INCA: qualitative reference framework for incentive mechanisms in P2P networks

Andrew Roczniak,* Abdulmotaleb El Saddik and Pierre Lévy

Multimedia Communications Research Laboratory and
Collective Intelligence Laboratory,
University of Ottawa,
Ottawa, Canada
E-mail: roczniak@mcrlab.uottawa.ca
*Corresponding author

Abstract: The existence of peer-to-peer networks is due to benefits brought by decentralisation of control and distribution of resources. It is expected that the usage of such networks will grow and provide support for a variety of applications, including collaborative environments. Since entities participating in those networks are 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 INCentive frAmework (INCA), a qualitative reference framework, highlighting essential elements and major design decisions in any implementation of incentive mechanisms. In the context of collaborative environments built on top of P2P architectures, the reference framework can be used in assessing the impact on the quality of experience of applications when incentive mechanisms are included.

Keywords: collaborative environments; peer-to-peer systems; incentives.

Reference to this paper should be made as follows: Roczniak, A., El Saddik, A. and Lévy, P. (2007) 'INCA: qualitative reference framework for incentive mechanisms in P2P networks', *Int. J. Computer Applications in Technology*, Vol. 29, No. 1, pp.71–80.

Biographical notes: Andrew Roczniak is pursuing his PhD in Electrical Engineering at the University of Ottawa. His research interests include peer-to-peer systems, multimedia communication systems and distributed computing.

Abdulmotaleb El Saddik is an Associate Professor in the School of Information Technology and Engineering (SITE). He is the Director of the Multimedia Communications Research Laboratory (MCRLab) and the Director of Information Technology Cluster, Ontario Research Network on Electronic Commerce (ORNEC). He received an MSc (Dipl-Ing) and a PhD (Dr-Ing) in Electrical Engineering and Information Technology from Darmstadt University of Technology, Germany in 1995 and 2001, respectively. He has authored or co-authored two books and more than 70 publications in the areas of software engineering development of multimedia artefacts and collaborative virtual environments. He is a Senior Member of IEEE and a recent winner of the prestigious Canadian Premier's Research Excellence Awards (PREA).

Pierre Lévy is a full-time Professor in the Department of Communication at the University of Ottawa, Fellow of the Royal Society of Canada, Canada Research Chair in Collective Intelligence and Director of the Collective Intelligence Laboratory at the University of Ottawa. He is the author of 12 books, including '*Collective Intelligence*' and numerous scientific papers.

1 Introduction

In contrast to a client-server model where servers provide services to clients, a Peer-to-Peer (P2P) model 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) (Adar and Huberman, 2000).

Although incentive mechanisms are well documented in business, economy and sociology literature (Bamberg and Spremann, 1989), their application to P2P 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 in P2P networks were recently addressed in Strulo et al. (2003), Senior and Deters (2002) and Hummel et al. (2003). The fundamental difficulty in deploying incentive mechanisms rests in ensuring that they are resilient to failures and robust against malicious entities (Feldman et al., 2004). Since entities are expected to be independent and rational, there is no guarantee that they will adhere to a given protocol or mandated strategy (Ng et al., 2003; Shneidman and Parkes, 2003). Ensuring that an incentive mechanism is not subverted is complicated by the nature of current P2P networks, where an entity can join a network with an easily obtainable identity (Douceur, 2002; Friedman and Resnick, 2001; Marti and Garcia-Molina, 2003).

Some file-sharing P2P networks do not deploy incentive mechanisms and yet are very successful. It is natural then to ask if incentive mechanisms are truly necessary and not purely of academic interest. It can be argued that the success of file-sharing applications also depends on other factors besides the level of free-riding. The cost of providing services (files) on existing file-sharing networks is close to zero and the expected quality of service correspondingly low. Coupled with the compelling nature of the exchanged files, this creates a very useful network for entities to join. However, if the cost of providing services or the demanded quality of service rises, it is unlikely that a P2P network without an incentive mechanism can be successful.

Large numbers of free-riders on the P2P network bring up questions about the viability of the network in general and about its usefulness to other participating entities in particular (Krishnan et al., 2004). As argued by Vahdat et al. (2002), benefits of P2P networks manifest themselves in massive diversity. When an entity free-rides, it contributes to the reduction of this diversity and therefore to the decrease of usefulness of P2P networks (Bhagwan et al., 2003). Whenever rational decisions made by an individual entity negatively impact the utility of other entities, incentives (i.e. reward or punishment) become necessary to ensure that the usefulness of a P2P network does not diminish. Practical implications of incentive mechanisms thus relate to increasing collaboration and availability of services and target applications such as file-sharing (Condie et al., 2004; Ma et al., 2004), distributed storage (Ngan et al., 2003) and streaming (Habib and Chuang, 2004).

In the context of collaborative environments built on top of P2P architectures, incentive mechanisms will therefore have an impact on deployed applications. A natural interest then is in the understanding of this impact and in deriving a quantifying measure of its effect. 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. We further outline one possible way of deriving a quantifying measure, based on the QoE of applications as perceived by the user.

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 evaluation of the framework based on existing incentive mechanisms is presented in Sections 4, and 5 concludes.

2 Related work

The body of literature describing frameworks for incentive mechanisms encompasses both quantitative and qualitative studies. In the latter case, (Obreiter and Nimis, 2003) considers distributed systems composed of autonomous entities cooperating to achieve their goals. Authors observe that incentive schemes play a central role and propose a classification of incentive patterns. Matching these incentive patterns with specific application environment allows for a more systematic approach to the design of incentive schemes.

