Sharing is Power: Incentives for Information Exchange in Multi-Operator Service Delivery

Poul E. Heegaard*, Gergely Biczók†, and Laszlo Toka‡
*NTNU, Dept. of Telematics, Trondheim, Norway, poul.heegaard@item.ntnu.no
†BME, MTA-BME Future Internet RG, Budapest, Hungary, biczok@tmit.bme.hu
‡BME, MTA-BME Information Systems RG, Budapest, Hungary, toka@tmit.bme.hu

***Abstract*—A majority of 5G verticals have the potential to generate large revenues, but are expected to have strict Quality of Service (QoS) guarantees, and are projected to be delivered as a service chain of multiple, independent operators. Such multi-operator service delivery requires a set of interdependent Service Level Agreements (SLAs) between operators. The amount and aggregation-level of information shared between stakeholders inside such SLAs will determine how efficient the coordinated traffic engineering between the operators will be. Sharing more details on one's network is uncommon in today's interactions due to the fear of losing competitive advantage and regulations with regard to national security. In this paper, we analyze the economic incentives for information exchange in the context of multi-operator service delivery. We show that the current practice of exchanging only highly aggregated information can lead to both significant under- and overestimation of the risk of not meeting user-facing Quality of Service guarantees. We also show that economic incentives for mutually sharing an optimal amount of information do exist, and optimal information exchange between operators is viable in the long run. Moreover, through a simple numerical example, we demonstrate how the mutually shared information and the resulting risk estimation affect the revenues of the operators from the end-user market. We believe this work opens up a new line of research connecting the economics of multi-operator service delivery and network performability.**

I. Introduction

The ICT service-provisioning infrastructure and related digital ecosystems are increasingly complex and evolving systems where services delivered involve multiple business entities (stakeholders) sharing the responsibility of providing robust, dependable and predictably high performance services. A textbook example for such a complex ecosystem is that of anticipated 5G verticals, which hold the business potential of being true value-added services [1]. Such verticals inherently require strict end-to-end Quality of Service (QoS) guarantees and the collaboration of mobile and fixed network operators, cloud infrastructure owners and Over-The-Top (OTT) providers, comprising a multi-actor value chain. This inevitably calls for multi-operator business, service and resource coordination, which is in the focus of the in-progress H2020 5G-PPP 5G Exchange project (5GEx) [2]. While emerging concepts such as the softwarization of the network control plane (Software Defined Networking, SDN) and the virtualization of resources and network functions (Network Function Virtualization, NFV) can serve as powerful technological enablers, the business framework including inter-stakeholder contracts, resource trading and coordination models require its own (but integrated) mechanisms.

Bilateral contracts between operators and between end-user and operator are referred to as Service Level Agreements (SLAs). SLAs give guarantees on the non-functional properties of the service provided including performance and dependability. The chain of SLAs (along the service chain) is the means of ensuring end-to-end QoS for the end-user, but it is insufficient in its current form to fulfill its objective. One problem is that the SLAs between the operators do not include sufficient information about QoS properties for appropriate coordinated resource management and traffic engineering, and the information that is included is highly aggregated. The reason behind this is two-fold. First, on the technical side, the current *de facto* inter-operator information exchange protocol, BGP (Border Gateway Protocol), is inflexible and limited: information is only the destination IP prefix, scope is only direct neighbors, and policies can only be expressed indirectly. Note that SDN can provide a richer set of capabilities for exchanging more detailed information better suited to traffic engineering purposes [3]. Second, on the business and regulation side, the fear of losing competitive advantage (willingness) and national security regulations with regard to protection of information on critical infrastructure (feasibility) leave operators with exchanging only a limited set of highly aggregated information. Thus, risk estimation along the service chain, and particularly at the user-facing provider is bound to be imprecise, potentially affecting revenues directly.

In this paper, we analyze the economic incentives for information exchange in the context of multi-operator service delivery. Our main contribution is twofold. First, we develop a simple model for estimating the risk of not meeting QoS guarantees under different information aggregation regimes. We show through a simple two-operator example that exchanging coarse grain information can lead to both significant under- and overestimation of risk. Second, we develop a game-theoretical model of information exchange between operators. We show that sharing the optimal amount of information is a sustainable equilibrium if there is a mutual, long-term cooperation among the operators. Furthermore, we integrate the risk and game models in a simple numerical example, and show how the information exchange affects the revenues of the operators from the end-user market. We conclude that given a platform enabling mutual trust and cooperation (such as the 5G Exchange proposed in [4]), operators benefit from exchanging sufficiently (but not excessively) detailed information and the resulting, more precise risk estimation.

