Towards Network Controlled IP Traffic Offloading

Jouni Korhonen
Nokia Siemens Networks
Linnoitustie 6
02600 Espoo, Finland
Email: jouni.korhonen@nsn.com

Teemu Savolainen
Nokia
Visiokatu 3
33720 Tampere, Finland
Email: teemu.savolainen@nokia.com

Aaron Yi Ding, Markku Kojo
University of Helsinki
Gustaf Hällströmin katu 2b
00014 University of Helsinki, Finland
Email: [yding, kojo]@cs.helsinki.fi

***Abstract*—Operator controlled IP traffic offloading in cellular networks has been a lively topic in both product and standard development during recent years. Specifically, 3GPP[1] has developed multiple solutions for their system architecture. However, most IP traffic offloading solutions are complex, requiring 3GPP specific modifications on the system and typically also in mobile devices. We argue that an adequate IP traffic offloading solution is achievable entirely at IP level, resulting to a light-weight access technology agnostic offloading solution, which still can utilize 3GPP system properties for operators to push offloading policies securely into mobile devices. We present three variations of our IP traffic offloading implementations and compare them against 3GPP specified solutions. Our hands-on experience in operator networks shows that such direction is feasible and promising.**

I. Introduction

Searching for IP traffic offloading solutions has become topical for most cellular operators during recent years. The drivers are simple. The combination of throughput increase and latency decrease in cellular networks, and the popularity of smart phones has changed the diurnal behavior of cellular data users. The trend is towards always on applications with extensive use of online social media associated with frequent video, audio, and photograph rich content. As a result the growth in both user IP traffic volumes and core network side signaling has been phenomenal. This leads to huge investment pressures on cellular operator's IP transmission, radio and cellular core infrastructure while profits do not necessarily grow at the same pace. Operators have identified offloading bulk Internet traffic to alternative access technologies as a viable solution to relieve the infrastructure investment pressure. In many visions, the alternative access technology is a (public) Wireless LAN (WLAN). The current trend to equip smart phones with an additional WLAN radio supports such visions.

3GPP have put considerable effort to standardize IP traffic offloading solutions for the evolved packet core (EPC). All approaches rely on tight cellular operator control and integration to 3GPP network architecture. A new access network discovery and selection function (ANDSF) [2] is required to provide desired operator policies to the mobile devices [4]. It is questionable whether major mobile device and general purpose operating system (OS) vendors are willing to implement a 3GPP specific extension just to support IP traffic offloading solutions that oftentimes also make use of 3GPP specific radio features. Furthermore, it is to be seen what is the level of management burden the operators are eventually willing to take to micromanage IP traffic offloading policies.

This paper discusses existing traffic offloading solutions and presents and evaluates three different IP traffic offloading solutions that aim to be *IPv6 friendly* and rely on standard IETF[2] defined TCP/IP protocol suite. The goal is to identify working solutions that do not break backward compatibility, work in the majority of the intended use cases, and would not require mobile devices to implement 3GPP specific extensions. Our emphasis is on IPv6, since we believe there is no room for sophisticated IPv4-based offloading solutions anymore and the future is in IPv6 networking.

The rest of the paper is organized as follows. In Section II we discuss the background and current 3GPP-specified solutions for cellular network IP traffic offloading. Section III describes three different offloading solutions using IPv6 and Section IV compares the three solution proposal against 3GPP-specified solutions. Finally we give our conclusions on IP traffic offloading solutions in Section V.

II. Background

A. Motivations in Cellular Networks

Originally the 3GPP GPRS[3] architecture adopted point-to-point approach for the network access. This is visible, for example, in how the mobile device - the User Equipment (UE), sees the network connection, the network interface and how the point-to-point connection between the UE and the network is realized. The point-to-point connection between the UE and the external Packet Data Network (PDN) is referred to as a *PDP context* or a *PDN connection*. PDNs are identified using Access Point Names (APN). A full APN is in fact a Fully Qualified Domain Name (FQDN) that points to a gateway node connecting to a PDN such as Internet.

