Impact of Incentive Mechanisms on Quality of Experience

Andrew Roczniak
Multimedia Communications Research Laboratory (MCRLab)
University of Ottawa
Ottawa, Canada
roczniak@mcrlab.uottawa.ca

Abdulmotaleb El Saddik
Multimedia Communications Research Laboratory (MCRLab)
University of Ottawa
Ottawa, Canada
abed@mcrlab.uottawa.ca

ABSTRACT

Since entities participating in P2P networks are usually autonomous and therefore free to decide on their level of participation, mechanisms to resolve conflicts between individual and collective rationality are needed. How can implementations of such mechanisms be compared? This paper introduces a qualitative reference framework, highlighting essential elements and major design decisions in any implementation of incentive mechanisms. In the context of multimedia applications built on top of P2P architectures, the reference framework can be used in assessing the impact on the quality of experience (QoE) when incentive mechanisms are included.

Categories and Subject Descriptors

H.5.1 [**Multimedia Information Systems**]: Evaluation, methodology; H.3.5 [**Online Information Services**]: Data sharing; C.2.1 [**Network Architecture and Design**]: Distributed networks

General Terms

Design, Economics, Performance

Keywords

Incentives, Peer-to-peer, Streaming

1. INTRODUCTION

In contrast to a client-server model where servers provide services to clients, a peer-to-peer model (P2P) dictates that an entity should act both as a server and a client with respect to other entities. If providing a service entails costs, then the entity acting as a server should expect some form of compensation, or more generally, should have an incentive to offer this service. Incentive mechanisms are however neither fully understood, nor effectively deployed on existing popular P2P networks, with the result that many entities choose to act solely as clients (free-riders). Application of incentive mechanisms in peer-to-peer networks is relatively new. Issues such as improvement of cooperation using techniques from economics and social sciences, applicability of various market models and viability of different revenue models were addressed in [14]. The fundamental difficulty in deploying incentive mechanisms rests in ensuring that they are resilient to failures and robust against malicious entities [5], or entities that do not adhere to a given protocol or strategy [10]. Ensuring that an incentive mechanism is not subverted is complicated by the nature of current peer-to-peer networks, where an entity can join a network with an easily obtainable identity [9]. Large numbers of free-riders on the peer-to-peer network bring up questions about the viability of the network in general, and about its usefulness to other participating entities in particular [8]. As argued in [15], benefits of peer-to-peer networks manifest themselves in massive diversity. When an entity free-rides, it contributes to the reduction of usefulness of peer-to-peer networks. Whenever rational decisions made by an individual entity negatively impact the utility of other entities, incentives become necessary. A lack of common reference point leads to difficulty in comparing characteristics and performance of incentive mechanisms. In order to compare these characteristics, a reference framework is presented in this paper. This framework presents elements common to all incentive mechanisms, and choices available for their design.

The rest of the paper is structured as follows. In section 2 we present a review of related work. We describe the framework in section 3. An example of the impact of incentive mechanisms on QoE is presented in section 4, and section 5 concludes.

2. REVIEW OF RELATED WORK

A qualitative study is presented in [11] where authors observe that incentive schemes play a central role, and propose a classification of incentive patterns. Quantitative studies typically analyze incentive mechanisms in the context of game theory. The performance of three different incentive schemes, token-exchange, peer-approved and service-quality, is analyzed in [12] by defining and using a model based on Multi-Person Prisoner's Dilemma. A formal game theoretic model for systems relying on centralized servers is presented in [6], and provides an analysis of equilibrium of user strategies for various payment schemes. Another study presents a discussion of Nash equilibrium for a non co-operative game between strategic players [2]. Assuming that the probability of obtaining a service from a producer depends on the contribution of a consumer (its production), this study addresses cases when the benefit to a peer from other entities' contribution is common to all, and when the benefit is specific to each peer. In networks with autonomous and transient entities, the supply and demand of services is likely to fluctuate. It would be desirable to design incentive mechanisms to take into account the dynamics of peer-to-peer networks, offering higher compensation for scarce services, and lower for abundant ones. More generally, for each entity having a set of strategies and preferred outcomes, the challenge is to design a utility function that at equilibrium yields a certain desirable network property [4]. This, in essence, is the domain of mechanism design. Work presented in [16] designs a utility function, and provides a solution using control theoretic approach.

