

Inter-domain Coordination Models

Eleni Agiatzidou¹, Costas Courcoubetis¹, Olivier Dugeon², Finn-Tore Johansen³,
George D. Stamoulis¹

¹Athens University of Economics and Business,
{agiatzidou, courcou, gstamoul}@aueb.gr

²France Telecom-Orange Labs,
olivier.dugeon@orange.com

³Telenor Digital Services,
finn-tore.johansen@telenor.com

Abstract. In order for the Network Service Providers (NSPs) to provide end-to-end Quality of Service (QoS) at the inter-domain level different coordination models have been proposed by ETICS project. In this work we present and analyze the plausible alternatives of those models and we compare them with each other in terms of information asymmetry issues. We show that different information sets affect the total service provision and we present a basic model analysis on information issues by means of game-theoretic models.

Keywords: Charging and pricing QoS, game theory, end-to-end services

1 Introduction

QoS enabled services such as Gaming as a Service require end-to-end (inter-domain) connectivity that can be achieved through the existence of agreements among the involved Network Service Providers (NSPs). Interconnection agreements are constantly changing to fulfill the needs of the new technologies affecting the Internet architectures [1]. The FP7 ETICS project [2] studies those Internet architectures and the collaboration models that are capable of sharing information about business and technical parameters to provide QoS-enabled end-to-end connectivity services.

According to the ETICS solution, the NSPs share information concerning the traffic source, destination and the statistics of the QoS-enabled service, as well as the price of their part of the service. The price affects the final service that will be provided to the buyer. The revelation of the buyer's willingness to pay may lead to advantageous positions for some providers. On the other hand the total price that the chain of providers asks for an end-to-end service may lead to failure of its provision. Also in cases of many paths between a source-destination pair the pricing strategy of an NSP may affect the final chain of NSPs that will serve the customer. In this paper, we study a variety of inter-domain coordination models defined in the framework of ETICS, and we address the information and coordination issues that arise. The knowledge of how the information affects the formation of the end-to-end paths is necessary for the right definition of pricing and revenue sharing models.

In section 2 of the paper we describe the ETICS framework while in section 3 we present the different inter-domain coordination models. In section 4, we present a comparison of the models based on information issues. In section 5, we formulate a game-theoretic model pertaining to one of the above models in its simplest form in order to draw some preliminary conclusions that verify our initial analysis.

2 ETICS Framework

The ETICS project [2] has introduced the notion of the ETICS Community. Fig. 1 depicts the interaction between the ETICS community and a customer, who asks for provision of a QoS service from its collaborating NSP inside the ETICS community.




Fig. 1. ETICS Community

The nodes that belong to the ETICS Community collaborate in order to provide QoS-enabled services to members and non-members. Those nodes represent edge, transit or transport NSPs and form paths between an S (source) and a D (destination) NSP. In our considered case, the objective of NSP S (which we will call “buyer” to differentiate from the “customer” which does not belong to the community) is to buy connectivity to an IP range in domain NSP D allowing its packets to reach the specific destination under the conditions desired. Besides the destination range, a Point of Interconnection (PoI) and specific requirements, such as maximum delay, minimum bandwidth, jitter, period of validity etc. have to be defined. The PoI describes where and how the traffic is exchanged between two NSPs. These parameters along with a price specify completely a service that is provided to the buyer, or in other words an SLA offer, which has to be created from sub-SLAs offered by the various NSPs in the path. For the provision of a bundled offer, the NSPs participating in the community have to share information about what they are willing to provide. This is the Publishing Phase. When a buyer is willing to purchase a QoS-enabled service, the Service Composition Phase is triggered to perform the abutment of the different sub-SLA offers in order to provide a global proposal to the buyer. In the next section we study the different scenarios of those two phases and the effect that different information sets may have on the pricing strategies of the participants.

3 Inter-domain Coordination Models

As studied in [3], the role of information asymmetry is critical when considering interconnection agreements. The way information about NSP network capabilities is

propagated affects the topology that each NSP will be aware of after the Publishing Phase. The Service Composition Phase can be performed by using Centralized and Distributed models. The Centralized models require the existence of a central entity, called Facilitator, which gathers information from the NSPs and assists the composition of the offer under agreed selection criteria. In the Distributed models the information is propagated between neighbors only. The Service Composition can furthermore be done by using Push and Pull models, which concern the creation time of the offer. In Push models, the offers are available before a buyer's request while in Pull models the offers are created upon a buyer's request. The combination of those models and the creation of two hybrid ones results in six different scenarios. For all these scenarios we investigate the minimum information set that each participant must have available in order for the model to serve its purpose. In all subsequent scenarios, we suppose that an ETICS customer issues requests to the NSP S about connectivity from NSP S to D . Even if we show only one intermediate NSP A for the clarity of the figures, all scenarios could be extended to many intermediate NSPs.

