Concept for Cooperative Traffic Management

Szilveszter Nádas, Attila Mihály
Ericsson Research,
Szilveszter.Nadas@ericsson.com, Attila.Mihaly@ericsson.com

***Abstract*—We envision that optimized resource sharing can be deployed in a way which benefits all players in the ecosystem. We argue that this requires cooperation between end-points and the network building on explicit communication. We demonstrate how solutions like this can be deployed, when the content and maybe also the transport protocols are encrypted.**

We argue for the need of incentive frameworks to achieve this cooperation. We introduce the concept of Trust and Policy Controller which can further ease the introduction of different incentive frameworks, giving the choice of QoS to the end-user, while minimizing user interaction during actual use of applications.

This position paper is intended to trigger discussion about these concepts and its main purpose is to provide a basis for exploration and discussions of this potential way of traffic management.

Keywords: Cooperative traffic management, incentive framework, QoE, Trust and Policy Controller

I. Traffic Management in Cellular Networks

Several cellular network operators have recognized that the growth of user data traffic outpaces the growth of cellular capacity [1]. Consequently it is more and more important to optimize the resource sharing in cellular bottlenecks. These bottlenecks are most often at the air interface, though these can also happen in the Mobile Backhaul connecting the Radio Base stations to the core network. In addition to resource sharing, Radio Resource Management algorithms can be optimized and also the redundancy between Transport Protocols and protocols in cellular network can benefit from that optimization [2].

Traffic management solutions can be cooperative and non-cooperative. Currently cellular networks predominantly deploy non-cooperative solutions. Cooperative solutions are less used today. They may need signaling between the end-points and the network, an example being the 3GPP Rx interface.

One category of non-cooperative traffic management is when there is no assumption of traffic characteristics and consequently the resources are shared among users in a fair (or predetermined) way in all situations. Examples include air interface scheduling, resource sharing control solutions in the mobile backhaul and Transport Protocol Performance Enhancing Proxies (TP PEP). TP PEPs may replace E2E congestion control (CC) with a CC suited for the given radio access. Many content unaware solutions can work the same way with encrypted traffic. However the settings of these might be still tuned by content of the flows and also some require access to higher protocol layers (e.g. TP PEP), which might be encrypted in the future.

Submitted to "Managing Radio Networks in an Encrypted World (MaRNEW) Workshop", organized by IAB and GSMA, Sept, 2015

There is a significant potential in optimizing the overall utility for both the network and the subscribers by allocating resources unevenly when there is a congestion event. For example, by giving more resources to a web download than to a background file transfer, the overall user satisfaction increases because the impact on user experience improves for the former without significantly affecting the latter. Current traffic management solutions that allow temporal deviations from the equal share of the different users generally use Deep Packet Inspection (DPI) to determine the content type and other meta-data of traffic flows. Such solutions may include an agreement with the content provider to transmit the data for lower cost (to the end-user) and/or with different QoS. Other solutions may be implicit based on the estimated QoS requirements of the content.

In addition to DPI being challenged by encryption there are several other challenges of the content aware solutions used today. DPI based traffic identification, even for non-encrypted content, might put some OTTs at disadvantage, because e.g. their traffic is not detected right or if its QoS requirements are not estimated correctly. This might happen, e.g. because that particular OTT is a smaller player, whose traffic was not tested when designing the DPI algorithm. Some of these improvements also raise net neutrality concerns: the implicit solutions are often not perceived by the end-users, or in a bad case these might be perceived as hostile due to the down-prioritization (or lack of up-prioritization) of a content actually deemed important to that particular end-user. Even the same solution can be perceived differently by different end-users, because they have different desires.

All above shows the need for traffic management solutions based on explicit cooperation.

II. Efforts targeting explicit cooperation

AEON/AECON (Application Enabled Collaborative Networking) was an attempt to start an IETF working group which can support "Identification and treatment of application flows" by flow related signaling. The three drafts [3] summarize several use cases and information elements also related to content aware QoS differentiation.

The goals of IAB IP Stack Evolution Program [4] include "Improving path transparency in the presence of firewalls and middleboxes: guidelines for the detection of and cooperation with these devices". The program resulted in several IETF WGs or WG proposals: TAPS, SPUD and HOPS.

The Substrate Protocol for User Datagrams (SPUD) is a WG proposal discussed at IETF. There is a wide range of opinions for the role of such a substrate protocol, in the

scale for "indication of session start and stop for NATs" to "possibility authenticated in-band signaling channels" with no clear consensus. The SPUD WG is not yet chartered.

When the question of content and/or treatment signaling was raised in the above activities people had several concerns. When signaling a treatment or a content type results in positive discrimination of the user, there is an incentive to lie, and it is hard to verify the statements, especially when the content itself is encrypted. Another concern often raised is that when the user is given choice regarding preferred treatment it requires too much interaction and decision form the user.