Monolithic UEs have a silo view of the network connectivity for a single application or a group of applications. Launching a new application either shares an existing PDN connection or creates a new parallel PDN connection with its own IP address, effectively turning the UE multi-homed. The host IP stack and the radio modem are typically tightly integrated. *Split-UEs* have a clear separation of the host IP stack and the radio

[1] The 3^{rd} Generation Partnership Project

[2] Internet Engineering Task Force

[3] General packet radio service

modem. They are usually even in physically separate devices. Hence, a Split-UE is just like any IP enabled host equipped with a wireless radio networking technology. All applications typically share the same single PDN connection. The split-UE model is increasingly becoming the dominant design.

The 3GPP Release-7 enhancements on the HSPA radio technology and architectural enhancements aiming for flatter network (e.g., direct tunneling) have significantly reduced the gap between the cellular and fixed broadband access. Moreover, the 3GPP Release-8 Long Term Evolution (LTE) for both radio and core network evolution has brought the cellular broadband access close to the mass market fixed broadband access (such as DSL) in terms of network latency and throughput. The increase in cellular broadband usability and the near flat rate billing models have had two notable outcomes: 1) consumers use cellular broadband in a similar way as fixed access and 2) the cellular Internet service providers experience an exponential IP traffic growth.

Heavy traffic growth poses investments pressures on radio access, backhaul and packet core capacity. Today 3GPP packet core network architecture and live network deployments tend to favor heavy IP traffic aggregation at Gateway GPRS Support Nodes (GGSN) or Packet Gateways (PGW) into relatively few sites. The IP traffic packet forwarding capacity of these gateway nodes might not sustain the traffic growth. In addition, the increased capacity investment requirements can be hard to meet due to the declining average revenue per user.

The issues discussed above have driven mobile operators to evaluate solutions to offload the bulk or low profit (Internet) traffic to alternative access technologies that are cheaper to deploy in dense hotspot areas and ideally would use someone else's backhaul. There has been a vision of using managed WLAN deployments or subscribers' home WLANs as alternative accesses. We can identify two technical motivations:

- *Compensate cellular radio coverage and access capacity* with a cheaper radio technology in a dense hotspot area but still routing the traffic through operator's packet core.
- *Bypass operator's packet core, cellular access and possibly backhaul* completely to maximize "savings".

In this paper we assume the latter one, although, our proposed solutions could apply for both.

The development on UEs strengthens these IP offloading visions, since even most mid range UEs have a WLAN radio. There are few issues left though. First, a UE may not allow operating multiple radios simultaneously, which effectively prohibits selective offloading of IP traffic between access technologies. In this paper we assume a UE can operate multiple radios in parallel. Second, how to determine which IP traffic to offload and which traffic to route through the mobile operator's core. Third, how to steer the offloading decision making from the network side. This can be challenging especially in the case of split-UEs. Fourth, how to minimize the impact on the operator network and especially in the UE.

B. Existing Solution Approaches

3GPP have worked on multiple offloading solutions for their network architecture, especially starting from the 3GPP Release-9:

- Multiple Access PDN Connection (MAPCON) [1] allows for a UE to connect to different APNs (and therefore PDNs) simultaneously via a 3GPP access (such as LTE) and a non-3GPP access (such as WLAN). MAPCON essentially enables traffic offloading from the cellular access to alternative radio accesses but the traffic still gets routed through operator's core network and GGSN/PGW. The routing rules and policies [4] for offloaded traffic can be achieved via ANDSF. MAPCON requires support from both core network and UEs.
- Selective IP Traffic Offload (SIPTO) [3] allows for a packet core to select a GGSN/PGW topologically or geographically close to a UE, thereby offering more optimized routing of IP traffic and hopefully less aggregation of traffic to few gateways. When the UE moves too far away from the associated GGSN/PGW the network may force re-establishment of the PDN connection and select a more optimal GGSN/PGW. The connection re-establishment causes disconnection in IP connectivity. SIPTO requires support from core network and optionally from UEs.
- Local IP Access (LIPA) [3] allows for the use of Local PGW (LGW) located in a Home (evolved) - NodeB (H(e)NodeB). The concept is similar to SIPTO. IP addressing used by the LGW is local e.g., to a corporate office premises, and managed by the LGW/H(e)NodeB rather than the home operator. The change of the LGW results in re-establishment of the PDN connection and a disconnection in the IP connectivity. LIPA requires support from core network and optionally from UEs.
- IP Flow Mobility (IFOM) [5] extends Dual-Stack Mobile IPv6 (DSMIPv6) to flow-based mobility. IFOM is a superset of MAPCON, and any IP flow can be moved selectively to any available access. The routing rules and policies [4] for offloaded traffic can be managed by ANDSF, for example. IFOM offers IP address preservation while switching the point of attachment to the network. IFOM requires support from both core network and UEs.
- Non-seamless WLAN offloading is the simplest form of 3GPP solutions. Certain traffic, which is possibly identified by the routing rules and policies [4] managed by the ANDSF, is directly routed to a WLAN access and allowed to bypass the operator packet core and the cellular access. This solution reflects what a multiple interfaces UE can do today.
- S2a Mobility based on GTP & WLAN access to EPC (SaMOG) [6] integrates managed WLAN access into the EPC. SaMOG essentially allows for offloading traffic from the cellular access to WLAN, while the traffic still gets routed through operator's core network and

