ACCEPTED FROM OPEN CALL

MONITORING EMERGING IPV6 WIRELESS ACCESS NETWORKS

PEDRO MARQUES, PORTO UNIVERSITY
HELDER CASTRO AND MANUEL RICARDO, INESC PORTO

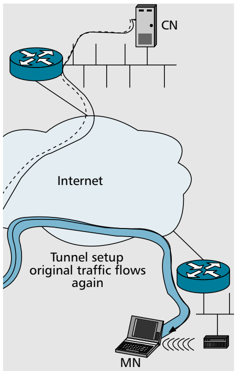

Foreseeing a future where IPv6 and mobile terminals play an important role in public access communication networks, the authors introduce a monitoring system capable of identifying relevant traffic flows and tracking them while terminal equipments move between network attachment points.

ABSTRACT

Foreseeing a future where IPv6 and mobile terminals play an important role in public access communication networks, this article introduces a monitoring system capable of identifying relevant traffic flows and tracking them while terminal equipment moves between network attachment points. The mobile flows are characterized and represented so that individual users and flows can perceive the quality of service they receive, and operators can have global traffic views of their heterogeneous access networks.

INTRODUCTION

The fourth generation (4G) of mobile communications is often associated with the *always best connected* concept, which is characterized by:

- Terminals equipped with multiple network interfaces that can include 802.11, General Packet Radio Service (GPRS), and Universal Mobile Telecommunications System (UMTS)
- Terminals capable of dynamically selecting the least costly network interface
- Solutions capable of enabling seamless handover of terminals

Some solutions have been proposed to support the integration of wireless LAN and GPRS/UMTS systems, known as tight or loose coupling [1]. A closer look into current research activities working with the loose coupling concept also shows that some projects design their mobility solutions based on Mobile IP and, in particular, Mobile IPv6 (MIPv6).

IPv6 [2] uses 128-bit addresses and enables stateless address auto-configuration. Using the latter, a terminal arriving at a new access network becomes automatically configured with a public network address. This address is formed based on the MAC address of the network interface card in use and the visited network prefix, announced in advertisement messages sent by access routers. An IPv6 packet contains a fixed size header, helpful for fast routing, and optionally some extension headers.

MIPv6 [3] improves IPv6 by enabling packets to always reach the terminal, independent of its location. For this reason, MIPv6 is often mentioned as capable of enabling vertical handovers (i.e., handovers between network access technologies or operators). Using MIPv6, TCP or Reat-Time Transfer Protocol (RTP) connections are not dropped, and packets are always delivered in both directions.

New wireless access networks based on IPv6 are also expected to implement the all-IP concept. Time-critical information, such as that produced by a video telephony service, can be conveyed in IP packets that, in turn, are transported along with IP packets transporting data from other services. This requires the IP layer to differentiate the transport service it provides to IP packets. IP quality of service (QoS) solutions, mainly those based on differentiated services (DiffServ), are being evaluated and combined with the layer 2 QoS solutions provided by wireless access technologies. QoS assurance in wireless heterogeneous networks is particularly difficult to achieve, mainly due to the high bit error rates of wireless links and handover periods. The former may force packet retransmission and higher delays; the latter may cause temporary losses of connection.

Network traffic is better understood, from the QoS provisioning and management points of view, if modeled as flows. A unidirectional packet flow may be defined as a set of packets observed close in time, and sharing common characteristics such as the source and destination addresses, next header protocol, and sometimes source and destination ports. A packet flow may also be classified as elastic or real-time. An elastic flow can be extended in time, and transports, for instance, a file by FTP or an HTML page; it is usually supported by TCP, which dynamically adapts the flow rate to the network and receiver levels of congestion. In a real-time flow, the time of packet arrival at the receiver is relevant: if a packet containing a piece of voice misses a predefined deadline, it may be of no use for the receiver; this type of flow is usually supported by RTP/User Datagram Protocol (UDP).