The rest of the paper is organized as follows. Section II introduces the structure of SLAs and the price of uncertainty in risk estimation. Section III describes a simple risk model, and presents, via a two-operator example where the ominous QoS guarantee involves maximum end-to-end delay, a risk calculation for not meeting the specific guarantee for different levels of transparency. Section IV develops a game model inspired by the well-known prisoner's dilemma, and presents its equilibrium analysis. Section V integrates the two models and presents a brief numerical study for the above-mentioned two-operator example concerning revenues from the end-user market. Finally, Section VI provides concluding remarks.

II. Service Level Agreements for Multi-Operator Service Delivery

In this section, we first take a look at SLAs, and then make the connection between the information shared in SLAs and the corresponding risk estimation at the user-facing operator.

A. Service Level Agreements

When an operator sells a service to end customers (partly) building on the infrastructure of another operator, the business relation between the operators (and between end-user and user-facing operator) is mutually agreed upon and is fixed in a Service Level Agreement (SLA) (see Figure 1). The SLA consists of: Service Level Objects (SLOs) defining technical minimums to be provided; the agreed cost of the service; the compensation, i.e., the fee to pay to the buyer if the SLOs are not met; and the contract period.

SLOs translate into traditional QoS definitions that describe performance and dependability of the overall system in which the service for the end customer is created. The granularity of information shared within SLOs, however, determines how well the mapping of SLOs into QoS towards the end customers can be performed. In general, SLOs contain numerical performance attributes (e.g., 10 ms of delay) *and* their probabilistic occurrence (e.g., met in 99% of packets). QoS guarantees towards end customers are composed of these SLOs in order to present a compact, high level, system-wide description, e.g., maximum end-to-end delay, that hides the potentially complex cascade of multi-operator infrastructure.

Network-related SLO *performance* metrics typically describe throughput, delay and packet loss with their mean, maximum, minimum or empirical distributions yielding the probability of measured values to be above maximum or below minimum. The *dependability* of the infrastructure can be given in terms of availability, max down time, number of failures, number of long outages, etc. For both types of attributes long-time measurements provide the input, and details of these measurements (methodology, duration, location, etc.) can be part of the SLOs. The risk and information sharing models in the paper are metric-agnostic.

The monetary sections of an SLA include the cost of service and compensation. In general a service with better/stricter QoS guarantees can be sold at a premium, given that there is a willingness-to-pay among the prospective buyers. Compensation is usually given as a reduction in the consequent billing period; typical compensation functions are step-wise

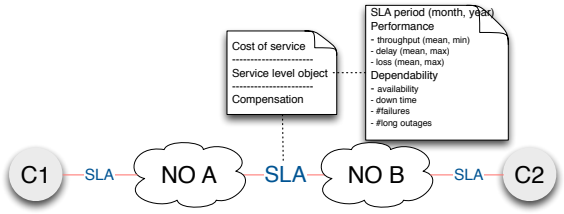

Fig. 1. Bilateral QoS information exchange through SLAs

and somewhat rudimentary [5]. We assume that receiving the service satisfies the customer more than getting even the highest possible compensation (e.g., the maximum compensation in the Google Compute Engine SLA is 50% of the monthly fee [6]).

If the business relation of two operators is symmetric, i.e., both operators sell services to each other, their respective SLOs provide mutual insight into each other's infrastructure, operation and management. In the expected 5G service ecosystem this will be the rule rather than the exception: operators have to collaborate in service delivery to satisfy end-to-end QoS, extend their geographic footprint or simply lease idle resources from each other. A simplistic setup is drawn in Figure 1, where both operator *NO A* and *NO B* have their customers paying for services that are provided with the support of the other operator, within the boundaries of their framework contract SLA. This is the setup we will assume throughout the paper.

B. The price of uncertainty

Operators design their services and sell them to end customers factoring in all the elements of cost of service. Apart from the cost of provisioning the service based on own investments and operations, if the service is partially built on another operator's infrastructure, then the price of using that is also added. Should the provider fail to meet the QoS guarantees advertised in the published service plans, (i) it has to pay compensation to its users, (ii) its reputation is damaged potentially materializing in users switching to other, substitute providers.

Every operator has to be able to estimate the probability of not meeting a certain QoS guarantee. Owing to the interdependence of operators in a multi-operator setting, estimating this risk can be challenging. SLAs (and the SLOs inside) are the only vessel for sharing information. Historically, SLOs lack details in order to avoid sharing in-house information with the competitor, to hide an operator's internal operations, network design and policies. However, as multi-operator services may become the norm in the near future, operators should find a balance between competition and cooperation. Estimating service-related risks based on SLOs without sufficient details, leaving potential physical and logical interdependencies (which might even vary over time) uncovered, leads to high uncertainty. If the operator faces high uncertainty, two sub-optimal outcomes are possible: (i) it underestimates the risk, and advertises stricter than realistic QoS guarantees, leading to higher compensation payment and reputation damage or (ii) it overestimates the risk and advertises looser than feasible QoS guarantees resulting in a lower end user market price. In

both scenarios the operator ends up with less income than the optimal.