GGSN/PGW. The integration is rather complex and requires 3GPP specific functions in the WLAN access network and core network gateways. SaMOG does not offer IP address preservation while switching WLAN and 3GPP accesses in 3GPP Release-11.
- S2b solution [2] allows for accessing EPC over an IPsec tunnel from any non-3GPP access and could also be used for MAPCON purposes. S2b essentially enables traffic offloading from the cellular access to WLAN, while the traffic still gets routed through operator's core network and GGSN/PGW. S2b solution requires support from both core network and UEs. S2b may or may not offer IP address preservation.

The offloading related Operator Policies for IP Interface Selection (OPIIS) [4] can be managed by the ANDSF, by a local configuration or by OMA Device Management (DM). Strictly from the IETF protocols point of view, enablers for IP traffic offloading has always been there when multiple interfaces/PDN connections and/or access technologies have been available. For example, there is no reason why IPsec traffic selectors could not also be used for offloading certain traffic to local access network when IPsec is used in untrusted non-3GPP access.

From research community, there are solutions that depend on new non-existing infrastructure support. Studies at MIT AI lab revealed that grassroots WLAN connections are viable for variety of applications [7]. Measurement studies in South Korea show that WLAN offloading without delayed transfer can offload 65% of total traffic [14]. The MADNet proposal utilizes both WLAN and Ad-hoc communication to enable flexible and efficient offloading in metropolitan areas [12].

C. Issues in Existing Offloading Approaches

We categorize the issues of the existing offloading approaches as follows:

- Heavy system specific standardization. This inevitably postpones deployment and ties it too tightly to a specific architecture.
- The micromanagement of the offloading policies. We believe that flow granularity leads to unnecessary management burden and a reliable identification of flows is challenging due encryption.
- 3GPP specific functionality in UEs. We believe that solutions should only be at the IP level and therefore independent of the access technology.

We advocate that the IP offloading should be considered just as a "normal" IP routing and next-hop selection issue with minimal management overhead. As long as the Internet connectivity is guaranteed, independent of the used network access, and the few required specific operator services can be reached, then it is possible to view the IP traffic offloading just as a normal IP level routing and source address selection problem that virtually every IP stack needs to support without access specific enhancements.

III. IP-Friendly Offloading Solutions

A. Introduction and Common Design Choices

We studied and implemented three IP-friendly approaches to achieve IP traffic offloading solution for multi-interfaced UEs with network side control for the offloading policies. The first approach builds on top of DHCPv6. The second approach builds on top of IPv6 neighbor discovery (ND) protocol and the third approach extends the second solution with IPv4 capabilities.

The IP-friendly solutions try to conform to a "pure IP" view and have specifically designed split-UEs in mind. None of the solutions aim at guaranteeing that the offloading policy provided by the network would work in all possible cases. More important is that the UE always has the Internet connectivity, meaning none of the used access networks shall be a walled garden.

The three solutions have several aspects in common:

- They rely on IPv6 features when possible. No 3GPP-specific extensions are required.
- They primarily target at UEs with cellular 3GPP access. The GGSN/PGW is used as the orchestra leader in the operator network.
- Cellular 3GPP radio is considered a *trusted* access and hence used to deliver offloading policies.
- Designed to benefit from multiple interfaces.
- The offloading policies are typically form of *offload everything except a few selected destinations*.