This article presents a new traffic monitoring system, targeted at emerging IPv6 wireless access networks and strongly oriented toward the capture and characterization of mobile flows. When deployed over an IPv6 access network, this monitoring system first gathers the IPv6 network topology. Then it detects the flows transported by the network and network terminals. When a terminal moves or a flow is rerouted, the monitoring system detects them in real time, displays

1536-1284/05/$20.00 © 2005 IEEE

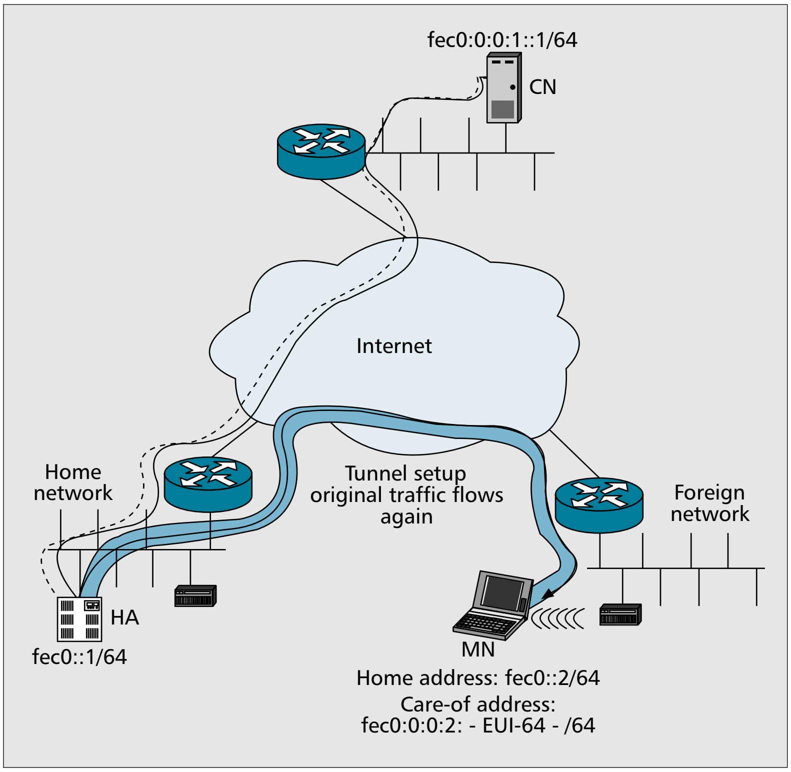

■ **Figure 1**. *Tunnel between the HA and the MN.*

these moves over the network topology, and gathers the relevant statistics.

This work is presented as a set of sections. First, Mobile IPv6 is reviewed. A flow is defined next and discussed from the mobility perspective. An overview of the traffic monitoring system follows. Its central components, the sampler and collector, are characterized, as well as the techniques they use to sample packets. Then some views of the traffic monitoring system are given based on scenarios used to validate the monitoring system. Finally, the issues requiring further work are characterized, and conclusions are drawn.

Mobile IPv6

The Internet Engineering Task Force (IETF) Mobile IP working group proposes an IPv6-based solution for supporting terminal mobility. When a mobile terminal, known as a mobile node (MN), moves from its home network to another network, it gets a second IP address on the foreign network, the care-of address (COA); then the MN registers the COA in a mobility agent located in its home network: the home agent (HA). Having noticed the MN movement, the HA starts to receive packets destined to the MN, and sends them through a tunnel to the MN, whose endpoints are the HA address and COA. Figure 1 represents packets sent by a node in the internet, the correspondent node (CN), to the MN. Since the MN is away from home, these packets are captured and tunneled by the HA to the MN.

Although conceptually simple, the tunnel concept leads to strong routing inefficiencies. In order to eliminate them, the MN may enter a return routing procedure and inform the CN about its COA. The CN, having this information, starts sending packets directly to the COA in the visited network.

Mobile Flows

Many definitions of packet flow exist [4] that include layer 3 or 4 flows, and may also consider flow labels or security associations semantics. In this work a simple packet flow definition commonly used in IPv4 is adopted. A packet flow is assumed to be an ordered set of packets traveling through the network in a given direction, close in time, and having the following characteristics in common: source address; destination address; transport protocol, for now limited to TCP or UDP; source port; and destination port. The first two fields can be found in the IPv6 packet header; the third field can be found in the IPv6 header or an IPv6 extension header; the fourth and fifth fields belong to transport (TCP, UDP) headers.

A flow can be observed when, for instance, a CN communicates with an MN located in its home network. When the MN moves to a foreign network, as shown in Fig. 1, the packets of a flow are placed in the tunnel, starting at the HA and terminating at the COA, collocated at the MN. Thus, between the HA and the MN, the original packet is placed inside another packet whose destination address is the new COA, source address is the HA address, and next header protocol field corresponds to the IP-in-IP protocol (tunnel).

