Future Network and MobileSummit 2012 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2012
ISBN: 978-1-905824-29-8

Socioeconomic Tussles Analysis of the ETICS Approach for Providing QoS-enabled Inter-domain Services

Costas KALOGIROS[1], Costas COURCOUBETIS[1], George STAMOULIS[1], Manos DRAMITINOS[1], Olivier DUGEON[2]
[1]*AUEB, Patission 76, Athens, 10434, Greece*
Tel: +30 210 8203653, Fax: +30 210 8203686,
{ckalog,courcou,gstamoul,mdramit}@aueb.gr
[2]*Orange Group, Pierre Marzin 2, Lannion, 22307, France*
Tel: +33 296 052880, Fax: +33 296 051470, olivier.dugeon@orange.com

Abstract:
We apply a tussle analysis methodology to two types of functionality offered by the ETICS framework. The first one is related to establishment of QoS-aware, inter-domain paths. We find that allowing fine-grained control over major properties of interconnection agreement (such as set of IP destinations) can help adoption of the ETICS system and sustain a healthy ecosystem for all participants. The second functionality investigated is service delivery with assured quality between multiple ISPs. Here we find that even if ISPs are honest in primary paths they may overbook backup paths. Thus, without adequate monitoring it is not possible to correctly assign responsibility for breaking the end-to-end SLA in case of failure along the path (where the backup should be used).

Keywords: Future Internet, Tussles, Design Principles, QoS, SLA

1 Introduction

Internet is a platform that can be studied as a system composed of multiple technologies and an environment where multiple stakeholders interact by using those technologies. At the socio-economic layer stakeholders have interests, expressed by making choices governed by laws of economics (e.g. supply and demand), sociology, psychology, etc. Such choices affect the Internet technology layer by specifying the technologies to be introduced, how these will be dimensioned, configured, and finally, used. This phenomenon makes Internet interesting not only to technology developers but economists and social scientists, as well.

For example the DNS (Domain Name System) associates domain names to network addresses in a distributed way, based on each administrator's configuration of the local DNS server's entries. Similarly, routers today forward data packets based on each ISP's configuration of the Border Gateway Protocol (BGP). To provide an example of how complex these interdependencies may be, content (and/or service) providers are usually multi-homed in order to assure redundant Internet connectivity, but also to optimize end-user's experience by performing advanced load balancing (e.g., at the application level by creating multiple DNS entries that each one directs user to different servers [1] and at the network level by announcing multiple, more specific BGP entries).

This natural quest of Internet stakeholders for achieving their interests defines a "tussle" instance [2]. Thus, a tussle does not involve the interests of the stakeholders only, but how these conflicting interests are expressed through the available technologies. The combination of actors' policies at a given point in time, which take into account other stakeholders' socio-economic decisions and the restrictions imposed by the current technology set, leads to a tussle outcome. These outcomes are rarely static and the emergence of new technologies and stakeholders, the adoption of new strategies, or introduction of new regulatory constraints can trigger the transition to new outcomes. All these interactions allow the Internet to evolve and act as a living organism.

There is a growing number of researchers in the Internet community suggesting that standardized technologies should have taken into account the incentives of all major stakeholders during the design phase. Clark et al. [2] have described two new design principles that implement the "Design for Tussle" goal. In a similar line of thought Kilkki [3] proposes three practical rules for avoiding the development of technologies that do not capture key aspects of their socio-economic environment.

SESERV coordination project [1] has defined a systematic approach for analyzing and assessing the importance of socio-economic tussles in the Internet [4]. The main idea is to make sensible predictions about the behaviour of major stakeholders in several scenarios, each scenario reflecting candidate implementations of the desired protocol functionality. We argue that selecting the features of a technology in a more holistic way, by taking into account the Internet socio-economic layer would lead to more attractive outcomes and increase the chances of that technology to be adopted in the long-term.

