to collect over $34GB$ of traffic traces. The traces have been carefully

---

[4] An open source porting of Linux for Palmtop devices

Table 1. Characteristics of Measurement Traffic.

Traffic Condition	**IDT [s]**	**PS [Bytes]**	**Bit Rate [Kbps]**
Low	1/100 s	$\in \{32, 64, ..., 1024, 1500\}$	$\in \{26.1, ..., 819.2, 1200\}$
Medium	1/1000 s	$\in \{64, ..., 512\}$	$\in \{512, ..., 4096\}$
High	1/10000 s	$\in \{64, 128\}$	$\in \{5120, 10240\}$

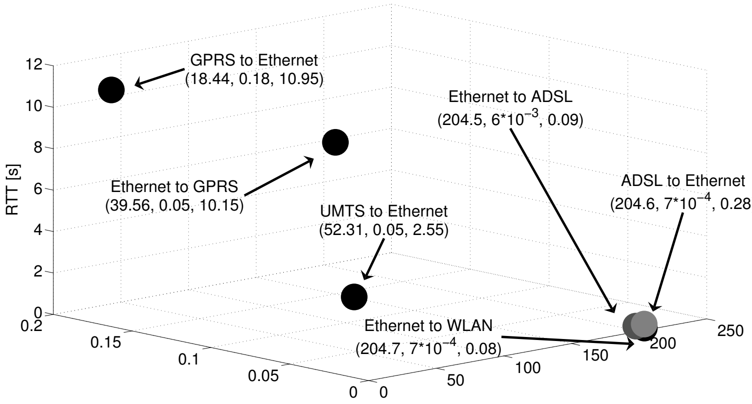

Fig. 7. Throughput, jitter, and round trip time over the small scale test-bed.

inspected and sanitized detecting and removing samples affected by errors. At [7] we made freely available several archives containing the outcomes of measurements over real networks (not only those we used in this work). Each archive contains files with samples of QoS parameters measured over several end-to-end paths.

Results

For the purpose of this chapter, it is interesting to report here a comparison of what we achieved with different network configurations. In Figure 7 we sketch a three-dimensional plot showing the average throughput, jitter, and round trip time we obtained in the *low* traffic condition with a PS equal to 256 Bytes (i.e. with a generated bit-rate equal to 204.8 Kbps). For the sake of clarity, we have selected 6 network path that present different characteristics only in terms of AN. This allows us to exclude the other variables (OS, EuD, ...) from the possible explanation of the results. This is a very important point when performing measurement on a such heterogeneous test-beds: it is necessary to vary only one variable at a time.

Figure 7 shows a clear separation between slow (i.e. GPRS and UMTS) and fast (i.e. ADSL and WLAN) Access Networks. The very different results obtained by the slow AN have the effect of making the fast AN appear as a single point in the graph. However, differences are noted are between them.

This figure allows to separate the ANs not only looking the the obtained throughput, which was expected, but also looking the the other statistics. This is useful because, in some cases, it is simple to obtain an estimate of the RTT and then of the jitter (e.g thanks to the TCP acknowledgments), while it can be difficult to estimate the throughput. Having such statistics and using these results, several applications can be devised. For example, in [12] and in [13] automatic identification of network characteristics is performed.

5 Conclusion

The Internet has evolved from a simple and purpose-specific network to the common infrastructure for global user communication. Its current shape was impossible to imagine for its designers. As a consequence, it has now become something very different from the initial plan. However, it still preserves some of the original protocols which causes different problems to both the users and the network administrators with respect to new services and applications.

To understand the benefit and the limitations of the current Internet, in this chapter we have analyzed the main causes of its heterogeneity. We have seen that such causes can be partitioned along different independent dimensions, and we have explored these dimensions. Moreover, we have identified the main challenges that this infrastructure poses with specific regard to the protocols.

Thanks to the use of two real life examples we have then observed how the heterogeneity can be studied. We have described a large and a small scale test-bed, both characterized by an high degree of heterogeneity. We presented some results obtained on the test-bed and showed how the obtained results, can be exploited for addressing the aforementioned issues.

Heterogeneity - in particular heterogeneity where multiple technologies can dynamically be chosen in parallel - raises fundamental questions: how to optimize a network, who decides, who pays. For this reason test-beds are needed to evaluate principles, to show tradeoffs and to create firm arguments.

Moreover, we believe that the heterogeneity may represent a problem; but, at the same time, it provides a great opportunity that must be exploited for network and service convergence: convergence of fixed, wireless and cellular technology increases the demand to support heterogeneity within a single Internet architecture.

Concluding, heterogeneity is a challenge, far from easy to solve at the Internet scale in a distributed fashion. Yet, heterogeneity needs to be addressed to simplify the use of different technologies and to provide unified services to users.

References

1. J. Saltzer, D. Reed, D. Clark, “End-to-end arguments in system design”, ACM Transactions on Computer Systems, 2(4):277-288, Nov 1984
2. D. Clark, J. Wroclawski, K. Sollins, R. Braden, “Tussle in Cyberspace: Defining Tomorrow’s Internet”, ACM Sigcomm, August 2002, Pittsburgh, PA
3. Vinton Cerf (June 19, 2000), “Visions of the 21st Century: What Will Replace the Internet”. TIME.com.
4. Kevin Kelly (August 2005), “We Are the Web”, Wired.com.
5. Graeme Wearden (December 2004), “Does the ’Halo 2’ effect threaten broadband?”, CNET News.com.
6. Sean McCreary, kc claffy (2000), “Trends in wide area IP traffic patterns - A view from Ames Internet Exchange”, ITC Specialist Seminar.
7. `http://www.grid.unina.it/Traffic` [Online]
8. R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, V. Paxson (2000), RFC 2960, “Stream Control Transmission Protocol”.
9. B. Ford, “Structured Streams: a new transport abstraction”, ACM Sigcomm, Aug 2007, Kyoto, Japan
10. S. Katti, S. Gollakota, D. Katabi, “Embracing Wireless Interference: Analog Network Coding”, ACM Sigcomm, August 2007, Kyoto, Japan
11. P. Key, L. Massoulie, D. Towsley, “Combining Multipath Routing and Congestion control for Robustness”, Proc. CISS, March 2006
12. A. Botta, A. Pescapé, G. Ventre, “Identification of Network Bricks in Heterogeneous Scenarios”, In Proc. of the First IEEE LCN Workshop on Network Measurements, November 2006, Tampa, Florida.
13. W.Wei, B.Wang, C.Zhang, J.Kurose, D.Towsley. “Classification of Access Network Types: Ethernet, Wireless LAN, ADSL, Cable Modem or Dialup?”. In Proc. of IEEE INFOCOM 2005, Miami, March 13-17, pp. 1060-1071.
14. V. Srivastava, M. Motani (December 2005), “Cross-layer design: a survey and the road ahead”, IEEE Communications Magazine, Volume 43, Number 12, pp. 112119
15. National Science Foundation (January 2006), “GENI: Global Environment for Network Innovations - Conceptual Design and Project Execution Plan”.
16. Zhenghua Fu, Petros Zerfos, Haiyun Luo, Songwu Lu, Lixia Zhang, Mario Gerla (2003), “The Impact of Multihop Wireless Channel on TCP Throughput and Loss”, INFOCOM 2003.
17. `http://www.grid.unina.it/software/ITG` [Online]
18. G. Iannello, A. Pescapé, G. Ventre, L. Vollero, “Measuring Quality of Service parameters over heterogeneous IP networks”, International Conference on Networking (ICN 2005), LNCS 3421, pp. 718-727, 2005.- April 2005, Reunion Island