B. New DHCPv6 options

IETF has recently standardized how multi-interfaced nodes can make decisions to which recursive DNS server DNS requests should be sent [15]. Furthermore, IETF has worked on DHCPv6 extensions to provide more specific route information [9]. While the IETF discussions were ongoing, we set up a test network shown in Figure 1 that provided dual-stack Internet connectivity via WLAN and cellular access through Nokia Siemens Networks laboratory.

Internet Systems Consortium's (ISC) DHCPv6 servers were installed on both access networks. The DHCPv6 server on the cellular access was modified to accept unicast DHCPv6 messages, as multicast packet delivery was not available. In commercial deployments the DHCPv6 server would reside at the GGSN and hence would be capable of multicast reception.

We used Linux-based Nokia N900 as the UE. On top of the IPv6 enabled N900 we ported and implemented support for the recursive DNS server selection and the DHCPv6 specific route rules. The list of changes is the following:

- Integration and modification of ICS's `DHCPv6 4.2.0` client to support unicast DHCPv6 messages for requesting the recursive DNS server selection option as specified in Internet-Draft *draft-savolainen-mif-dns-server-selection-04* (predecessor of [15]) and DHCPv6 route options as specified in Internet-Draft *draft-dec-dhcpv6-route-option-05* (predecessor of [9]).

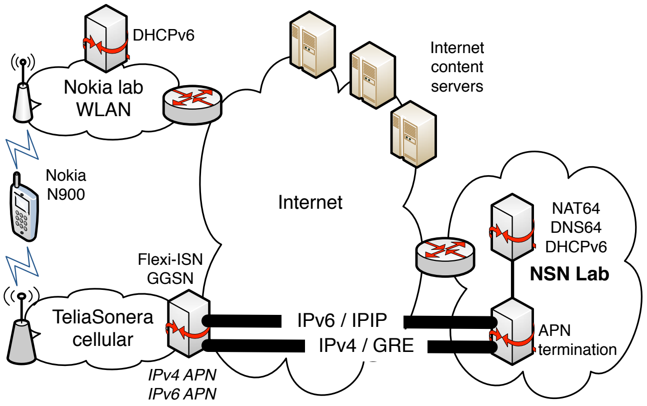

Figure 1: Cellular and WLAN test network architecture

- We replaced the DNS resolver (`dnsmasq`), which comes with N900, with more capable ICS's `BIND 9.7.1-P2` DNS resolver. All required behavior was available through configuration file (`named.conf`) modifications.
- We implemented `DNS Server Selection Configuration Generator` to create `named.conf` for BIND based on the DHCPv6 recursive DNS server option.
- We implemented `IP Route Configuration Generator` to install IPv6 routes based on the DHCPv6 route option.
- Set of scripts were created to place N900 into multi-interfaced state with both cellular and WLAN simultaneously enabled, as by default only one interface is active at a time. These scripts also triggered activation of BIND, DHCPv6 client, and Configuration Generators.

The 'Improved Recursive DNS Server Selection for Multi-Interfaced Nodes' recommends using secure and trusted channel and/or DNSSEC to counter against possible attackers [15]. In this prototype we did not include support for DNSSEC, as the cellular access can be considered both secure and trusted.

At the end, the implementation on the N900 was a rather simple task. Figure 2 illustrates a packet capture of a DHCPv6 Information Reply message with the new options (codes 92 and 93 were selected for prototyping, as official codes were not available at the time).

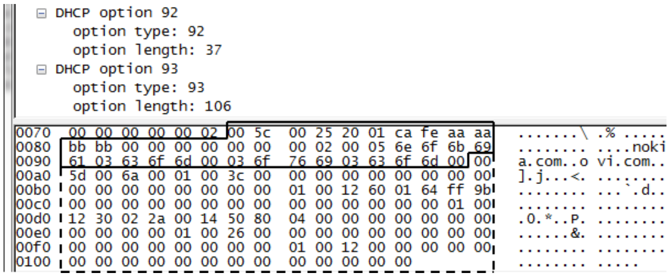

```
⊟ DHCP option 92
    option type: 92
    option length: 37
⊟ DHCP option 93
    option type: 93
    option length: 106
0070  00 00 00 00 00 02 00 5c  00 25 20 01 ca fe aa aa   .......\ .%.....
0080  bb bb 00 00 00 00 00 00  00 02 00 05 6e 6f 6b 69   ........ ....noki
0090  61 03 63 6f 6d 00 03 6f  76 69 03 63 6f 6d 00 00   a.com..o vi.com..
00a0  5d 00 6a 00 01 00 3c 00  00 00 00 00 00 00 00 00   ].j...<. ........
00b0  00 00 00 00 00 00 00 00  01 00 12 60 01 64 ff 9b   ........ ...`.d..
00c0  00 00 00 00 00 00 00 00  00 00 00 00 00 00 01 00   ........ ........
00d0  12 30 02 2a 00 14 50 80  04 00 00 00 00 00 00 00   .0.*..P. ........
00e0  00 00 00 00 01 00 26 00  00 00 00 00 00 00 00 00   ......&. ........
00f0  00 00 00 00 00 00 00 00  01 00 12 00 00 00 00 00   ........ ........
0100  00 00 00 00 00 00 00 00  00 00 00 00 00            ........ .....
```