The purpose of the paper is to apply this methodology using as a case study the ETICS FP7 European research project [2], in order to clarify how each step may be executed and demonstrate its value. Section 2 provides a short introduction to the tussle analysis methodology and Section 3 describes the scope and basic services of the project being our case study. The main part of this paper consists of Section 4 and Section 5, where tussle analysis is performed for the establishment of QoS-aware, inter-domain paths and service delivery, respectively. Finally, conclusions are drawn in Section 6.

2 A tussle analysis methodology

The proposed methodology is visualized in Figure 1. Each step is shown as a horizontal rectangle with arrows denoting transitions. All steps are applied in the context of one, or more, functionalities (rounded vertical rectangles).

The first step suggests identifying and studying the properties of the most important stakeholders (their interests, technologies used, etc). Identifying alternative technology schemes will be useful in performing the second step, which refers to identifying tussles among the set of stakeholders. More specifically when a conflict of interest is found to exist among some stakeholders, we should seek for policies enabled by the technologies that these rational entities would select in order to meet their goals. The third step of the methodology aims to evaluate each tussle outcome from the perspective of each stakeholder (in order to infer the stability properties of the functionality under investigation) and understand its effects on the stability of other functionalities. In the ideal scenario of a tussle outcome, we have an equilibrium point where a) all stakeholders of that functionality derive a *fair* payoff (thus no one will select another policy) and b)

[1] http://www.seserv.org
[2] http://www.ict-etics.eu

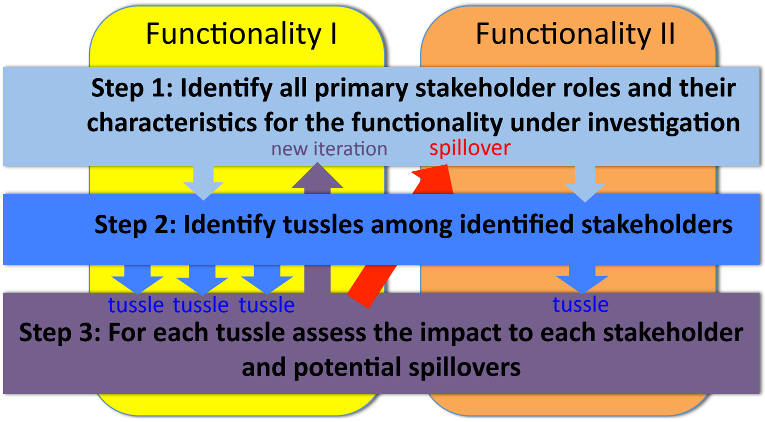

Figure 1: High-level view of tussle analysis methodology

no stakeholder of another functionality, who was receiving a fair payoff before, gets an unfair payoff after the new tussle equilibrium has been reached.

If both conditions hold, then the analysis of this particular tussle is completed and we can move on to the rest tussles of Step 2. In case condition (a) is not met, a new iteration of the methodology must be performed by making assumptions on the most probable policies adopted by unhappy stakeholders. Similarly, a new iteration must be performed for each spillover to other functionalities, when condition (b) is not met.

When designing Internet technologies, this methodology can be useful for understanding the expected impact to the stability and efficiency of that particular functionality and possible spillovers to other spaces. Each of the three steps can be performed by several techniques. For example, the third step could be based on game-theoretic models for determining the existence and evaluating the socio-economic properties of candidate equilibria.

3 Introduction to the ETICS approach

The ETICS solution enables the support of the emerging Internet QoS-sensitive high-performance services (e.g. HD video streaming, tele-presence, e-health) through network interconnections of assured quality, ensuring that the applications' Quality of Service (QoS) constraints will be met. Technically, ETICS automates the support of end-to-end (e2e) QoS guarantees across multiple networks. Economically, it serves as a market enabler for services that require QoS assurance, also allowing for a fair distribution of revenue sharing among the market stakeholders. In particular, ETICS supports premium inter-domain connectivity services by stitching or nesting connectivity agreements - called Assured Service Quality (ASQ) agreements or goods - from several ETICS providers.