A third situation appears when, away from home, the MN has seen its route optimized to the CN. Now an extension header is added to the IPv6 packet. A packet sent in the CN→MN direction has a routing header extension added by the CN; this extension header carries out the MN home address, while the destination address in the IPv6 packet header contains the current COA. A packet sent in the MN→CN direction has a destination options header extension added by the MN; this extension header carries out the MN home address, while the source address in the IPv6 packet header contains the current COA. In cases where both nodes are mobile, both headers may be required.

Flow packets, as seen, may be represented differently in IPv6 networks supporting mobility. The central problem consists in inspecting a packet and deciding with which flow it shall be associated.

Traffic Monitoring System Overview

The traffic monitoring system was designed to satisfy six requirements:
- Be capable of detecting MIPv6 traffic flows
- Detect the mobility of flows and terminals in real time
- Represent individual flow information from a QoS perspective
- Resolve the access network topology
- Represent information from users and operator points of view
- Use sampling techniques

When deployed over an IP access network, the monitoring system first gathers the IPv6 network topology. Then it detects the flows currently transported by the network and terminals. When a terminal moves or a flow is rerouted, which implies MIPv6 address changing, the monitoring system detects it in real time and displays these moves over the network topology. Individual flows are represented as arrival and departure curves, complemented with statistics. Using them, information about the debits, delays, and losses of individual flows and terminals can be known. The operator has also aggregated information on network interface debits, link delays, and losses.

From the architectural point of view, the traffic monitoring system was devised to work with generic access network topologies, and independent of the wireless technologies in use. Packet capturing is required, and can be performed at any IP network element, such as the edge router, internal router, and even mobile terminals.

Two type of components are used in the traffic monitoring system: the Sampler and the Collector. The Sampler is the entity in charge of gathering information about the traffic traversing a network element; it copies and filters packets headers in network element interfaces, storing relevant information about them. The Collector is a central and unique unit, in charge of processing the information provided by Samplers. Being the access network traffic modeled as a collection of unidirectional flows, each Sampler forms in real-time a flow information structure, and flushes it regularly to the Collector. When receiving this information, the Collector builds a global information structure, where all information is combined.

The Sampler

Any IP access network element may have a sampler installed. The sampler takes advantage of existing network element configuration tools, such as the *route* and *ifconfig* commands in UNIX-based routers, to detect local network interfaces, and IPv6 addresses and masks assigned to each network interface. Using this information, a sampler knows in which network and element it is installed. Then the sampler contacts, authenticates, and registers with a collector, to which it transfers the network information; in normal operation, a sampler captures and processes packets.

Packets are captured in a network interface basis. In a UNIX-based router, for instance, the *libpcap* library is used, and is enabled to get a copy of selected information for relevant packets. The packets are selected using filters defined based on: IP host addresses (source and destination); network and mask addresses (source and destination); source and destination ports; and transport layer protocol.

A packet is then processed. This processing includes the identification of mobility extension headers, flow identification, and packet sampling. Two IPv6 extension headers related to mobility need to be inspected: the Destination Options and Routing headers. As mentioned above, a packet sent in the MN→CN direction includes a Destination Options header, while a packet sent in the CN→MN direction includes a Routing header. If an IPv6 Destination Options header is found in a packet (direction MN→CN), its home address option is used to help solve the MN source home address. In normal network interfaces, the Destination Options header contains the MN home address, and the packet source address contains the MN COA. However, these values are switched by the CN to which the packet is destined: Destination Options will contain the MN COA, and the packet source address will contain the MN home address. In this way, TCP connections can be maintained. The Routing header used in the CN→MN direction also has to be parsed according to these extension header rules, which also interchange packet destination address with home address at the MN.

A flow is represented in a sampler by its real source and destination addresses. By resolving these values using the rules mentioned, a packet can be associated with a flow. With each packet is associated information that includes its sequence number or the acknowledgment (ACK) sequence number in a TCP flow, or a hash value for a UDP flow; its source and destination COAs are also kept. This information is stored in a data structure organized on a flow basis. Periodically, the flow structure is converted into a NetFlow9-based [5] format and sent to the collector.