3. REFERENCE FRAMEWORK

3.1 Motivation

Understanding why an entity interacts and collaborates with another entity is paramount to building effective incentive mechanisms. In effective incentive mechanisms, an entity should expect to either give or receive some form of compensation (see Fig 1). An obvious form of compensation is an exchange of money for the service. Payment systems assume the existence of a currency that can be used outside of the network. If entities are capable of making properly informed choices in deciding on the appropriate level for the price of a service, then free riding should not exist. If there is no payment system, then entities can be motivated to participate in a network by either "greed" or "fear", or possibly a combination of both. On one hand, an entity may be enticed to contribute to the network by threats of sanctions and restrictions. Methods in this branch would work best in the context of networks where membership is restricted and participants have a well-known identity, such as communities or clubs. On the other hand, an entity's intrinsic need or desire (eagerness) to obtain a service may be leveraged. Another form of compensation can be based on trade, or payment-in-kind. A producer node can be compensated with another service or a promise of a service in the future. This promise must be worthless outside of the network. The motivation element can therefore be based on some concept of "proxy" currency, on barter or some combination of those characteristics.

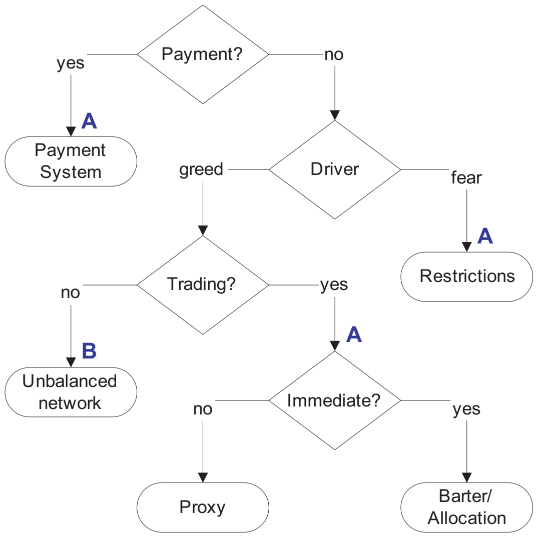

Figure 1: Motivation. Points "A" is where an entity can be expected to be motivated to collaborate in a network. Point "B" illustrates two common problems, tragedy of the commons and free-riding.

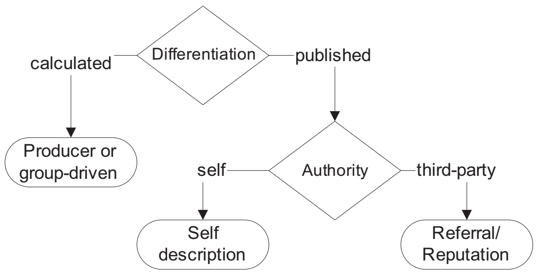

Figure 2: Differentiation

3.2 Differentiation

The ability of a network to provide differentiation of services, of quality of service or between entities is necessary for designing viable incentive mechanisms. The concept of differentiation is not dependent on what exactly is chosen as basis to assess differences. Two main choices exist for differentiation (see Fig. 2). The first, places the burden of describing a service and its cost on producers. In the simplest form, an entity is responsible for making its own claim. Claims can also be made by a third-party, which can be trusted to varying degrees. If the third-party is not fully trusted, incentive mechanisms can rely on referral or reputation infrastructures [3]. If the third-party is fully trusted, design of incentive mechanisms becomes easier but at the expense of shifting complexity to the design and deployment of the trusted third-party itself. The second involves entities calculating what class of service to offer to other entities based on factors such as credentials, the result of optimization of an objective function, or a combination of those two methods. This choice implements an explicit reward and punishment mechanism where entities make decisions based either on their personal goals, or on some common group-specific goals.

3.3 Peer Selection

Peer selection element is the "glue" that holds the differentiation and utility calculation elements together, and indeed, it is sometimes subsumed by one of those neighbors. Why is the peer selection element important then? Considering the transient nature of entities in a peer-to-peer network, we immediately see that some producers who are providing a service to consumers can leave the network, or new producers with possibly better qualifications can join the network. It would be beneficial for consumers to be able to monitor and seek out better partners for interaction and collaboration. This is also the case when producers select consumers with whom to interact. In instances where a producer optimizes an objective function, it might obtain better results if it replaces some consumers by others. The goal of this element is the selection of an entity most likely to lead to a successful interaction. The basic design choice defines whether an entity's interaction with other entities is static or dynamic.

3.4 Utility Calculation

Selection of a utility function is one of the fundamental decisions for an incentive mechanism as it will influence all interaction and collaboration between entities. Each entity could define its own utility function, but for the sake of clarity, we will assume that such functions will be defined in terms of specific applications' requirements, or based on well defined goals. In [1], for example, this function is defined in terms of cost and requirements as specified by parallel downloading or streaming of audio/video objects. Despite the potentially unlimited number of utility functions, the basic design choice of the utility calculation element defines whether an entity explicitly acts in its own interest or the interest of some group when optimizing an objective function.