Figure 2: Wireshark capture of recursive DNS server selection (solid line) and specific route (dottet line) DHCPv6 options

C. Default Router and More-Specific Route Selection

The IETF standard *"Default Router Preferences and More-Specific Routes"* [8] (RFC4191) extends IPv6 ND with two router preference flags in the Router Advertisement (RA) message header, and a *Route Information Option* (RIO). The former allows for a simple three step prioritization of default routers in host's default router list (*LOW, MEDIUM - the default, and HIGH*). The latter allows for an RA emitting router, even if not willing to be included into host's default router list, to mark up to 17 IPv6 destinations that the router wants serve as the first-hop. RFC4191 also can be deployed in a multiple interfaces scenario. The only consideration is to limit the number of interfaces accepting RFC4191 extension to one interface that is both trusted and centrally managed by the operator. In our case the 3GPP cellular connection fulfills these requirements. Several mainstream OSes already implement RFC4191, including Linux, BSD variants and Microsoft Windows starting from XP. The interface(s) accepting RFC4191 extension can be specified just using host side configuration.

3GPP architecture relies on IPv6 Stateless Address Auto Configuration (SLAAC) for its (un)trusted 3GPP access. The GGSN/PGW must always send RAs for SLAAC purposes and therefore using the 3GPP access as a command channel for ND-based offloading purposes is a small enhancement to the GGSN/PGW functionality. Furthermore, RAs can always be sent unsolicited. The ND-based solution is both extremely lightweight and allows for on demand push mode of operation from the cellular operator point of view.

Since modifying a live network GGSN/PGW was not an option, we implemented required tools in the APN termination router (see Figure 1). When the APN is terminated to an external router, the GGSN/PGW essentially becomes a bridge and the APN terminating router is the first-hop router for UEs. We implemented a `ndsend` tool that allowed us to send RAs with RFC4191 support, among other various ND messages, from the network to a specific UE.

Host OSs typically prefer WLANs over cellular access. For example, Linux implicitly prefers WLAN access over cellular access, thus prioritizing the first-hop router(s) on WLAN access over the first-hop router on a cellular access. For consistency we should always use *LOW* default router priority on the cellular access. Since the default router priority on other interfaces is implicitly *MEDIUM*, the host IP stack will prefer any other interface for default destinations than cellular. When the cellular operator wants to route certain traffic over the cellular, it only needs to send an RA with an RIO containing the IPv6 prefix(es) for those destinations (e.g., prefix(es) used to number operator's own services). The default router and default address selection algorithm [11] in the host IP stack will take care of selecting appropriate interface for new IPv6 connections (e.g., existing TCP connection will not move). When the model of operation is *"offload everything except specific destinations"* the number of routing rules can be kept low.

D. Enhanced Neighbor Discovery with IPv4 Support

To overcome the limitations of the IPv6-only RFC4191 approach (see Section III-C), we propose new IPv4 traffic specific RA options [13]. The new options enable access routers to convey IPv4 default gateway address and more-specific IPv4 routes.

Following our design principle to minimize the impact on host systems [10], we implemented a system prototype based on Linux with kernel version 3.0. The system architecture is illustrated in Figure 3. We extended the existing kernel implementation of *ndisc.c* by adding an intercepting hook function (*kernel offload hook*) and one module (*kernel offload module*) to push more-specific IPv4 route and the IPv4 default gateway address RA options from kernel to user space via the *sysfs interface*. In the user space, an *IPv4 offload daemon* handles main tasks for IPv4 traffic offloading by manipulating IPv4 routing tables. Figure 4 shows a captured RA message, where the RIO option carries an IPv4-mapped IPv6 address (`::ffff:203.178.141.0/120`) to route an IPv4 subnet to the RA originating router. The RA also carries the default IPv4 gateway address of the dual-stacked router as the last option (`10.6.6.6` i.e. `0a 06 06 06`).