C. Prior work

Relevant literature can be divided to four categories. First, the role and impact of SLAs in federated environments are studied with relation to early IT outsourcing [7], open federated cloud computing [8] and QoS in federated cloud-based software defined networks [9]. Second, economics of SLAs are studied with regard to SLA negotiation for service composition in [10] and maximizing profits stemming from SLAs in [11]. Third, in the economics literature, specifically in the field of supply chains (which show structural similarities to service chains in multi-operator service delivery), contract design and information sharing is studied when there are competing supply chains in [12]; moreover, the coordination of information sharing is analyzed in [13]. Finally, risk-aware networking is scrutinized in great detail in [14]. Due to space constraints we do not attempt to present a complete literature review. For further related work we refer the reader to the above-mentioned papers and their references. Each of these papers make observations or propose mechanisms related to parts of our work; however, none of them combine SLA design, risk estimation and incentives for information exchange in a multi-operator environment.

III. Risk of breaking QoS guarantees

In this section we discuss information aggregation levels, define our use of the term *risk*, specify an estimator of the risk, and demonstrate the effect of exchanging information with different aggregation levels on risk estimation.

A. Risk estimator

We define risk simply as the *probability that the service in not delivered according to the guaranteed specified QoS attribute value*. The definition of risk is similar to the definition of *unavailability* in [15], but we emphasize the guarantee aspect, and relate it to both performance and dependability. In [14] risk assessment and modeling, risk response, and risk monitoring in communication networks are discussed in more details.

Let X be a stochastic variable that represents a QoS attribute value, then the *risk* is (θ is the threshold, max or min, of performance attribute X):

$$\mathcal{R}(X_i, \theta_0, \theta_1) = 1 - \int_{\theta_0}^{\theta_1} f_i(x)dx = 1 - F_i(\theta_1) + F_i(\theta_0) \quad (1)$$

where f_i is the probability density function of X_i and $F_i(x) = P(X_i \leq x)$. For example, if the guarantee is a maximum threshold, $\theta_{\max}$, then the corresponding risk $\mathcal{R}(X_i, 0, \theta_{\max}) = 1 - F_i(\theta_{\max})$. The definition in (1) also includes risk evaluation of guarantees in an interval $(\theta_{\min}, \theta_{\max})$ with $\mathcal{R}(X_i, \theta_{\min}, \theta_{\max}) = 1 - F_i(\theta_{\max}) + F_i(\theta_{\min})$

When a service is delivered over multiple domains, the guarantee to the end users is a functional combination of the X_is from each of the n domains i ($i = 1, \cdots, n$). Consider the following examples of guaranteed threshold over n domains

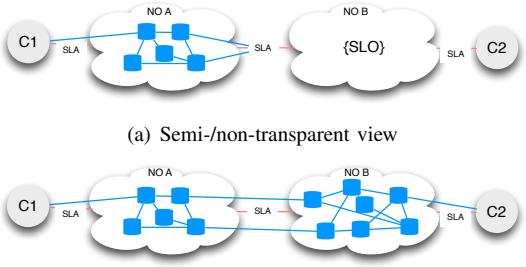

(a) Semi-/non-transparent view

(b) (Fully) transparent view

Fig. 2. Illustration of information exchange options where (a) network operator A has limited access to information about network domain B, and (b) network operator B exposes all available information

(where f_i of all n domains are known, and θ_1 is the maximum acceptable value of Y):

(i) *sum* (e.g. end-to-end delay): $Y = \sum_{i=1}^{n} X_i$

$$\mathcal{R}(Y, 0, \theta_1) = \int_{\theta_0}^{\theta_1} f_1 \oplus \cdots \oplus f_n(x)dx$$

(ii) *max* (e.g. max down time): $Y = \max_{\forall i} X_i$

$$\mathcal{R}(Y, 0, \theta_1) = 1 - \prod_{i=1}^{n} F_i(\theta_1)$$

Assume that only n_1 out of n domains will share information about their probability density function f_i. Without loss of generality, we can enumerate these domains $i = 1, \cdots, n_1$, and enumerate the remaining domains where only the mean values $\bar{X}_i$ are shared by $i = n_1+1, \cdots, n$. The risk estimation is changed (exemplified by the maximum value of a sum of stochastic variables):

$$\mathcal{R}(Y, 0, \theta'_{\max}) = \int_{\theta'_{\max}}^{\infty} f_1 \oplus \cdots \oplus f_{n_1}(x)dx \quad (2)$$

where $\theta'_{\max} = \theta_{\max} - \sum_{i=n_1+1}^{n} \bar{X}_i$.

B. Information exchange

In order to assess the risk across several domains it is important that sufficient information is exchanged between domain operators. In Figure 2 a small example is illustrated with two network domains A and B. The two domains are managed by one operator each, and there exists a mutual SLA between the operators. The two network operators provide services that depend on resources from both network domains. They both need to estimate the risk of not delivering the guaranteed QoS.