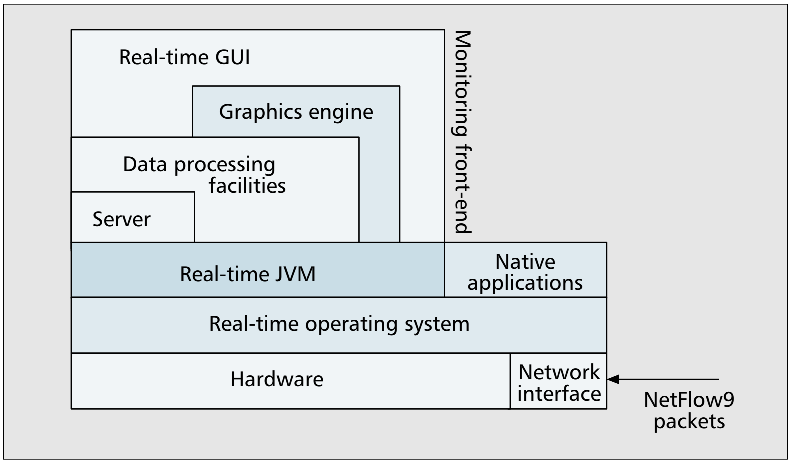

Figure 2. *Collector architecture.*

Sampling Techniques

Samplers may experience difficulties in reporting to the collector information about all the packets they capture; sampling is employed to surpass this problem [6]. By selecting one of N packets, the information sent by the samplers reduces to $1/N$th, and errors can be estimated for a given confidence interval. By increasing the sampling rate, the measurement error can be reduced. However, by being oriented to flow and QoS, the monitoring system may also need to address trajectory sampling: all the samplers traversed by a flow must sample *the same packets* so that, for instance, delays and losses can be evaluated. The monitoring system implements different techniques for TCP and UDP flows.

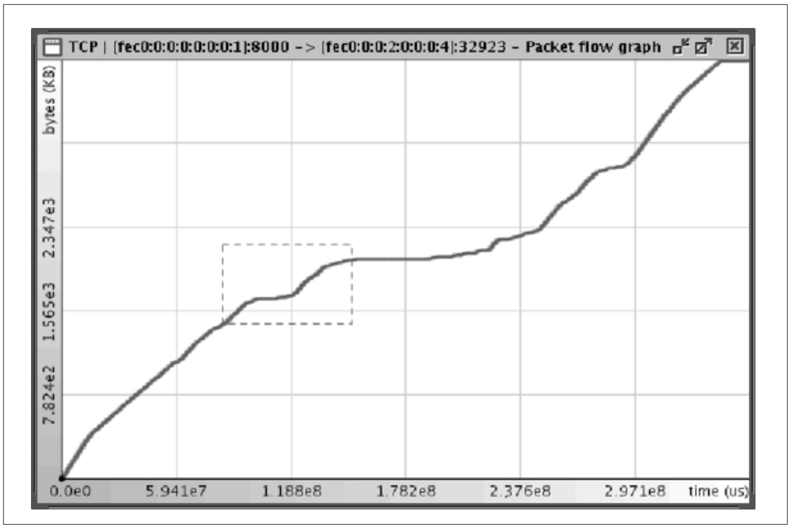

Figure 3. *Arrival/departure curve of a TCP flow.*

For TCP flows three sampling techniques were developed:

- Time-based selection: A packet is sampled whenever a timer expires.
- 1 in *N* systematic selection: The first of *N* captured packets is sampled.
- Synchronous stream selection: Packets containing a payload byte multiple of *N* are sampled.

This last sampling technique takes advantage of TCP flow packets that carry information about the stream octets transported: sequence number in the TCP header, payload length in the IPv6 header; so all the samplers can decide about whether or not to sample a packet.

For UDP flows, one sampling technique was developed: hash-based selection, in which packets with a common hashing pattern are sampled. For UDP flow packets, further processing is required since the stream sequence number concept does not exist. Now, UDP packets are hashed with a common algorithm, SHA1 or MD5, resulting in a small length value that becomes the packet signature. There is a probability of the wrong packet being sampled. If a 32-bit hash length is used, this probability is $1/2^{32}$, which is low. Needless to say, for both TCP and UDP flows the samplers must be preconfigured with rules describing the sequence numbers and hash values to be sampled.

Both synchronous stream selection for TCP and hash-based selection for UDP enable trajectory sampling.

THE COLLECTOR

The collector's role in the monitoring system is to gather, process, and store data received from samplers. It also provides a graphical interface, developed in Java language, to help a user configure the system and analyze the results.