The Internet currently offers Best Effort end-to-end services only, based on two interconnection market products (namely *peering* and *transit*). Each NSP accepts only BGP information from its neighbors (thus a single route per destination can be used for all traffic types) and classifies all incoming packets to low priority by ignoring any ToS (Type of Service) information in packet headers and RSVP signaling. This reflects both the lack of high level of trust among NSPs and the lack of compensation that should be expected for the extra effort of the NSP to provide such a premium service.

On the other hand, ASQ products provide tangible QoS assurance in terms of reliability, bandwidth, delay, jitter, etc over a certain ASQ path; this path is selected independently of the BGP in order to meet the QoS constraints demanded by the customer. Furthermore, although an ASQ is always associated to a specific service type (e.g. video) or a given destination / source, multiple ASQ products can be offered for the same destination. Thus, in contrast to existing solutions for premium connectivity services that are tailored to the needs of specific applications (such as IPX [3] and IMS [4] in the case of VoIP), ASQ goods provide more flexibility to NSPs.

In the ETICS context, service composition is the process of establishing the e2e path based on the technical parameters of the associated Service Level Agreements (SLAs) over a chain of ETICS network operators. In order to do so, participants must be aware of the available services/products and all the necessary information (prices, availability, etc.), which takes place through a service discovery mechanism. Other types of functionality explored are admission control for service establishment and SLA monitoring mechanisms for collecting the necessary information during service provisioning to validate conformance to contract terms.

SLA monitoring is another functionality where different NSPs currently find difficulties in collaborating. Apart from adopting incompatible monitoring technologies, NSPs may be interested in measuring different sets of QoS metrics. Furthermore, they are very reluctant to allow third parties have access to sensitive information about network topology, current network conditions, etc.

In what follows we provide examples of using the tussle analysis framework in the context of ETICS to predict and analyze potential tussles.

4 QoS interconnection and transit competition

The basic idea is that two competing Internet Service Providers in a regional market (e.g. the same country) may benefit by establishing an ASQ agreement appropriate for transporting and terminating each others (customers) traffic, since this will allow more lucrative e2e services requiring QoS to be provisioned. But this cooperation conflicts with an ISP's goal to improve the value of services to its own end-customers and becoming more competitive in the provisioning of such e2e services. We analyze this techno-economic phenomenon and using the principles for design for tussle show that if enough control is available in the definition of the ASQ good, then the spillover effects can be reduced and hence the ASQ agreement can function as originally envisaged.

Based on today's Internet market structure, let us suppose that a large operator, called ISP-1, has attached the cache of a popular Content Provider to its network and no financial transactions take place (similar to a peering agreement). Furthermore it has a *peering* link with ISP-2, which allows them to exchange their customers' traffic for free. A third operator, called ISP-3, buys *transit* connectivity from ISP-1 as a result of the higher quality connectivity to the Content Provider. All ISPs have a number of end customers, but ISP-1 has the largest market share, followed by ISP-2.

Performing the first step of the methodology for the QoS provisioning functionality (first iteration), would identify the ETICS Communication Service Providers (or ETICS ISPs for short) and Content Providers as the main set of stakeholder roles. Other in-

[3] www.gsma.com
[4] http://www.3gpp.org

volved roles are consumers of ICT services and Regulators, but due to space constraints we will focus on the first set only.

In today's Internet *peering* links are usually under-dimensioned. So in the second step we would find that ISP-1 has no incentive to upgrade the capacity of the peering link in order to maintain its competitive advantage over ISP-2 for communication providers that buy transit services. On the other hand, ISP-2 will try to improve the quality of the services offered to its customers by performing traffic engineering.

Thus, in the third step we would conclude that a desirable stable tussle outcome - the ability to offer e2e QoS services for any type user traffic - will not be reached unless the peered ISPs are symmetric in terms of volume exchanged and networking services supported. The problem is that transactions in the Internet connectivity market are characterized by severe information asymmetries. Asymmetric information in peering agreements gives rise to opportunistic behavior in different guises, such as "backbone free riding" [5] and "hot-potato routing". We will denote this initial state of QoS functionality in Figure 2 using a blue circle with dotted-line border to indicate that, in absence of an incentive-compatible QoS mechanism, it is an unstable one.