Example of information exchange options can be the following:

(i) *Non-transparent (high aggregation, Figure 2(a))* : first and second order statistics, such as median, mean, max, min, variance, standard deviations, quantiles, etc.

(ii) *Semi-transparent (medium aggregation, Figure 2(a))*: higher order statistics, such as correlation, probability distributions (of end-to-end QoS attributes), etc.

(iii) *Transparent (low aggregation, Figure 2(b))*: information about resources (network elements capacity and load) and structure (topology),

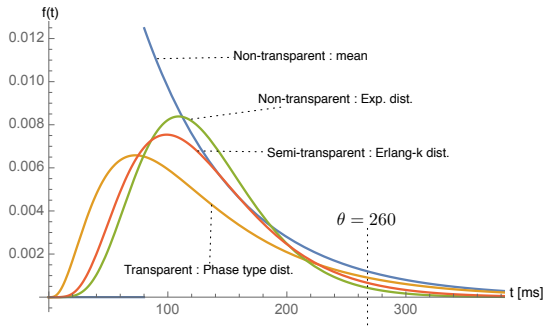

Fig. 3. End-to-end delay distributions over network domains A and B with different information available from domain B

(iv) *Fully transparent (no aggregation, Figure 2(b))*: as discussed above, this will be never exposed. Included only for reference. This will include everything that is necessary for a precise estimation of the performance metrics, i.e. all statistics (also distributions) per network element, topology, network element capacities, path routes (number of hops, or specific hops), traffic load on network elements, network element failure rates, measurement logs, etc.

C. Risk of delay not according to SLA

Here we consider an example with single path over two domains where the QoS guarantees are given as a maximum end-to-end delay. The effect of sharing information on different levels of details is demonstrated though a numerical example.

Information exchange aggregation between domains. The risk assessment of the end-to-end delay over two network domains will change depending on how much information is exposed, from one network domain operator to the other. Let $f_X(t)$ be the delay distribution of network X in Figure 2, $X = A, B$. We want to estimate the end-to-end delay distribution over network A and B; $f(t) = f_A \oplus f_B(t)$. In this example the following information exchange options are considered (options related to operator B in Figure 2):

(i) *Non-transparent* : mean value, X_B
(ii) *Non-transparent* : number of hops is not known, then assume one hop and therefore $X_B \sim \text{Exp}(1/X_B)$.
(iii) *Semi-transparent* : number of hops is k, but only the overall load is known, assume $X_B \sim \text{Erl}{-}k(1/X_B)$
(iv) *Transparent* : both k, the individual traffic intensities Γ_i, and service rates μ_i per hop i are know, then distribution over the domain is formed by convoluting the distributions over each hop, $f(x) = f_i \oplus \cdots \oplus f_k(x)$, and $f_i \sim \text{Exp}(\mu_i - \Gamma_i)$, phase type distribution, $X_B \sim \text{PH}$

Numerical values. Figure 3 shows end-to-end delay distributions over network domains A and B with different information available from domain B. Since the examples are just for the illustration of the aggregation's effect, for the sake of simplicity we assume that the delay of each hop in a path is exponentially distributed. In Table I you find the risk estimates for maximum thresholds $\theta_{\max} = 208, 260, 312$, with $E[X_A] = 50$ and $E[X_B] = 80$. The transparent case will give the most precise risk estimate out of the four options considered since this option contains more detailed information. The main observation from the table is that for all three threshold values considered we get both over- and underestimations of the risk. In the non-transparent case where we do not know the number of hops, we overestimate the risk. In the semi-transparent case where we know the correct number of hops but do not know the individual routing and traffic loads, the risk is underestimated.

TABLE I
END-TO-END DELAY RISK ESTIMATION FOR DIFFERENT INFORMATION AGGREGATION FOR INFORMATION EXCHANGE BETWEEN TWO DOMAINS, WITH $E[X_A] = 50$ AND $E[X_B] = 80$

aggregation		$\theta = 208$	$\theta = 260$	$\theta = 312$
non-trans: $\bar{X}_B$		0.201897	0.105399	0.055023
non-trans: X_B Exp	$\sim$	0.152858	0.080047	0.041820
semi-trans: X_B Erl$-k$	$\sim$	0.078358	0.018001	0.003534
transparent: X_B PH	$\sim$	0.106239	0.035683	0.011175

In Figure 4 the risk estimates of the four different information exchange options are plotted as a function of the threshold value θ. The plot also zooms in on the range $208 - 312$ which was used in Table I. The plots show that the relative ranking of the risk for the four options changes as the threshold changes. This means that if the operator does not have sufficient information, the risk can be both under- and over-estimated (with the same information exchange option) depending on the threshold value considered.