From the architecture point of view, the collector consists of four modules, shown in Fig. 2: server, data processing facilities, graphics engine, and graphical user interface (GUI). The *server* module enables the samplers to concurrently communicate in a client-server service model, and use authentication; mainly, the server receives data from the samplers and builds the global information structure. *Data processing facilities* are placed above the server; they are responsible for:

- Maintaining the global flow structure
- Calculating on-demand performance metrics
- Processing the global flow information in real time
- Session logging and synchronization

The *graphics engine* is used to help the display of the overall interface and manage the drawing process. The GUI provides an interface to users.

From the functional point of view, the collector is responsible for a set of features that includes UDP packet pseudo sequencing, real-time information update, flow arrival/departure curve plotting, topology view, and metrics and statistics computation.

UDP packet pseudo sequencing — The collector assumes UDP flows as a byte stream and creates pseudo sequence numbers for the UDP packets collected. This enables the plotting of true arrival/departure curves for UDP IPv6 flows, with *in loco* packet loss visualization. For each flow, this curve depicts the accumulated bytes observed at a network interface and represent thems along time. TCP flow does not require pseudo sequence numbers, since sequence numbering is already produced by the TCP flow source machine.

Real-time information update — The collector displays the network activity per flow, and detects real-time flow mobility. For this to be possible, the collector needs to know at any time which network interfaces are active and what flows are passing through. Also, network interface throughputs, packet statistics, active COAs, COA history, and active link delays need real-time update.

Flow arrival/departure curve plotting — These curves are built based on trajectory packet capturing, which allows a packet to be identified throughout all the interfaces it traverses and mapped to the same curve level. It includes a facility to zoom out curve details, and gather information such as instantaneous throughput, bytes transferred, and delays.

Topology view — This enables the real-time presentation of the IPv6 network topology and flows traversing it. The network is modeled as a set of subnetworks, each characterized by an IP address and mask size. Samplers are connected to subnetworks by one or more interfaces. Simply said, a subnetwork is represented as a horizontal line. Each sampler is then connected to its networks by non-horizontal lines, as shown in Fig. 4. In order to keep the topology clean, samplers are split in two groups:

- Samplers with multiple interfaces, connected to multiple subnetworks and capable of forwarding traffic
- Samplers with only one interface, connected to a single subnetwork

The topology view adapts dynamically to network changes. It also detects host and flow mobility, and can display the path of all the active network flows.

Metrics and statistics computation — The collector can evaluate:

- The number of collected and transmitted packets per flow
- The average payload size per flow
- The packet loss per UDP flow
- The packet retransmissions per TCP flow
- Instantaneous data rates and packet delays, per flow and per sampler interface
- The real-time L2 + L3 header overhead
- The total number of captured TCP and UDP packets
- Average, current, and maximum interface throughputs, aggregated per sampler, per interface, and per TCP and UDP flows
- Maximum, average, and minimum packet delays between all sampler pairs, alongside delay standard deviation
- Source to destination delay, plotted with time
- Connection times, COAs history, and active link delays for each MN in the access network

Monitoring System Evaluation

In order to evaluate the monitoring system, a wireless access network supporting MIPv6 was tested in some scenarios. The flows captured were audio or video streaming, transported by TCP (using the HTTP protocol) or UDP (using the RTP/RTSP protocols). In these scenarios, the MN communicates with an infrastructure network through wireless 802.11 access points.

Scenario 1: The MN Moves Away from Home and Loses AP Connectivity — TCP Flow

In this case, the MN moves away from its home and eventually loses contact with the access point. Then it comes back and regains contact with the network. Figure 3 shows the TCP flow behavior, and represents the sequence numbers against time. In the red square, the TCP layer reduces the data rate when the MN gets far from home — high frame error ratios, and TCP closes its congestion window. When the MN regains connectivity, the data rate returns to normal value.