Furthermore, the inability to compose network services with QoS features has a spillover to the routing functionality (shown with a dotted red arrow). Performing the tussle analysis methodology for the routing functionality would identify the same set of stakeholders (Step 1) entering a loop of routing decisions, which utilize the peering link in a selfish way (Steps 2,3).

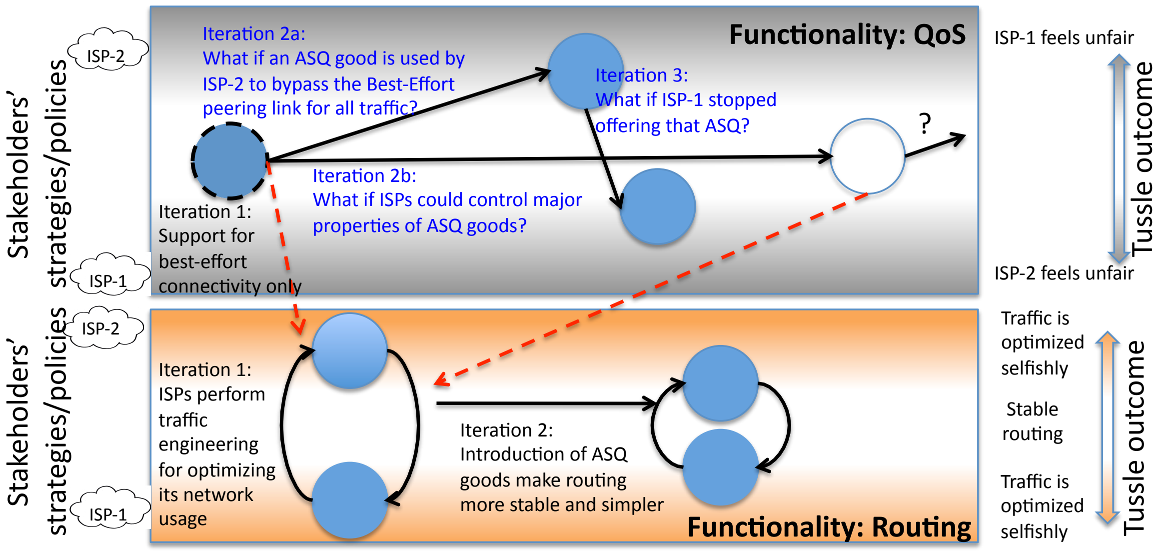

Figure 2: Tussle evolution during network service composition between competing ISPs

Now let's consider the impact of two ASQ goods on ISP-1 and ISP-2 regarding the QoS functionality, depending on what SLA properties can be configured. We do this by performing more iterations of the tussle analysis methodology.

In the first case an ASQ good is assumed to describe QoS-related properties only, like bandwidth, delay, jitter, etc. If such a *hypothetical* ASQ good has been setup between ISP-1 and ISP-2 then the former would increase the quality that its customers perceive when interacting with customers of the latter for services like video etc. Similarly, ISP-

2 would get premium quality connectivity for both the Content Provider, as well as the rest of ISP-1s end-customers, without increasing its cost. Step 3 would conclude that this new situation is advantageous for ISP-2 in competing with the ISP-1 for end customers. The reason is that customers of ISP-3 can access a popular destination (the Content Provider) with similar quality across both transit providers ISP-1 and ISP-2. As we described above this tussle outcome is not desired by ISP-1 and thus is not stable (shown as a blue circle).

One way for ISP-1 to deal with this tussle would be to stop offering that ASQ good and exchange their customers' traffic through a transit provider (we assume that the peering link no longer exists). In this outcome both providers would experience increased cost in relation to the initial state and would not be satisfied. Similarly, ISP-2 would find it beneficial to peer with the Content Provider for free. If ISP-1 had performed a Return-on-Investment (ROI) analysis before adopting the ETICS solution then, under certain assumptions related to its effect on demand for other market segments, this would lead to an unstable outcome.