IV. THE INTER-OPERATOR INFORMATION SHARING GAME

One the one hand, risk estimation depends on the amount and nature of information exchanged between operators. On the other hand, the level of information sharing is a strategic choice for participating actors. We model their interaction as a two-player information sharing game; we use a simple, one-shot matrix game for motivation, and extend it in two steps to account for more realistic conditions. Basic game-theoretical concepts used in this section can be found in [16].

Binary sharing game G_0. Let us assume that two network operators rely on each other to a certain extent in providing a service to their customers (just like the simple setup of

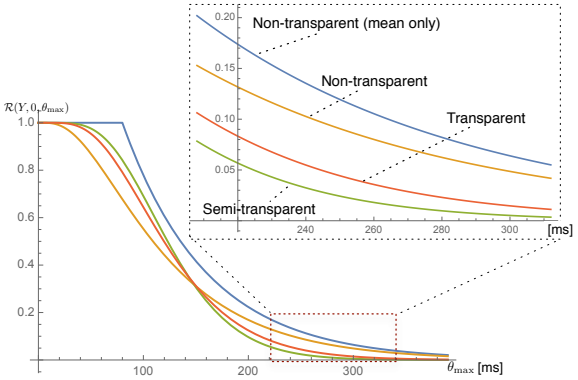

Fig. 4. Risk for different levels of transparency

TABLE II
PAYOFF MATRIX FOR G_0

		NO B: share	*NO B*: don't share
NO A	share	$(u-c, u-c)$	$(-c, u)$
	don't share	$(u, -c)$	$(0, 0)$

Figure 1). Let us also assume that the extent is the same: they equally rely on their mutual SLAs. Furthermore, let us also assume that they are equally sensitive in sharing information about their infrastructure with other operators, i.e., both see the same amount of threat (cost) in disclosing any given amount of data of aforementioned network topology, network element performance, etc. With these assumptions we arrive to a symmetric game model with opposing interests; for demonstration purposes we first assume each operator can decide between two strategies: *share* or *do not share* information.

Before drawing the game G_0 in its matrix-form, we have to establish the payoffs for each player in each outcome, i.e., combination of strategies played. Let us denote the cost of sharing information with c and the utility of getting information from the other player as u for both players. We assume that both not getting and not sharing information leads to 0 payoff. Then, the symmetric payoff matrix of G_0 is given in Table II. If $u - c > 0$, then G_0 is a special case of the much-studied prisoner's dilemma called the "donation game" [16], where cooperation equals to offering the other player a benefit at one's own cost. Such a game can be solved via the *dominant strategy* concept: here, each player prefers not sharing to sharing, irrespective of what the other player plays. Consequently, *mutually not sharing is the only strong Nash Equilibrium (NE, a strategy profile from which neither player wants to deviate unilaterally)* of the game G_0.

Continuous sharing game G_1. In a more practical scenario, operators may define different levels of aggregation for the information to be shared, such as the ones presented in Section III. In order to allow for multiple aggregation levels, we define a two-player game G_1. For the sake of brevity, we assume that the strategy space is continuous. Therefore, $s_i = I$, while $I \in [0, 1]$, where $I = 0$ means not sharing at all, and $I = 1$ means sharing all available information, relating to the fully transparent view in Section III. We define the payoff function of an operator (similarly to G_0) as the utility u of getting information on the other provider's network minus the cost c of giving information to the other provider on its own network, i.e., $\pi = u(I) - c(I)$.

Here, we make two assumptions; first, we assume that $u(I)$ is monotone increasing and concave, i.e., marginal utility is decreasing as the amount of information received from the other operator is increasing. This corresponds to an adapted form of the law of diminishing returns: there are other factors influencing the income (such as customer population, geographic limitations, willingness-to-pay, etc.), so improving on a single factor (information on the other operator's network) has diminishing benefits [17]. Second, we assume that the cost function $c(I)$ is monotone increasing and convex. Note

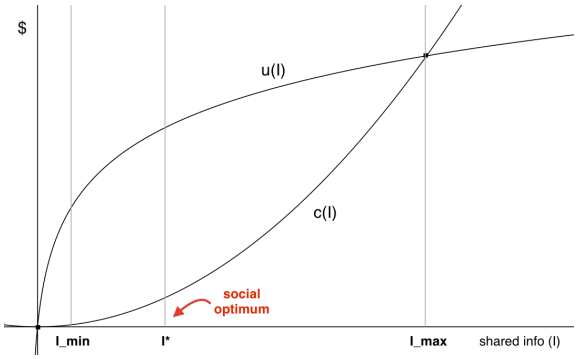

Fig. 5. Utility, cost and social optimum in G_1

that $c(I)$ only corresponds to the cost directly related to sharing information. Even in a co-opetitive environment such as the foreseen 5G ecosystem, operators are likely to keep certain internal information to themselves either for maintaining competitiveness (e.g., exact make and model of network devices, flow-level traffic statistics, resources reserved for joint services with other operators) or adhering to national security requirements (e.g., policies for traffic prioritization and deep packet inspection, device-level topology). In order to account for the increasing sensitivity of the nature of information shared, the marginal cost of sharing is increasing in I.