All this was needed because currently Linux IP stack sends only few selected RA options into the user space, RFC4191 options are processed within kernel and IPv4-mapped IPv6 addresses in RIOs do not affect IPv4 routing, and we had no other way of changing the IPv4 default gateway on demand. Note that our proof-of-concept implementation was based on the earlier version of [13]. The current version of [13] is independent of RFC4191 and therefore removes the need for our own hook function and module. The existing kernel method for pushing RA options into the user space could be reused with few lines code change. We recognize that our RFC4191 extension for IPv4 more-specific routes would need a major push in standardization to reach wider acceptance.

IV. Solution Comparison

Table I summarizes the key characteristics of the 3GPP-based solutions and our three IP Friendly offloading solutions.

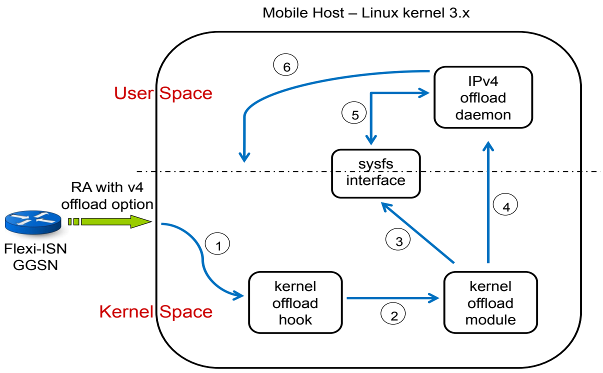

Figure 3: System overview for IPv4 offloading

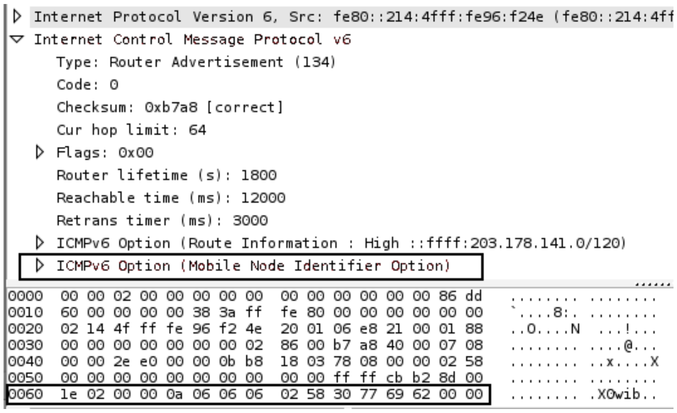

```
▷ Internet Protocol Version 6, Src: fe80::214:4fff:fe96:f24e (fe80::214:4ff
▽ Internet Control Message Protocol v6
    Type: Router Advertisement (134)
    Code: 0
    Checksum: 0xb7a8 [correct]
    Cur hop limit: 64
  ▷ Flags: 0x00
    Router lifetime (s): 1800
    Reachable time (ms): 12000
    Retrans timer (ms): 3000
  ▷ ICMPv6 Option (Route Information : High ::ffff:203.178.141.0/120)
  ▷ ICMPv6 Option (Mobile Node Identifier Option)
0000  00 00 02 00 00 00 00 00  00 00 00 00 00 00 86 dd   ........ ........
0010  60 00 00 00 00 38 3a ff  fe 80 00 00 00 00 00 00   `....8:. ........
0020  02 14 4f ff fe 96 f2 4e  20 01 06 e8 21 00 01 88   ..O....N  ...!...
0030  00 00 00 00 00 00 00 02  86 00 b7 a8 40 00 07 08   ........ ....@...
0040  00 00 2e e0 00 00 0b b8  18 03 78 08 00 00 02 58   ........ ..x....X
0050  00 00 00 00 00 00 00 00  00 00 ff ff cb b2 8d 00   ........ ........
0060  1e 02 00 00 0a 06 06 06  02 58 30 77 69 62 00 00   ........ .X0wib..
```