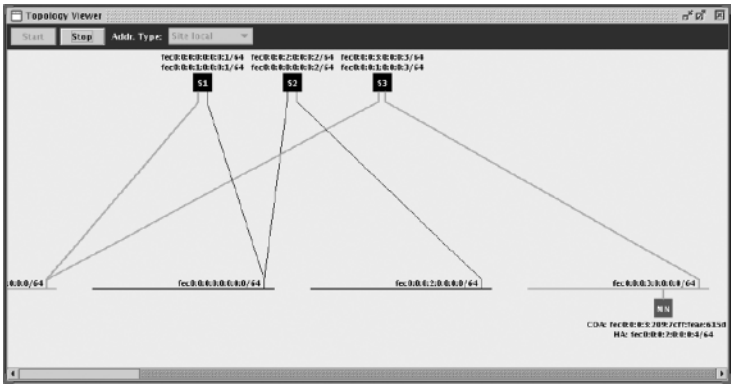

Figure 4. *Network topology: the MN in a foreign network (red rectangle).*

Scenario 2: The MN Roams to a Foreign Network — UDP Flow

When the MN roams to a foreign network an increase in header overhead is detected. Now, two extension headers circulate with the IPv6 packets — Routing and Destination Options. The topology viewer detects the flow and host mobility, and displays them accordingly on the

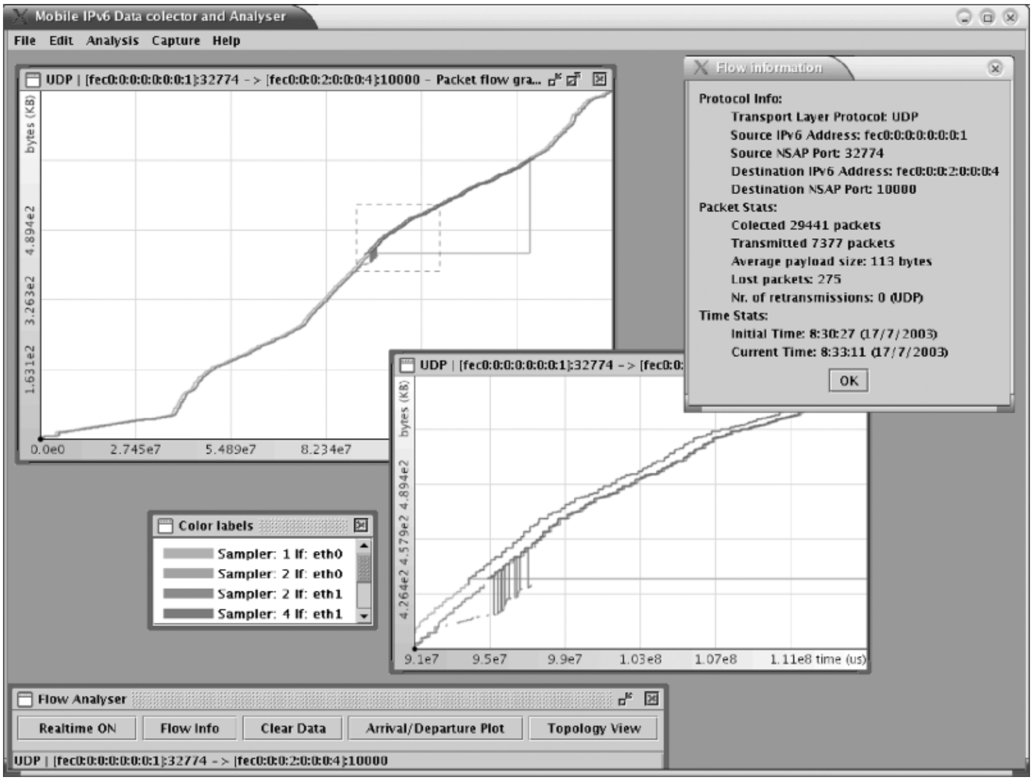

Protocol Info:
Transport Layer Protocol: UDP
Source IPv6 Address: fec0:0:0:0:0:0:0:1
Source NSAP Port: 32774
Destination IPv6 Address: fec0:0:0:2:0:0:0:4
Destination NSAP Port: 10000
Packet Stats:
Colected 29441 packets
Transmitted 7377 packets
Average payload size: 113 bytes
Lost packets: 275
Nr. of retransmissions: 0 (UDP)
Time Stats:
Initial Time: 8:30:27 (17/7/2003)
Current Time: 8:33:11 (17/7/2003)