Keep in mind that operators only care for a positive payoff in this context, i.e., $u(I) > c(I)$ Moreover, a minimum amount of information $I_{\min} > 0$ is required for establishing a contract (equivalent to the mean or maximum value of a QoS metric as seen in the non-transparent case). Thus, the operating region is characterized by $I_{\min} \leq I < I_{\max}$) as seen in Figure 5. The existence of a social optimum can also be noticed: the aggregate payoff of the two operators is maximal when both of them choose to share I^* taking advantage of the largest distance between utility and cost curves. However, it is easy to see that the only NE of G_1 is $(I_{\min}, I_{\min})$. None of the operators wants to unilaterally share more information as it would result in an increased cost c but the same utility u. Also, neither of them wants to share less information as it would result in no contract being established.

Infinitely repeated continuous sharing game G_r. In the ICT service industry, reputable companies do business on a longer timescale. SLAs usually have a fixed contract period (e.g., a month or a year), which is much shorter than the business lifetime of operators. Moreover, emerging 5G services and the demand for flexibility in network operations could even result in on-demand SLAs with much shorter contract periods. Thus, the act of information exchange happens repeatedly over time calling for a repeated game model. Luckily, a famous result in game theory hints that mutual cooperation could still prevail in the iterated prisoner's dilemma game, under certain conditions [16].

Let us define G_r as the infinitely repeated extension of G_1. In each round r of G_r, the two operators play G_1, but keeping the history of $h|^{r-1}$ in their consideration. In such a game, the social optimum (highest total payoff, at I^*) can be enforced via the so-called *grim trigger strategy* s_g $(g = 1, 2)$. This means that *NO A* plays the optimal $s_1^* = I^*$ until *NO B* deviates

from the social optimum by playing $s_2 \neq I^*$ in round r; then it punishes by playing $s_1 = I_{\min}$ from round $r+1$ forever on (and vice versa), which is decreasing the payoff of *NO B*. This essentially means that the NE strategy s_1 of the stage game G_1 serves as a threat to enforce cooperation in G_r (also known as Nash reversion [16]).

In order to prove this and show the necessary conditions, we apply the one-step deviation principle [16]. The principle states that a strategy profile is a *sub-game perfect equilibrium* (SPE, a strategy profile that represents a Nash equilibrium of every sub-game of the original game) if and only if there exists no profitable one-step deviation. Since we have to look into future payoffs, we use the common method of discounting: payoffs are discounted at step r with a discount factor $\Theta < 1$. This discount factor quantifies the expression of the traditional time preference (balancing the ongoing interest rate), patience (how strategic an operator is with regard to long-term income) and the uncertainty about the length of the game (an operator could go bankrupt). Since the discounted payoff at any step is bounded by $u(I_{\max}) - c(I_{\min})$, the game is continuous at infinity; hence, we can use the one-step deviation principle. *NO B* does not deviate from the optimum (and thus cooperation) if

$$\sum_{i=r}^{\infty} \Theta^i \pi_2^{\text{coop}} > \Theta^r \pi_2^{\text{dev}} + \sum_{i=r+1}^{\infty} \Theta^i \pi_2^{\text{eq}} \tag{3}$$

where π_j^s is the payoff of player j ($j = 1, 2$) employing the strategy s ($s = \{$*coop*eration, *dev*iation, Nash *eq*uilibrium of $G_1\}$) for the given round i. After solving (3) for Θ we get

$$\Theta > \frac{\pi_2^{\text{dev}} - \pi_2^{\text{coop}}}{\pi_2^{\text{dev}} - \pi_2^{\text{eq}}}. \tag{4}$$

The most beneficial one-step deviation at step k is realized by playing $s_2 = I_{\min}$ resulting in a payoff of $\pi_2^{\text{dev}} = u(I^*) - c(I_{\min})$ for NOB. Putting this back to (4) along with $\pi_2^{\text{coop}} = u(I^*) - c(I^*)$ and $\pi_2^{\text{eq}} = u(I_{\min} - c(I_{\min})$, we get:

$$1 > \Theta > \frac{c(I^*) - c(I_{\min})}{u(I^*) - u(I_{\min})}. \tag{5}$$

If (5) holds, *mutually playing grim trigger is SPE for* G_r. Note, that the same discount factor Θ is applicable for NOA due to the symmetric nature of the game.