3.5 Metric Scope

This element includes the definition of parameters used in the utility calculation element. Quite often, it is very helpful to think of this layer as defining a "proxy", or an intermediate store-of-value used in valuation of services exchanged between entities. This occurs for example in reputation frameworks, where an entity is judged by others depending on their reputation. Reputation itself can change depending on the amount and quality of supplied services. Although many metrics can be used, the basic characteristic of this element is the scope of such metrics within peer-to-peer networks. A group in a network can share the same metric, or each entity can create and use its own. Referral systems are examples of systems that need to carefully manage the semantics of metrics, while simple machine-learning implementations use metrics that have local meaning only.

4. IMPACT OF INCENTIVES ON QOE

Addition of incentive mechanisms to multimedia communication systems built on peer-to-peer architecture has the potential to increase the robustness of various applications. To compare relative qualities of incentive-enabled systems various QoS parameters can be used, but ultimately, the level of satisfaction of a user will determine if a system is good or not [13]. The experience of a user while interacting with a particular multimedia communication system can be quantified in a single measure, the Quality of Experience (QoE). Incentive mechanisms will impact the user's perception of the interaction in various ways. In order to start defining quantitative parameters, four questions can be asked. Does the user's benefit increase with increasing contribution? Do users with lower level of contribution receive lower benefit? Is a change in contribution reflected on the user's benefit in a timely manner? Is the received benefit stable if the contribution level is also stable? In general, if the answer is negative to those questions, the incentive mechanism does not improve the QoE as perceived by the user. The following discussion presents an incentive mechanism and how its characteristics may impact QoE.

4.1 Incentives in Media Streaming

Work presented in [7] describes an incentive system for peer-to-peer media streaming applications. Based on the insight that random peer selection provides random quality, the authors propose a system in which collaborative nodes are rewarded with higher flexibility in peer selection.

Motivation The incentive system is based on participants' score which increases when providing a service. Nodes with a higher score have more flexibility when selecting a producer, which leads to higher streaming quality. Motivation is thus based on proxy trading.

Differentiation Entities publish their scores. In order to mitigate potential cheating and collusion, the authors suggest using reputation systems. Producers also differentiate themselves by publishing their availability and offered bandwidth rates which are also required to predict the expected streaming quality.

Peer selection Consumers locate producers willing to provide them with a service. Those producers are classified as active and standby, and the consumer can select producers that best match its requirements at any time. Peer selection is therefore dynamic. Although not explicitly stated, peer selection appears to be dynamic at the producer as well. A producer could for instance stop providing services to an existing consumer if a more suitable consumer made a request.

Utility calculation Entities choose their contribution level in order to maximize their own utility. The utility function is defined as $U_i(x_i) = a_i Q(x_i) - b_i C(x_i)$ where Q and C are the quality and cost functions dependent on the contribution level x. The constrained optimization is: $max\, U_i(x_i)$ subject to: $U_i \geq 0$ for $x_i \geq 0$ The cost function is based on bandwidth and storage usage $C(x_i) = (c_L + c_T)B_{out}$ where c_L, c_T are the unit storage and unit transmission costs respectively, while the quality function is defined as a monotonic non-decreasing in user score, asymptotically reaching a value of Q_{MAX} and having a non-negative initial value $Q_{BestEffort}$. User's contribution first gives a certain score. That score is then compared to the scores of other entities, yielding a certain percentile ranking R_i. The resulting quality function is: $Q_i(R_i) = Q_{MAX}$ when $R_i \geq a$, and $\frac{R_i}{a}(Q_{MAX} - Q_{BestEffort}) + Q_{BestEffort}$ otherwise. By increasing its contribution, a producer expects to increase its ranking as compared to other producers, and thus its quality function.

Metric scope The metric is each entity's score, which is globally known. The score can be computed in different ways, but is a function of the entity's contribution. Supply/demand conditions on the network could be reflected in variable rewards for the same contribution, or a punishment applied for refusing to honor a request. Furthermore, scores can be subjected to decay over time.

4.2 Analysis

An entity will co-operate until it obtains a ranking that makes its quality function equal to Q_{MAX}: increasing costs would make its utility smaller if it continues to co-operate. As other entities co-operate, their score and thus their ranking increase, lowering the relative ranking of the non co-operating entity. By increasing its co-operation, an entity can then maximize its utility. In cases where other entities are also co-operating, an entity might not be able to increase its ranking substantially, while still incurring costs associated with providing services. A situation would then arise where an entity, even willing to co-operate, is better off by not co-operating even when its quality function is close

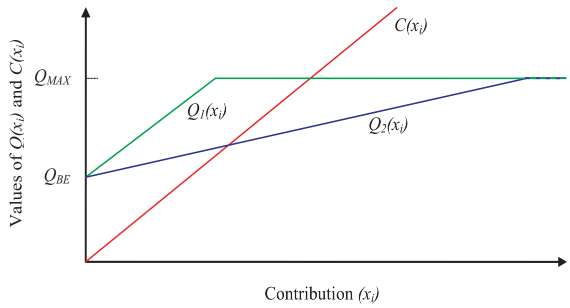

Figure 3: Score inflation illustrates the problem when costs rise faster than ranking

to $Q_{BestEffort}$. This situation of 'score inflation' could be caused by malicious nodes and is illustrated in Fig. 3.