Fully Centralized Push Model. During the Publishing Phase all NSPs in the community inform the Facilitator about their sub-SLA offers. For the rest of our analysis we assume that these sub-SLA offers contain: the logical point(s) of interconnection (PoI), the destination(s) network(s) prefix, the QoS characteristics, the expiration time of the offer and a price. The Service Composition Phase begins with the buyer's request. The buyer (S) communicates with the Facilitator and reveals only the PoI (where its traffic is taken care of the ETICS Community), the destination and the QoS characteristics that it wishes to have for the specific end-to-end connectivity (SLA Request). The Facilitator, based on the knowledge that it already has due to the Publishing Phase and the SLA request, combines sub-SLA offers of those NSPs that form a chain from the source NSP to the destination NSP(s). If there are multiple possible offers, the model is flexible regarding the selection criteria of the Facilitator. Indeed, it can choose the best offer from a social welfare or the buyer's point of view or even let the buyer decide. The buyer may accept or not the price of the bundled SLA offer(s) presented by the Facilitator (Fig. 2).




Fig. 2. Fully Centralized Push




Fig. 3. Fully Centralized Pull

Fully Centralized Pull Model. The Publishing Phase of this model is slightly different from the previous one. The Facilitator is aware of the Network Capabilities (PoIs,

and QoS characteristics but not the final price offer) of all NSPs of the community. Based on this information, the Facilitator computes one or multiple NSP chains that could potentially handle the buyer's request for PoI, destination and QoS characteristics. When it receives a request from the buyer it sends specific sub-requests to the NSPs in those chains. After receiving the sub-SLA offers, it combines them. The best offer is chosen under one of the criteria that previously mentioned (Fig. 3).

Distributed Push Model. In this model the information that is available to each NSP depends entirely on the Publishing Phase. Each NSP in this phase publishes its sub-SLA offers to all of its neighbors, to a subset of them or to the whole community. In the first two cases each NSP may combine its own offers with all the offers that it is already informed about and propagate the combined offers to other NSPs in order for the information to be diffused to other participants. Thus, each NSP may be notified about already bundled sub-SLA offers that are available in the community in a cascading way. The buyer will negotiate with its neighbor in order to buy a bundled offer (Fig. 4).




Fig. 4. Distributed Push Model




Fig. 5. Distributed Pull Model

Distributed Pull Model. During the Publishing Phase the NSPs exchange information about their Network Capabilities. As in the previous model, the NSPs may decide what kind of information they publish and propagate to each of their neighbors. The Composition phase begins with the buyer's request for connectivity from its neighbor(s) in order to reach a destination (Fig. 5). The neighbor accepts (or rejects) this SLA Request by providing a price offer. In case of accepting, it extracts the part of the requesting SLA that corresponds to its Network Capabilities and adds a price. It subsequently propagates the price and the remaining Network Capabilities requirements to its neighbor(s). After the destination NSP adds a price to the SLA request, the complete offered price is propagated back to the buyer in the reverse cascading way, which then accepts or rejects the final offer. Clearly, the position of an NSP in the chain can affect its strategic power. Such issues can be analyzed by means of game-theoretic models; as discussed in Section 5.

Per-NSP Centralized Push Model. This model is a hybrid of the Centralized and the Distributed Push models. The NSPs create sub-SLA offers (including prices) that are published to catalogues that are available to the community. In the Composition Phase the buyer combines offers and creates a bundled one that fulfills its needs. As opposed

to the Distributed model, here the buyer buys the SLA offers from each NSP even if it is not directly connected to it (Fig. 6).




Fig. 6. Per-NSP Centralized Pull Model




Fig. 7. Per-NSP Centralized Push Model

Per-NSP Centralized Pull Model. The buyer is informed about the Network Capabilities that each NSP is willing to provide during the Publishing Phase. During the Composition Phase it asks each NSP in the chain for an SLA offer (which is a part of the whole ASQ good that it wants). In this model (Fig. 7) the buyer communicates separately with each NSP in the chain. The apportioning of price and/or QoS is done according to the buyer's own strategy who leads the composition.

4 Information Issues on the Inter-domain Coordination Models

The total price of the service, the revenue sharing among the NSPs (affected by means of the prices assigned to their sub-SLA offers) and the selection of the NSPs that will participate in providing a service differ across the models. These differences have their origin in the information set that is available to each NSP and is formed during the Publishing Phase, as well as in the fact that a different entity makes the decision of the optimal offer in each model.