It is easy to see that $\frac{c(I^*)-c(I_{\min})}{u(I^*)-u(I_{\min})} < 1$ always, as $I^* = \arg\max_I(u(I) - c(I))$ is realized where the derivatives of the concave $u(I)$ and the convex $c(I)$ are equal: $u'(I)|_{I=I^*} = c'(I)|_{I=I^*}$, and $I_{\min} < I^*$. Therefore, it is the relative steepness of the utility versus the cost curve (in the range $[I_{\min}, I^*]$) that determines the exact value of the discount factor. The larger the relative steepness, the lower the bound on the discount factor, meaning that even fairly short-sighted operators (those who strongly discount future payoffs) can sustain an optimal level of information sharing in the long run. For a broad set of utility and cost functions, this operation point is characterized by non-excessive, mutual generosity.

V. A NUMERICAL EXAMPLE

Let us take a closer look at the two-operator example presented in Section III-C with the setup presented in Figure 2, where the SLOin contains a maximum end-to-end delay guarantee θ. We map the example to the repeated continuous information sharing game G_r from Section IV. Recall that $u(I)$ is monotone increasing in I. Since the increase in utility comes from the more precise risk estimation for NO A enabled by more/finer-grained information shared by NO B we can write

$$u(I) = u(\hat{R}(I)). \tag{6}$$

where $\hat{R}$ denotes the estimated risk. Let R^* denote the exact value of risk, then we define the error in risk estimation as

$$\Delta R(I) = \hat{R}(I) - R^*. \tag{7}$$

It is intuitive to assume that $\Delta R(I)$ is also monotone increasing in I (we can do at least as well with more information than with less). As a consequence, both $\Delta R(I^*) < \Delta R(I_{\min})$ and $u(\Delta R(I^*)) > u(R(I_{\min})$ hold.

The impact of information sharing (through proper risk estimation) materializes at the end-user market. Turning to the end-user *C1* of *NO A*, their SLA contains an SLO (delay smaller than θ in our case), the monthly cost of the service (end-user price p, $p = p(\theta)$), the compensation to be paid by *NO A* in case of breaking the SLO (assume this is a constant fraction of p to be credited to *C1* in the following month [6]), and the contract duration (which we ignore for now). It is of key importance for the operator to set an optimal θ, which maximizes the income by setting p properly and avoiding having to compensate more than a few end-users. Observe that with a fully transparent operation mode (all relevant information is shared), the exact risk evaluation R^* of the corresponding correct threshold θ^* and price p^* can be calculated, assuming similar demand across end-users; however, full transparency is not realistic (this is captured in the cost function $c(I)$).

Overestimating the risk. If NOA's error in risk estimation $\Delta R(I) > 0$, i.e., it overestimates the risk. As a consequence, it establishes a loose $\hat{\theta} > \theta^*$ and thus a low $p < p^*$, resulting in lower than optimal revenues. Alternatively it may contract less users resulting in sizable unused capacity and similar low revenues. In order to give a numerical evaluation, we refer back to Table I. Let us assume that an acceptable level of risk for *NO A* is $R = 0.1$. If *NO B* shares only the most basic information on its network ($I_{\min}$, corresponding to being non-transparent), *NO A* will establish an SLO with the maximum delay being $\theta = 260$ ms. However, if mutuality results in a sharing more detailed information (closer to $I_{\min}$, corresponding to, e.g., being semi-transparent) and a smaller $\Delta R(I)$, *NO A* can establish an SLO with $\theta = 208$ ms. This difference can amount to a price (and thus income) decrease between 5% (assuming $p(\theta)$ is increasing logarithmically [17]) and 25% ($p(\theta)$ increasing linearly), given the same user demand.

Underestimating the risk. If *NO A*'s error in risk estimation $\Delta R(I) < 0$, i.e., it underestimates the risk. As a consequence,

it establishes a too strict $\hat{\theta} < \theta^*$ it cannot keep, resulting in having to pay compensation to many users leading to a decreased revenue and a loss of reputation (alternatively, given a high demand for *NO A*'s services, it may contract too many users with a lower θ yielding a similarly result). In order to give a numerical evaluation, we refer back to Table I. Let us assume that an acceptable level of risk for *NO A* is $R = 0.1$. If *NO B* chooses to be semi-transparent, *NO A* will settle for $\theta = 208$ ms. However, if *NO B* chooses to share more information amounting to being transparent, *NO A* obtains a more accurate risk estimation $\hat{R}|_{\theta=208} = 0.11$; therefore, it settles for $\theta = 260$. For *NO A*, this different SLO will result in an expected saving of $1 - 0.03/0.11 = 73\%$ in compensation payments the next month, given similar demand across potential end customers.

It is important to emphasize that the calculations above are meant to numerically illustrate the impact of information sharing on operator revenues. However, as we have made several assumptions and simplifications (costs, revenues, supply/demand, acceptable risk, etc.), the nature of the numerical results are much more important than their absolute values. Furthermore, the precise mapping of semi-transparent and transparent information exchange alternatives to the scale of I (or vice versa) depends on the service context (each of the two could map to I^* for different service deployments) and the exact shape of $u(I)$ and $c(I)$.