In the second case a mechanism is introduced for determining the set of IP addresses serviced by an ASQ agreement. This would allow ISP-1, for example, to setup an ASQ agreement for the Content Provider's range of IP addresses by asking a fee and another one for the rest of its customers for free (Step 2). This is in line with the existing situation where peering links are not suitable for moving data sensitive to congestion effects. It is important to note that in either case, the introduction of prices helps the parties involved to find an equilibrium that is fair for both of them.

What is interesting is that the new outcome of the network service composition functionality has a positive impact on routing and traffic engineering functionality (shown with a dotted red arrow). The reason is that ISPs do not have to perform the complex traffic engineering anymore to improve the QoE (Quality of Experience) of their customers. This could have led to a stable outcome, but we expect that some ISPs with less spare capacity would still rely on traffic engineering for meeting their SLAs. Thus, routing instabilities may still exist but these should have smaller impact on other ISPs.

5 SLA monitoring and incentives for backup ASQ provisioning

SLA monitoring functionality is considered to be important even when trusted operators (as in the case interested in ETICS) must collaborate in order to provide QoS-assured services end-to-end. This is because such premium transport services are secured by SLA terms and can trigger payments to the customer in case some of the ISPs in the ASQ path fail to meet its requirements. On the other hand, the monitoring solution must be carefully designed in order to keep capital and operating expenses low.

In this case study we look at another interesting implication of monitoring, related to the reservation of backup capacity for those ASQ agreements, whose contract specifies the reliability parameters. Although not directly mentioned in the SLAs, ISPs are expected to keep backup capacity available in case the original path used by the ASQ agreement has a failure point. If a failure occurs in its network, the ISP will need to reroute the traffic of the ASQ agreement which might cause a QoS degradation and hence an SLA violation. Of course whether this happens depends on the amount of backup capacity available in the network of both ISPs.

Monitoring can help provide the right incentives for keeping backup capacity since it enables finding the ISPs who can be considered responsible for the QoS failure. In simple terms, if no adequate monitoring is in place to identify the ISP who caused the rerouting and the QoS violation, then the penalty of the violation will be assigned to the service originator (the first ISP that interfaces to the customer of the ASQ agreement), or can be divided equally among the ISPs. One can easily see that at the equilibrium, ISPs will pick a strategy to provide minimal backup. There is obvious free-riding since the effects of low backup provisioning are shared among a large number of ISPs.

We will apply the tussle analysis methodology to one of the candidate schemes examined by the ETICS project, the more likely to be deployed. In the distributed hierarchical scheme, each ISP collects raw data by specialized border routers, called probes. In order to keep the operational cost tractable, data sampling is performed. Furthermore, sampling suggests that monitoring data stored by each ISP along the path refer to the same packets; otherwise not all SLAs and their metrics (for example the e2e one-way delay for short time contracts) may be checked. The monitoring data are stored at dedicated databases, called proxies, operated by each ISP to overcome confidentiality issues. In the case of an SLA violation ticket, a collector queries all relevant proxies and compares retrieved data in order to check the validity of each SLA violation ticket. ISPs or trusted third parties can act as Brokers operating collectors.

The first step would reveal that the main set of stakeholder roles includes ETICS Communication Service Providers and more specifically Edge and Transit ISPs. Again, other involved roles such as Content Providers, Consumers of ICT services and Regulators have been excluded from this exemplary analysis.