Figure 5. *Arrival and departure plots, flow statistics, and metrics.*

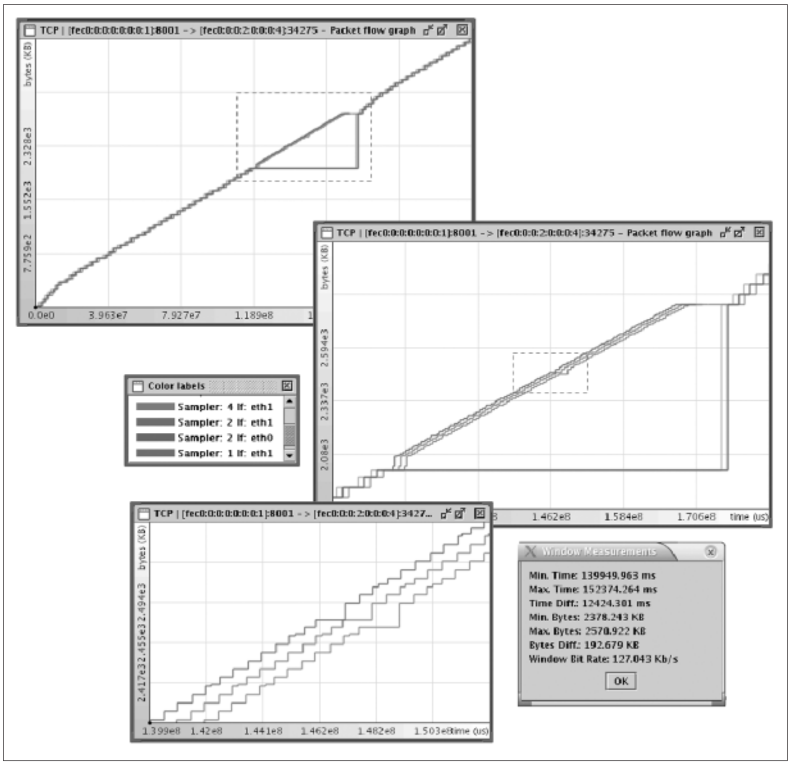

Min. Time: 139949.963 ms
Max. Time: 152374.264 ms
Time Diff.: 12424.301 ms
Min. Bytes: 2378.243 KB
Max. Bytes: 2570.922 KB
Bytes Diff.: 192.679 KB
Window Bit Rate: 127.043 Kb/s

■ **Figure 6**. *Arrival/departure curves for a sampled TCP flow.*

topology panel, shown in Fig. 4.

Figure 5 shows the collector GUI, including UDP flow arrival and departure plots. When the first handover occurs, traffic is redirected through other sampler interfaces (the red curves), which continue transporting the flow. A zoom has been performed to show a peculiar operation of UDP with IPv6 mobility. Just after the handover, when the source interface red line appears, a grey line wrinkle appears in the MN interface curve. This happens because before the binding update to the correspondent node (in this case the streaming server), the home agent gathers the packets destined to the MN and tunnels them to the foreign network; the problem is that after the binding update (route optimization) the packets are sent directly to the MN, but not all of the tunneled packets have arrived at the MN yet. This means that out-of-order packets will be received at the destination; these are represented by the lower ends of the grey wrinkle lines below the red and light green lines. This phenomenon appears only before route optimization. The GUI also presents flow-specific information, such as the number of transmitted packets, and UDP packets lost due to the handover.

Scenario 3: The MN Roams to a Foreign Network — TCP and UDP Flow Sampling

The curves shown in Fig. 6 represent a TCP flow detected using synchronous stream sampling. As shown, packet sampling adapts to variable throughputs, mobility conditions, and trajectory sampling. In this case, only packets having byte order number 50,000 or multiples of this value were sampled. The flow data rates represented are equivalent to those obtained without sampling. There is no noticeable loss of result significance.

A similar scenario was applied to a UDP flow using packet sampling based on a 5-bit-length hashing pattern, which means that only 3 percent of the packets will be used. All the samplers traversed by a specific flow were found to be synchronized and sampling the same packets.

Signaling Overheads

The sampler periodically sends flow information to the collector using the NetFlow9-based [5] format. This information is sent in a packet containing a fixed 80-byte overhead, plus 32 bytes per IP packet monitored. Let us assume that a 64 kb/s TCP flow is conveyed as 1500-byte IP packets, which gives 5.3 packets/s. If a sampler sends data to the collector every 5 s, 1.48 kb/s of monitoring data will be generated by each sampler, which represents 2.3 percent of the monitored flow. On the other hand, if flows are transported as very small IP packets (e.g., voice over IP), significant overheads may be observed that in the worst cases may amount to the values of the monitored traffic.

The solution envisaged to overcome this problem is the adoption of precisely the sampling techniques mentioned above. Using them, the signaling can be reduced to values as low as 3 percent of the original signaling traffic, which makes the technique also appropriate for small packet flows.