VI. Concluding remarks

In this paper we have studied the impact of information exchange through SLAs in multi-operator collaborative service delivery. We have developed a simple risk model with regard to violating the end-to-end QoS guarantees towards the end-user of the service. We have demonstrated how different levels of aggregation for information received from other operators, taking part in the same service chain, can result in both over- and underestimation of the risk. Motivated by this result, we have shown with the help of game-theoretical modeling that incentives for mutually sharing more information do exist, if operators think strategically in the long run. By integrating the risk and the game model, we have demonstrated numerically that the reciprocal exchange of finer grain information is mutually beneficial for all involved operators. Specifically, through a two-operator service chain example with end-to-end delay as the key QoS attribute, we have shown how the lower aggregation level of shared information results in more precise risk estimation, more optimal definition of the user-facing SLA, and thus higher revenues for both operators. A key enabler needed is a platform enabling the desired trust and long-term cooperation among operators as proposed in the 5GEx project [4].

Future work. We believe this work has the potential to open up a new line of research connecting the economics of 5G service delivery, SLA design and network performability. Potential future work may include: the design of architectural enablers for multi-operator service orchestration, information sharing and SLA monitoring; quantifying the impact of different coordination models (federation, alliance) on information sharing; a more realistic game-theoretical (multiple players, Stackelberg game) and risk model (dependability aspects, adding cloud operators); and a proper business impact assessment incorporating market dynamics and monetary values.

Acknowledgement

This work is partly funded by the research lab on *Quantitative modelling of dependability and performance*, NTNU QUAM Lab (http://www.ntnu.edu/telematics/quam). This work has been partly performed in the framework of the H2020-ICT-2014 project 5GEx (Grant Agreement no. 671636), which is partially funded by the European Commission. G. Biczók and L. Toka have been supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences.

References

[1] N. Alliance, "5G white paper," *Next Generation Mobile Networks, White paper*, 2015.

[2] "Website of the 5G Exchange project." [Online]. Available: http://www.5gex.eu

[3] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett, "SDX: A Software Defined Internet Exchange," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14. New York, NY, USA: ACM, 2014, pp. 551–562. [Online]. Available: http://doi.acm.org/10.1145/2619239.2626300

[4] C. J. Bernardos, O. Dugeon, A. Galis, D. Morris, C. Simon, and R. Szabó, "5G Exchange (5GEx) - Multi-domain Orchestration for Software Defined Infrastructures," in *Proceedings of EUCNC*. IEEE, 2015.

[5] "Amazon EC2 Service Level Agreement." [Online]. Available: https://aws.amazon.com/ec2/sla/

[6] "Google Compute Engine Service Level Agreement." [Online]. Available: https://cloud.google.com/compute/sla

[7] P. Bhoj, S. Singhal, and S. Chutani, "SLA management in federated environments," *Computer Networks*, vol. 35, no. 1, pp. 5–24, 2001.

[8] B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. M. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Caceres *et al.*, "The reservoir model and architecture for open federated cloud computing," *IBM Journal of Research and Development*, vol. 53, no. 4, pp. 4–1, 2009.

[9] A. Stanik, M. Koerner, and L. Lymberopoulos, "SLA-driven Federated Cloud Networking: Quality of Service for Cloud-based Software Defined Networks," *Procedia Computer Science*, vol. 34, pp. 655–660, 2014.

[10] J. Yan, R. Kowalczyk, J. Lin, M. B. Chhetri, S. K. Goh, and J. Zhang, "Autonomous service level agreement negotiation for service composition provision," *Future Generation Computer Systems*, vol. 23, no. 6, pp. 748–759, 2007.

[11] Z. Liu, M. S. Squillante, and J. L. Wolf, "On maximizing service-level-agreement profits," in *Proceedings of the 3rd ACM conference on Electronic Commerce*. ACM, 2001, pp. 213–223.

[12] A. Y. Ha and S. Tong, "Contracting and information sharing under supply chain competition," *Management science*, vol. 54, no. 4, pp. 701–715, 2008.

[13] J. Zhang and J. Chen, "Coordination of information sharing in a supply chain," *International Journal of Production Economics*, vol. 143, no. 1, pp. 178–187, 2013.

[14] P. Chołda, E. L. Flstad, B. E. Helvik, P. Kuusela, M. Naldi, and I. Norros, "Towards risk-aware communications networking," *Reliability Engineering & System Safety*, vol. 109, pp. 160 – 174, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0951832012001627

[15] A. Avizienis, J. C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *Dependable and Secure Computing, IEEE Transactions on*, vol. 1, pp. 11–33, 2004.

[16] M. J. Osborne and A. Rubinstein, *A course in game theory*. MIT press, 1994.

[17] H. R. Varian and J. Repcheck, *Intermediate microeconomics: a modern approach*. WW Norton & Company New York, NY, 2010, vol. 6.