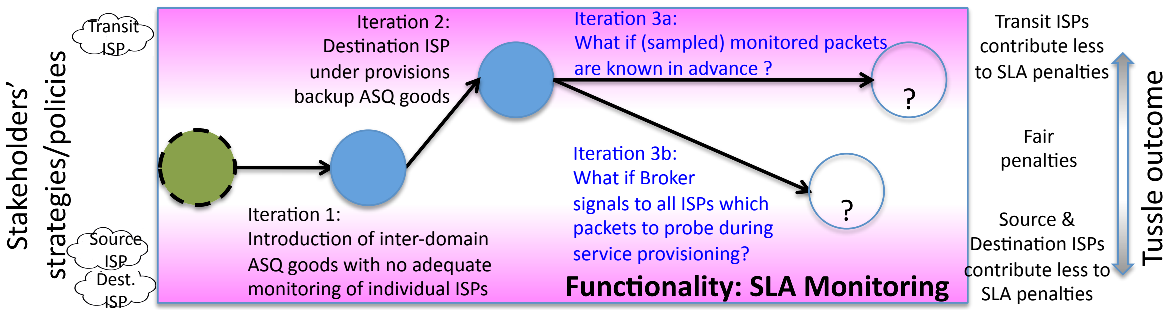

Figure 3: Tussle evolution during service delivery with assured quality between multiple ISPs

Figure 3 shows a possible evolution of the tussle described above between a Transit ISP (ISP-2 in our example) and Source, Destination ISPs (such as ISP-1, ISP-3 respectively). Investigating the Routing functionality, as long as Best Effort is the only traffic class available on the Internet, no SLA monitoring is needed and thus we assume that in the beginning we have a stable outcome (green circle).

The introduction of inter-domain ASQ goods creates the need for backup paths to be used in case of a sudden failure. All ISPs however have the incentive not to announce sensitive information such as network topology and dimensioning. Furthermore, they tend to keep backup capacity low to avoid unused and therefore unbilled capacity. This means that the new tussle outcome is not stable at the second phase, but no SLA violation has been reported and thus the outcome is still fair.

Let us examine the case where a trusted third party (a broker) implements the hierarchical SLA monitoring mechanism and ISPs agree to allow the collector to access data from their proxies. Such a technology, together with an incentive mechanism for calculating a fair allocation of the compensation to the ISPs, could lead to a stable tussle outcome. However, depending on the implementation of the SLA monitoring mechanism, identification may not always be feasible. We identify two cases for the SLA monitoring technology for illustrative purposes.

In the first case where sample packets were known in advance, Transit and Destination ISPs could forward probing packets preferentially. Thus the responsible ISPs may not always be identified. This means that the total payments made for customer SLA violation is analogous to the number of end-customers that ISPs have. Assuming that a Transit ISP has fewer customers than an Edge ISP, we can conclude that this tussle outcome is more beneficial to the former one.

The second case could be that Brokers signal to all ISPs along the path which packets to probe during service provisioning and they have to save a timestamp for those packets at the egress router. Following a secret algorithm for selecting those packets, a Broker could make harder the expedite forwarding of sample packets only, giving the right incentives for dimensioning of backup paths and lead to a stable outcome.

6 Conclusions

In this paper we applied a tussle analysis methodology for investigating potential solutions for two case studies, allowing us to assess what effect those potential solutions would have on major stakeholders and other, related functionalities. Our preliminary qualitative analysis provides evidence for the technical solutions that should be preferable in the case of ETICS network service composition and candidate mechanisms for the case of network service delivery. In the future, we plan to perform this methodology to a wider set of case studies and possibly use analytical techniques (such as game theory) to backup our claims.

7 Acknowledgments

These results were supported by the SESERV project (EC FP7 258138) and the ETICS project (EC FP7 248567).

References

[1] P. Vixie, “What DNS is not,” *Queue*, vol. 7, pp. 10:10–10:15, November 2009.

[2] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, “Tussle in cyberspace: defining tomorrow’s internet,” *IEEE/ACM Trans. Netw.*, vol. 13, pp. 462–475, 2005.

[3] K. Kilkki, “Sensible design principles for new networks and services.,” *First Monday*, vol. 10, no. 1, 2005.

[4] C. Kalogiros, C. Courcoubetis, G. D. Stamoulis, M. Boniface, E. T. Meyer, M. Waldburger, D. Field, and B. Stiller, “The future internet,” ch. An approach to investigating socio-economic tussles arising from building the Future Internet, pp. 145–159, Berlin, Heidelberg: Springer-Verlag, 2011.

[5] H. Ergas, “Internet peering: A case study of the ACCCs use of its powers under part XIB of the trade practices act, 1974,” in *Proc. 3th Internet Economics WS*, 2000.