Figure 4: A packet capture showing a RA with RIO carrying a IPv4-mapped IPv6 address and the IPv4 default gateway address (which Wireshark interprets as Mobile Node Identifier Option since we just reused one unimplemented option type)

In the table '*IPv4/IPv6*' indicates the IP version the offloading tool applies to. '*Dynamic*' characterizes the nature of offloading and policy information dynamics, indicating whether they are easily updated during the ongoing session. '*Push/Pull*' indicates whether the network pushes or the UE pulls the policy down to the UE. '*UE*' indicates whether the solution involves UE functions that are specific to 3GPP architecture. '*Core*' indicates whether 3GPP packet core nodes and 3GPP protocols are impacted or not. The notion 'IP' tells that only user layer IP is impacted but obviously something has to be implemented in the first hop router i.e., the GGSN/PGW. Last, '*Offload*' shows what the offloading targets at: 'R' means the radio access and 'C' means the core network.

Table I: Solution comparison

Tool	IPv4/v6	Dynamic	Push/Pull	UE	Core	Offload
SIPTO	Both	No	Push	Yes	Yes	C
LIPA	Both	No	Push	Yes	Yes	R,C
Non-seamless	Both	No	Push	No/yes	No	R,C
MAPCON	Both	No	Push	Yes	Yes	R
SaMOG	Both	No	Push	No/yes	Yes	R
IFOM	Both	Yes (DSMIP)	Pull(/push)	Yes	Yes	R
S2b	Both	Yes (IKEv2)	Push	Yes	Yes	R
DHCPv6 route	IPv6	UE inits	Pull	No	"No"	R,C
RFC4191	IPv6	Yes	Push	No	IP	R,C
RFC4191+IPv4	Both	Yes	Push	No	IP	R,C

As of today, the use of ANDSF cannot be considered a dynamic interface for frequent policy updates on ongoing sessions. Therefore, all tools that are dependent on the ANDSF (and DM) function are not considered dynamic.

Non-seamless WLAN offloading is not really 3GPP specific as already mentioned in Section II-B. However, if the offloading decision depends of the presence of the ANDSF function, it is considered 3GPP specific. SaMOG-based offloading is the same to the non-seamless WLAN offloading from the UE perspective. In case of IFOM the traffic policies are part of the DSMIPv6 protocol signaling, and policies are mainly requested by the UE, thus *pull*. However, if the ANDSF (and

DM) is present, then policies can also be pushed from the network instead of just authorizing them. The S2b solution is not really an offloading solution, however, technically it could be used for that purpose.

In the case of DHCPv6, the DHCPv6 client in the UE has to initiate a DHCPv6 protocol exchange in order for the network to push new routes and policies to the UE. However, if network side reconfiguration feature is enabled, then the network can trigger the UE to perform a generic reconfiguration of its IP configuration. Since both ends have to agree on the reconfiguration support, we still consider the DHCPv6 solution as *pull*. If a GGSN/PGW were to implement a DHCPv6 relay, then there in no impact to 3GPP packet core.

Both RFC4191 and RFC4191+IPv4 solutions build on top of the ND protocol, and piggyback policies in RAs. The RFC4191-based solutions are most lightweight and simplest to implement, specifically when the network side policies are also taken into account. The RFC4191 support can already be found in most mainstream OSes. Specifically, it is a good match in 3GPP networks, since SLAAC is the only mandatory IPv6 configuration method for 3GPP accesses. The real downside for RFC4191-based solutions is that there is no way for the network to know whether the UE supports the feature or not.

We can summarize all solutions as follows: the non-seamless offloading would be the most lightweight solution, since it is basically already implemented in most OSes. However, the dependency on 3GPP ANDSF makes its dynamics and simplicity questionable. DHCPv6-based solutions lack proper "push" features for the dynamic policy updates and also in 3GPP architecture DHCPv6 is not a mandatory function. RFC4191-based solutions integrate easily into 3GPP architecture and the network side policy push is easy to implement and especially lightweight. However, like in the case of DHCPv6, the RFC4191-based solution has no support for IPv4 traffic. Extending RFC4191 with IPv4 knowledge is easy on the network side but, unfortunately, requires UE side changes and successful standardization in order to get mainstream OS vendors to support it. Naturally, none of our solutions have an existing network management interface for delivering policies.