In case of a Centralized model the Facilitator has all the available information during the Publishing Phase. In the Distributed models the information must be propagated through the NSPs. The simplest way is to propagate information to the direct neighbors who undertake to bundle it with their own and redistribute it. However, this approach allows for strategic behaviors. For example, an NSP may decide not to propagate the information of all of its neighbors. This may result in different visible options for topologies due to different sets of information available. To overcome this possibility, rules may be enforced in the community in order for the information to be propagated equally. For example a rule could be enforced that the information must be flooded to all community members. Per-NSP Centralized models suffer from the same problem as the Distributed ones. Even if NSPs publish their information in catalogues which can be accessed by every member and thus allowing equal information to all, a minimal set of information about the topology and how the NSPs can be reached has to be available to all participants in the community. In order for all models to be equivalent in terms of the topology that each NSP is aware of, the respective information in the Distributed models must be flooded to all NSPs in the community.

The entity that makes the choice about the optimal offer that will be purchased may be the Facilitator or the buyer. One of the approaches applying to the Centralized

models is for the Facilitator to propagate the available choices to the buyer and let it decide; alternatively, the Facilitator may decide according to its own criteria on the final choice, or on a subset of them which are subsequently passed to the buyer for the final selection. In such a case the Facilitator can choose to maximize the expected social welfare, thus improving the overall well-being of the community of customers. Related criteria may be the maximization of the number of buyers that can be satisfied, or of the utilization of the network. Alternatively, the Facilitator can choose the optimal offer by applying potential buyers' criteria, typically the lowest price.

In the Distributed (and Per-NSP) models the buyer accepts or not the final offer(s) that are reaching it and thus it has the final choice. However, each NSP may influence the price of the bundled offer. Putting a high price in its part of the service an NSP increases both its profits if this offer is indeed purchased and the risk of rejection of the offer. Authors in [4] present similar problems and show that non-cooperative pricing strategies between providers may lead to unfair distribution of profit. Moreover the position of an NSP in the path to the destination may influence the price that it claims for its part of the service. Thus, certain NSPs may have an advantage in this model due to their position; e.g. the first NSP in the chain, or possibly bottleneck or central NSPs. As this constitutes a problem for all models where the NSPs publish their prices sequentially, and thus most of our models depending on their implementation, we illustrate in the following section the analysis of a simple Distributed Pull model that highlights the effect of the advantageous position of an NSP.

5 Analysis of a Distributed Pull Model

The effect of having an advantageous NSP position in the chain on pricing strategy and on the price of the whole service is presented through a theoretical analysis of a simple Distributed Pull model. We compare the model's efficiency in terms of the probability of the service being offered to the basic collaborative model where the NSPs share the price of the buyer so that they have equal profits.

We assume that we have only three NSPs in the chain; thus, the buyer S requests from A an SLA to connect to D (see Fig. 5). We assume that the Publishing Phase has already been completed and all NSPs have the same set of information about the network capabilities of the other NSPs. The Composition Phase is initiated by S who sends its request for a particular SLA. After deciding on its price P_A , A propagates the SLA and P_A to D . (Note that the results are still applicable if we reverse the roles of A and D , with D choosing first.) D decides its own price P_D and they propagate $P_A + P_D$ to the buyer which takes the offer or not. For A and D to agree to make the transaction their prices have to exceed their costs (C_A, C_D). Also $P_A + P_D \leq P_S$ and $C_A + C_D \leq P_S$, where P_S is the price that the buyer is willing to pay, should apply in order for the service to be ultimately provided. Also, both the costs are here assumed to be known to both the providers who only know the distribution of P_S . The problem that the NSPs A and D have to solve is:

$$\max E[(P_*(C_*) - C_*) * 1(\text{Service is provided})]$$

The interaction of A and D corresponds to a Stackelberg game. The optimization problem for D is:

$$\max[(P_D(C_D) - C_D) * \Pr[P_A + P_D(C_D; P_A) \leq P_S]]$$

where P_D is a function of C_D and P_A , s.t. $P_D \geq C_D$. Knowing that D will act this way, the optimization problem for A becomes:

$$\max[(P_A(C_A; C_D) - C_A) * \Pr[P_A(C_A; C_D) + P_D(C_D; P_A) \leq P_S]]$$

where P_A is a function of C_A , and C_D s.t. $P_A \geq C_A$.