A way to address these issues is described in [5]. There are three constraints proposed on the exposed information. (1) Information exposure is declarative, (2) all entities may trust but verify and (3) information must be incrementally useful. Another constraint related to these was mentioned several times: (4) the exposed information shall not change the total share of the user (from a bottleneck resource). Constraint (4) highly decreases any incentive of the user to cheat, but it also makes QoS solutions much less efficient, because the networks must meet the QoS demand of the most demanding service for all users all the time. This might be possible in some networks (e.g. fixed line), but this is much more challenging for cellular networks.

III. Middlebox communication as part of the IP stack

As we saw earlier, when most of the content and of the protocols become encrypted, implicit traffic management will have a hard time to optimize the utility for the end-users and OTTs. In the cooperative scenarios traffic management must add demonstrated value to the end-hosts in exchange for the information provided. Also it must solve issues raised in the previous chapter. We designed this solution in a way that works with end-to-end content encryption and it minimizes the information leaving the device.

A way to demonstrate this value is to provide the right service for the right price. More demanding QoS and resource requirements can be met this way (compared to BE access), however there is a consequence e.g. in pricing or in general in usage policy. This consequence shall discourage users to not request demanding services, when these are not really needed, and at the same time it may allow new services and/or higher Quality of Experience for existing ones. Similarly, background transfer with smaller resource demands might be incentivized by more favorable usage conditions in that case.

The paper [6] provides recommendations how to design protocols and interfaces to design for the Tussle between actors, instead of designing for a desired outcome. The proposed list includes "Visible exchange of value", "Exposure of cost of choice", "Visibility (or not) of choices made" and "Tools to resolve and isolate faults and failures". We emphasize the importance of designing for the Tussle. That design results cleaner interfaces and is more transparent for all outcomes.

We demonstrated concepts for incentive frameworks which can help this cooperation [7], [8]. One such concept is to slightly change the widely used monthly cap concept and introduce several service levels with different effect on the monthly cap. Another one is to reward voluntary down-prioritization by enabling up-prioritization of the critical user traffic, but not changing the monthly cap. All of these frameworks require frequent decisions about the service level. The final decision about service level shall be made by the one paying for the service, which is often the end-user. In all of the framework proposals a possible strategy is to not provide any metadata and ignore any received state from the network.

The end-user should thus have the possibility to control which applications and when to use a specific service option. However, the end-user shall not be bothered too often by configuration and especially not by making decision during the use of applications. We propose an application named Trust and Policy Controller (TPC) to fulfill these apparently contradicting requirements, which can be configured to receive service and state information from other application, receive information from the network path and based on these decide and signal selected service levels. The rules of selection can be described by a database and it is the users choice to select (or create/modify) the database most fitting to his preferences. Such a database can be provided by the network operator, by the OS vendor, by the application store or by the community (similar to e.g. AdblockPlus filter databases [9]).

We believe that the incentive frameworks and the Trust and Policy Controller introduced in this chapter may potentially provide a way forward in this area. The demonstrated incentive frameworks are not meant to be comprehensive at the moment. They always have to be adapted to the specifics of networks, local regulations and popular apps. Like all QoS solutions, these frameworks raise net neutrality concerns. While regulation, including net neutrality related rules, is not only the task of the engineering community, we believe that engineers can help in this by demonstrating the advantages and fairness of these solutions and proposing ways to keep this fair and transparent.

TPCs will also not likely be deployed first; instead applications taking advantage of this functionality may implement related setting and behavior themselves. We envision that on the longer run however it is more advantageous to coordinate this from a dedicated application. This would result in advantages similar to those of [10] in the area of middlebox cooperation and QoS. Legacy applications could also be configured in TPC and information send over interfaces to TPC could be more easily adapted to the incentive framework used by the network the device is connected to. At the same time the interfaces to TPC can be quite verbose to be able to make informed decision, but at the same time the metadata the TPC communicates to the network could include the minimal amount of information needed to choose the right service treatment.

The richness of the metadata used during the decision and the minimal amount of information sent to the network based on this metadata is intended to minimize privacy impacts of the solution. The TPC in this sense is similar to firewall software on PCs today: it regulates what information can leave the device and it is critical to have a secure and trusted implementation in the long run to keep the privacy sensitive information at the device. In the Tussle the TPC is the agent

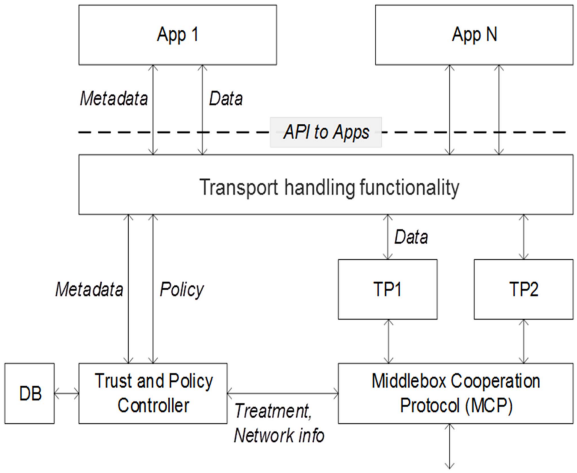

Fig. 1. Example device architecture

of the end-user, the objectives of the network operator shall be considered in the middleboxes receiving information from the TPC.