Since each entity is responsible for maximizing its own utility, the following situations are valid. An entity with high score can reply to an entity with any score - including lower score. Similarly an entity does not have to honor any request, even if it comes from an entity with higher score. The result is that unless entities have a social conscience and give preference to higher scores, higher scores do not mean that a producer will honor requests for service, or that it won't serve another consumer with a lower score instead, undermining the basic premise of the motivation element.

Since each entity reports its own score, the temptation to misrepresent it in order to gain an advantage are high since tangible benefits are at stake. Additionally, the risk of collusion between malicious entities needs to be taken into account. The authors identify reputation management systems as a possible solution to those problems. All the scores are then weighted by the degree of trust and/or reputation of the reporting and referring entity. Since both the score and trust-worthiness of an entity reflect its degree of co-operation, collecting scores appears then to be redundant.

5. CONCLUSION

This paper presents a qualitative reference framework for incentive systems. All such systems can be deconstructed into five elements, incentive, differentiation, peer selection, utility calculation and metric scope. Each element presents some design choices that will define characteristics of an incentive mechanism. This reference framework can be used in assessing the impact of incentive mechanisms on QoE of multimedia applications.

6. ACKNOWLEDGMENTS

The authors acknowledge the financial support of Natural Sciences and Engineering Research Council of Canada, LORNET and the Government of Ontario (PREA).

7. REFERENCES

[1] M. Adler, R. Kumar, K. Ross, D. Rubenstein, D. Turner, and D. Yao. Optimal peer selection in a free-market peer-resource economy. In *Second Workshop on the Economics of Peer-to-Peer Systems, Harvard University*, 2004.

[2] C. Buragohain, D. Agrawal, and S. Suri. A game theoretic framework for incentives in p2p systems. In *IEEE International Conference on Peer-to-Peer Computing, Linköping, Sweden*, 2003.

[3] P. Dewan. Peer-to-peer reputations. In *12th International Workshop on Parallel and Distributed Real-Time Systems, Santa Fe, USA*, 2004.

[4] J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 1–13. ACM Press, New York, 2002.

[5] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *5th ACM conference on Electronic Commerce, New York, NY, USA*, pages 102–111, 2004.

[6] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for sharing in peer-to-peer networks. In *ACM Conference on Electronic Commerce, Tampa, USA*, 2001.

[7] A. Habib and J. Chuang. Incentive mechanism for peer-to-peer media streaming. In *12th IEEE International Workshop on Quality of Service, Montreal, Canada*, 2004.

[8] R. Krishnan, M. D. Smith, Z. Tang, and R. Telang. The impact of free-riding on peer-to-peer networks. In *Hawaii International Conference on System Sciences*, 2004.

[9] S. Marti and H. Garcia-Molina. Identity crisis: Anonymity vs. reputation in p2p systems. In *IEEE International Conference on Peer-to-Peer Computing, Linköping, Sweden*, 2003.

[10] C. Ng, D. Parkes, and M. Seltzer. Strategyproof computing: Systems infrastructures for self-interested parties. In *Workshop on Economics of Peer-to-Peer Systems, Berkeley, USA*, 2003.

[11] P. Obreiter and J. Nimis. A taxonomy of incentive patterns, the design space of incentives for cooperation. In *International Workshop on Agents and Peer-to-Peer Computing, Melbourne, Australia*, 2003.

[12] K. Ranganathan, M. Ripeanu, A. Sarin, and I. Foster. To share or not to share, an analysis of incentives to contribute in collaborative file sharing environments. In *Workshop on Economics of Peer-to-Peer Systems, Berkeley, USA*, 2003.

[13] L. A. Rowe and R. Jain. ACM SIGMM retreat report on future directions in multimedia research. 2004.

[14] B. Strulo, A. Smith, and J. Farr. An architecture for peer-to-peer economies. In *IEEE International Conference on Peer-to-Peer Computing, Linköping, Sweden*, 2003.

[15] A. Vahdat, J. Chase, R. Braynard, D. Kostić, P. Reynolds, and A. Rodriguez. Self-organizing subsets: From each according to his abilities, to each according to his needs. In *International Workshop on Peer-to-Peer Systems, Cambridge, USA*, 2002.

[16] W. Wang and B. Li. To play or to control: A game-based control-theoretic approach to peer-to-peer incentive engineering. In *Lecture Notes in Computer Science Vol. 2707*, pages 174–192. ACM Springer-Verlag, 2003.