In order to proceed with the analysis of a simple yet illustrative case, we assume that P_S follows a uniform distribution and thus $P_S \sim U[0, P_{Smax}]$. Since A chooses its price first it will choose P_A such that $0 \leq P_A + C_D \leq P_{Smax}$. It can avoid the case of $P_A + C_D \geq P_{Smax}$ (where the service is not provided) since it is assumed to know C_D too. (Note that A can be assumed to have some knowledge on the cost of D if e.g. they have similar infrastructure. Of course, the assumption that this knowledge is accurate is only adopted for preserve simplicity of the model.) Once P_A is announced to D , then the only feasible and at the same time meaningful choice for P_D is to satisfy: $P_A + P_D \leq P_{Smax}$. Solving the maximization problem for D we obtain the optimal choice for him (P_D^*). Subsequently, we solve the maximization problem of A . The value of P_D^* is not known to A . It comes as a result of the choice of P_A by A . However, A can make use of the expression of P_D^* and then calculate P_A on this basis. A will choose P_A such that $P_D^* + P_A \leq P_{Smax}$. Through those optimal choices we can calculate the actual profits that A and D in the cases that the service is indeed provided. In the table below, we show some numerical values of those profits and we compare them to the collaborative scenario. As presented and also proved in closed form the profits of A in the Distributed Pull model are always higher than in the collaborative one and also double those of D . Also the third row shows a case of a failure of service provision under the Distributed model in contrast to the Collaborative one. The last row shows a failure of both models due to a low price of the buyer.

Table 1. Comparison of Profits of A and D

(C_A, C_D, P_S)	Collaborative	Distributed Pull	Distributed Pull
	Actual Profits	Actual Profits of A	Actual Profits of D
(0.043, 0.169, 0.649)	0.218	0.394	0.197
(0.113, 0.030, 0.939)	0.387	0.417	0.208
(0.075, 0.054, 0.531)	0.200	-	-
(0.547, 0.138, 0.149)	-	-	-

Also we can evaluate this model by showing the loss in efficiency. This loss can be quantified by comparing: a) the probability $\Pr[P_D^* + P_A \leq P_S]$ (see eq. 1) that the service is achieved under this model, where A and D follow their own optimal strategies (P_A^* , P_D^*); and b) the probability $\Pr[C_D + C_A \leq P_S]$ (see eq. 2) that a service is achieved collaboratively. Thus, when the NSPs act selfishly (Distributed Pull model)

there is a huge loss in efficiency, resulting in a reduction of the probability that the service is offered by a factor of 4.

$$\Pr[P_D^* + P_A^* \leq P_S] = \Pr\left[\frac{3P_{Smax} + C_A + C_D}{4} \leq P_S\right] = \frac{P_{Smax} - C_A - C_D}{4P_{Smax}} \quad (1)$$

$$\Pr[C_D + C_A \leq P_S] = 1 - \Pr[P_S \leq C_D + C_A] = \frac{P_{Smax} - C_A - C_D}{P_{Smax}} \quad (2)$$

6 Conclusions

In this paper we have studied the different coordination models proposed by the ETICS project and the way that different information sets affect the service provision under those models. Our research (only partly reported in this paper) has shown that for the Distributed models, either in Push or Pull ones, NSPs may gain an advantageous position against others by excluding a neighbor from having information of other NSPs and thus resulting in a topology different than that in a Centralized model. On the other hand the adoption of a Centralized model implies the concentration of all the information from the NSPs in one single entity which has to be unbiased and meet an optimization goal agreed by all community members. Therefore, future work on the ETICS architecture should consider and evaluate the alternative of using appropriate rules to promote collaboration and deter aggressive selfish pricing strategies that would mitigate the system's performance.

Acknowledgements. The research work of E. Agiatzidou has been co-financed by the European Union (European Social Fund- ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) – Research Funding Program – Heracleitus II. Investing in knowledge society through the European Social Fund. Also, this work has been performed in the framework of the EU ICT Project ETICS. The authors would like to thank all ETICS partners for their useful discussions on the subject of the paper and specially M. Dramitinos, N. Le Sauze, H. Pouyllau, P. Zwickl and A. Cimmino.

References

1. Faratin, P., Clark, D., Gilmore, P., Bauer, S., Berger, A. and Lehr, W.: "Complexity of Internet Interconnections: Technology, Incentives and Implications for Policy". 35th Annual Telecommunications Policy Research Conference, George Mason University, September 2007.
2. 2010 ETICS project consortium. (2010). ETICS. Available: <https://www.ict-etics.eu/>. Last accessed: 27th February 2012
3. Constantiou, I.D. and Courcoubetis, C.: "Information Asymmetry Models in the Internet Connectivity Market". 4th Internet Economics Workshop, Berlin, May 2001.
4. He, L. and Walrand, J.: "Pricing and revenue sharing strategies for Internet service providers", IEEE JSAC, 24:5, May 2006, pp. 942-951.