Example architecture in devices is shown on Figure 1. The applications (App) have an API towards a Transport Handling Functionality, which is very similar to the one discussed in [10]. This interface supports sending user data and metadata. The Transport Handling Functionality is responsible for selecting the right Transport Protocol(s) for the Applications based on the metadata. The Trust and Policy Controller (TPC) receives all Metadata from the Applications and, based on the one or more configured Databases (DB) that contain the rules of service selection described above, it communicates the traffic handling requirements using the Middlebox Cooperation Protocol (MCP) with Middleboxes. Middleboxes may also send e.g. Network information to the TPC. Based on this the TPC may further query applications about their state. In summary, the Metadata exchanged between TPC and APP through the Transport Handling Functionality may contain general session information (most likely at setup), APP state information and network state information. The TPC is responsible for removing all privacy sensitive Metadata before determining e.g. the preferred treatment of the session sent through MCP. The TPC may also aid the Transport Handling Functionality in Transport Protocol selection through the Policy API. The application and the TPC (and also the Middlebox) is free to discard any information or query received and they are also free to not send any initial information. The behavior of the TPC may be highly configurable to meet the desires of the different end-users.

The architecture on Figure 1 is an evolved state. In the first deployments of an MCP it is likely that all boxes will be implemented inside the application supporting MCP. We emphasize the importance of allowing user configuration also in this case (the analogy of the TPC DB). On the longer run the depicted architecture is advantageous because it provides a single consistent management opportunity, which may also support legacy applications not implementing any new API.

IV. Conclusion

In an end-to-end encrypted Internet a potential way to perform traffic management for congestion handling that benefits all players (end-users, cellular access providers, and content providers) is via a cooperative approach. This brings up the need for new solutions in the area of meta-information exchange: communication protocols, incentive frameworks to achieve this cooperation, trust and policy control in both sides.

In this paper we show how solutions like this can be deployed based on the current ecosystem. These solutions empower the end-user by giving him choice regarding requested treatment and allow using rich metadata during this decision, but not communicating privacy sensitive metadata to the middleboxes. The middleboxes shall take into account this requested treatment and service policy when configuring QoS solutions in the domain. The network operator also has a choice in this cooperation whether to accept the "requested treatment" or to choose different treatment due to policy reasons and to determine the usage policies governing the consequences of the different choices made. We introduce the concept of Trust and Policy Controller, an agent of the end-user in the Tussle, which can further ease the introduction of different incentive frameworks, giving the choice of QoS to the end-user, while minimizing user interaction. We exemplify an evolved device architecture embedding the TPC.

The solutions outlined in this paper are neither unique not complete. There are different evolution paths to arrive at similar architectures and we encourage the different actors (OS developers, (cellular) access providers, and content providers) to continue discussion and experiments in this area.

References

[1] 3GPP Work Item Description, User Plane Congestion management, S2-140513, 20-24 January 2014, http://www.3gpp.org/ftp/tsg_sa/wg2_arch/TSGS2_101_Taipei/Docs/S2-140513.zip
[2] Middleboxes in Cellular Networks, position paper at IAB SEMI WS, January, 2015, https://www.iab.org/wp-content/IAB-uploads/2014/12/semi2015_nadas.pdf, https://www.iab.org/activities/workshops/semi/
[3] https://tools.ietf.org/html/draft-conet-aeon-problem-statement-01, https://tools.ietf.org/html/draft-conet-aeon-use-cases-01, https://tools.ietf.org/html/draft-conet-aeon-gap-analysis-01 checked 2014-12-01
[4] https://www.iab.org/activities/programs/ip-stack-evolution-program/ checked 2015-06-10
[5] Substrate Protocol for User Datagrams (SPUD), Brian Trammell, slides at IETF92, 2015 http://www.ietf.org/proceedings/92/slides/slides-92-spud-1.pdf
[6] David D. Clark et all, Tussle in Cyberspace: Defining Tomorrows Internet, IEEE/ACM TRANSACTIONS ON NETWORKING, VOL. 13, NO. 3, JUNE 2005, https://impact.asu.edu/cse534fa06/reading/p462-clark.pdf
[7] Attila Mihály, Szilveszter Nádas, Enablers for Transport Layer Protocol Evolution, Internet-Draft, March 9, 2015, https://tools.ietf.org/html/draft-mihaly-enablers-for-tlp-evolution-00
[8] Attila Mihály, Szilveszter Nádas, Middlebox Communication Enabling for Enhanced User Experience, Internet-Draft, July 6, 2015, https://tools.ietf.org/html/draft-mihaly-spud-mb-communication-00
[9] Adblock Plus, https://adblockplus.org/, checked 2015-06-15
[10] IETF TAPS WG, https://datatracker.ietf.org/wg/taps/charter/, checked 2015-08-06