Using the Monitoring System

This monitoring system is aimed to be used in emerging 4G networks, such as that prototyped in [7]. It can be used as a traditional monitoring tool or as a 4G network operations tool.

When used as a traditional monitoring tool, the monitoring system will help operators to learn about their networks; 4G network traffic is unpredictable, and the first network elements and links will be deployed based on assumptions. The careful monitoring of traffic will help in understanding the nature of the traffic and evolve these networks by reviewing the number and locations of access points and access routers.

As a 4G network operations tool, the monitoring system can be in charge of detailed accounting functions, since it is capable of differentiating traffic by individual or aggregated flows; it can also be used to verify service level agreements between subscribers and their network providers.

Future Work

The centralized nature of the monitoring system, with a single collector, may lead to little fault tolerance or bottlenecks. In order to prevent them, the monitoring system needs to be improved with respect to two main aspects: a distributed collector structure and dynamic sampling.

In the distributed collector structure envisaged, a sampler is connected to a primary collector, but is also configured with a secondary collector, to be used if the primary collector fails. Collectors, on the other hand, can be intermediate or final in the data processing tree; in this approach, the number of samplers per collector can be kept small, if required, thus overcoming eventual scalability problems. Besides, the collector will be designed with the capability to take over the processing of another collector in case of failure.

Dynamic sampling may also be a relevant added value; under the explicit request of its controlling collector or when detecting overloads, samplers may be able to reduce the flow sampling rate, thus helping to avoid network congestion.

Support for IPv6 security headers and monitoring of signaling protocols, such as NSIS or Resource Reservation Protocol (RSVP), used to reserve resources for flows are also expected to be developed.

Conclusions

New wireless heterogeneous access networks may be unreliable because traffic constantly switches from path to path while the available data rates are limited by high frame error ratios and number of users. These networks tend to limit the quality of service offered and deserve special attention.

This article describes a new system capable of detecting and monitoring IPv6 mobile flows in these access networks. The system gives network operators useful information, including metrics and statistics. The monitoring system comprises sampler units distributed inside the access network and a central collector unit that processes the measurement data received from the samplers.

Major innovations of this system are: its capability to track real-time MIPv6 flows; network topology discovery along with dynamic flow path display; synchronized IPv6 flow trajectory sampling, independent of transport protocol; UDP packet sequencing; and graphical tools to analyze network performance.

This article describes a new system capable of detecting and monitoring IPv6 mobile flows in these access networks. The system gives network operators useful information, that include metrics and statistics.

References

[1] M. Buddhikot *et al.*, "Design and Implementation of a WLAN/CDMA2000 Interworking Architecture," *IEEE Commun. Mag.*, Nov. 2003.

[2] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," Dec. 1998, IETF RFC 2460.

[3] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6," IETF draft, June 2003.

[4] N. Brownlee and K. C. Claffy, "Understanding Internet Traffic Streams: Dragonflies and Tortoises," *IEEE Commun. Mag.*, no. 10, Oct. 2002, pp. 110–17.

[5] Cisco Systems, "Cisco IOS NetFlow Version 9 White Paper," May 2003.

[6] T. Zseby *et al.*, "Sampling and Filtering Techniques for IP Packet Selection," IETF draft, June 2003.

[7] V. Marques *et al.*, "An IP-Based QoS Architecture for 4G Operator Scenarios," *IEEE Wireless Commun.*, June 2003, pp. 54–62.

Biographies

Pedro Marques (pmarq@fe.up.pt) received a diploma degree in electrical and computer engineering from Porto University, Portugal, in 2003. In the same year he was a research collaborator at INESC Porto, working in next-generation wireless networks. He is currently pursuing a Ph.D. degree in electrical and computer engineering, also at Porto University. His main research interests are broadband mobile communications and MIMO channels.

Helder Castro (hcastro@inescporto.pt) received a diploma degree in electrical and computer engineering from Porto University in 2003. Currently, he works for INESC Porto under an R&D grant.

Manuel Ricardo (mricardo@inescporto.pt) received diploma (1988), M.Sc. (1992), and Ph.D. (2000) degrees in electrical and computer engineering from Porto University. Currently, he is an assistant professor at the Faculty of Engineering of Porto University, where he gives courses in mobile communications and computer networks. He also leads the Communication Networks and Services Area of INESC Porto.