V. Concluding Remarks

We discussed 3GPP-specified IP traffic offloading solutions and presented three IP-friendly offloading variations and implementations that are intentionally made to operate only at the IP level and make use of IETF protocols. A cellular operator can take advantage of the cellular network connection as a secure command channel to push offloading policies into the UE, while still only using standard IETF protocols. We also compared our DHCPv6 and IPv6 neighbor discovery protocol based solutions against 3GPP standardized offloading solutions. The implementation experience and the comparison shows that IP level solutions using IETF-only technologies are feasible and lightweight to deploy both on the network side, and specifically on the mobile device.

One of the major challenges we faced was the support for IPv4 traffic offloading in operating systems. While modern IP stacks offer a rich feature set for IPv6 to implement offloading in a multiple interfaces device, there is no clean solution available for IPv4. Another significant challenge is the resistance faced in the IETF community for using DHCPv6 or RAs for delivering offloading policies, and the persistent battle between DHCPv6 and IPv6 neighbor discovery protocol for generic host configuration.

We believe that the final deployed IP traffic offloading solution will likely be a mixture of existing technologies standardized in 3GPP, in IETF and what modern IP stacks can do. It is unlikely that mainstream operating systems' IP stack would implement 3GPP specific technologies - some 3^{rd} party dialer software may then add those missing elements. For the future, however, we believe that it would be useful to do further research on how dedicated routing protocols could be adapted to provide routing information for the end nodes. Use of routing protocols might provide more scalable architecture, perhaps even providing improved multi-homing properties for the end nodes.

References

[1] 3GPP. Multi Access PDN connectivity and IP flow mobility. TR 23.861, 3rd Generation Partnership Project (3GPP), Feb. 2010.

[2] 3GPP. Access to the 3GPP Evolved Packet Core (EPC) via non-3GPP access networks; Stage 3. TS 24.302, 3rd Generation Partnership Project (3GPP), Sept. 2011.

[3] 3GPP. Local IP Access and Selected IP Traffic Offload (LIPA-SIPTO). TR 23.829, 3rd Generation Partnership Project (3GPP), Oct. 2011.

[4] 3GPP. Operator Policies for IP Interface Selection (OPIIS). TR 23.853, 3rd Generation Partnership Project (3GPP), Aug. 2011.

[5] 3GPP. IP flow mobility and seamless Wireless Local Area Network (WLAN) offload; Stage 2. TS 23.261, 3rd Generation Partnership Project (3GPP), Mar. 2012.

[6] 3GPP. Study on S2a Mobility based On GTP & WLAN access to EPC (SaMOG). TR 23.852, 3rd Generation Partnership Project (3GPP), July 2012.

[7] V. Bychkovsky et al. A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks. In *Proc. of ACM MobiCom 2006*.

[8] R. Daves and T. Thaler. Default Router Preferences and More-Specific Routes. RFC 4191, November 2005.

[9] W. Dec, T. Mrugalski, T. Sun, and B. B. Sarikaya. DHCPv6 Route Options. Internet-Draft draft-ietf-mif-dhcpv6-route-option-04, Internet Engineering Task Force, Feb. 2012. Work in progress.

[10] Y. Ding, J. Korhonen, P. Hui, T. Savolainen, S. Tarkoma, and M. Kojo. NAO: A Framework to Enable Efficient Mobile Offloading. In *Proceedings of ACM Middleware PDT Workshop*, December 2011.

[11] R. Draves. Default Address Selection for Internet Protocol version 6 (IPv6). RFC 3483, February 2003.

[12] B. Han, P. Hui, and A. Srinivasan. Mobile Data Offloading in Metropolitan Area networks. *ACM SIGMOBILE Review*, 14(4), 2010.

[13] J. Korhonen, T. Savolainen, and Y. Ding. Controlling Traffic Offloading Using Neighbor Discovery Protocol. Internet Draft draft-korhonen-mif-ra-offload-04, IETF, March 2012. Work in progress.

[14] K. Lee et al. Mobile Data Offloading: How Much Can WiFi Deliver? In *Proc. of ACM CoNEXT 2010*.

[15] T. Savolainen, J. Kato, and T. Lemon. Improved Recursive DNS Server Selection for Multi-Interfaced Nodes. Internet Draft draft-ietf-mif-dns-server-selection-09, IETF, May 2012. Work in progress.