Quantitative studies typically analyse incentive mechanisms in the context of game theory (Ross, 2004). The performance of three different incentive schemes, token-exchange, peer-approved and service-quality, is analysed by Ranganathan et al. (2003) by defining and using a model based on Multi-Person Prisoner's Dilemma. A generalised form of Evolutionary Prisoner's Dilemma model is used to study cooperation in P2P networks (Lai et al., 2003). Authors conclude that even with unfettered and decentralised creation of identities, adaptable incentive mechanisms can produce systems that converge to full cooperation. Furthermore, scalability and security of networks is addressed in the context of private and shared history of entities' actions. A formal game theoretic model for systems relying on centralised servers is presented by Golle et al. (2001) and provides an analysis of equilibrium of user strategies for various payment schemes. With the view to assess and evaluate incentive systems in P2P networks, (Antoniadis et al., 2004) presents a model of utilities obtained by entities and various schemes to address the lack of incentives. Analytical expressions are provided for cases where there is a lack of complete payoff information and an inability to individually compensate entities. Authors' main findings relate to the size of the network and variability of entities' utility functions. Another study presents a discussion of Nash equilibrium for a non-cooperative game between strategic players (Buragohain et al., 2003). 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 P2P 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 (Feigenbaum and Shenker, 2002). This, in essence, is the domain of mechanism design. Work presented by Wang and Li (2003) designs a utility function and provides a solution using control theoretic approach.

3 Reference framework

3.1 *Reference framework elements*

All the incentive mechanisms have the following elements. Four of those elements are essential and one is optional.

Motivation: there must be a clear motivation for entities to interact and collaborate. Usually, motivation is provided by a prospective gain for an entity from joining and participating in a network. Since entities produce and consume services, motivation needs to be assessed from both consuming and producing perspective. This element is therefore concerned with defining the motivation for entities to participate in a P2P network.

Differentiation: for an incentive mechanism to function properly there must be a concept of differentiation, either of services, quality of service or between entities. Just as we can select a bank with the best terms for a line of credit and banks can choose to whom and at what rate they extend a line of credit, participants in a P2P network must be capable of observing and acting upon perceived differences. This element is therefore concerned with the manner in which differentiation is defined and made accessible to participating entities.

Peer selection: the goal of this element is the selection of an entity (or entities) most likely to lead to a successful interaction. In certain applications it is possible, indeed desirable, to obtain services from multiple entities concurrently or to dynamically seek out better entities to interact with. In those cases, this element acts as 'glue' for the differentiation and utility calculation elements. In other cases, it might not be necessary to include this element in incentive mechanisms.

Utility calculation: defining an entity's perception of derived utility from producing or consuming services is one of the fundamental issues for an incentive system. It is however possible to come up with potentially infinite number of definitions since the exact utility derived by an entity can depend on targeted application, various assumptions about available knowledge or conditions or even intangible factors. This element is therefore concerned with the characteristics of objective functions rather than their specific definitions.

Metric scope: some parameters are necessary for the preceding element. Incentive mechanisms may choose various parameters as metric, depending on the target application or availability of variables. Combining different parameters into a composite parameter is also possible. The metric can also be thought of as a proxy or store of value, for services rendered or consumed. The main concern of this element is the scope of the metric within the P2P network and not its specific definition.

Elements describing characteristics of incentive mechanisms are presented in Figure 1. They are represented as layered one on top of another to highlight their interdependency.

Figure 1 Framework elements

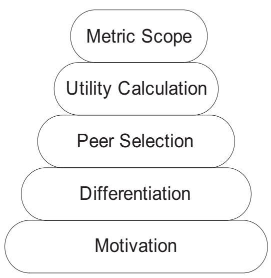

3.2 *Motivation*

Understanding why an entity should interact and collaborate with another entity is paramount to building effective incentive mechanisms.

Two popular file-sharing P2P networks, Kazaa (Leibowitz et al., 2003) and BitTorrent (Cohen, 2003) allow users to exchange files. Their protocols for doing so are different, but in both cases we can think of the offered service as the permission to upload contents from a storage device using some (preset) bandwidth. In the first network, an entity can access offered services while not offering any services itself. An attempt to distinguish between classes of entities based on their participation level was defeated by the release of a software client implementation (e.g. KazaaLite) which always reported its participation level as the maximum defined. In the second network, entities are rewarded with faster download speed in proportion to how fast they allow other entities to upload from them. An entity however can easily receive more services than it actually produces by disconnecting from the network after receiving the requested file.

It can be argued that these two networks do not have robust enough incentive mechanisms and a quick look at supply and demand curves illustrates the problem. Considering the price of consuming a service to be close to zero, the demand for this service can grow arbitrarily high. Since the price of supplying a service remains constant regardless of the quantity supplied, the equilibrium point would then be at an arbitrarily high level of the quantity of service. Assuming that the number of entities willing to supply a given service for free is limited, a gap between the quantity of service supplied and demanded will result. The net effect is that a number of consumers will not be able to fulfil their need for a service. Furthermore, every time a service is provisioned, the server entity incurs a cost that decreases its ability to keep offering services. If the P2P network does not continuously add such producer entities, then the network's capacity to offer services also decreases. It then follows that each entity consuming a service reduces the utility of the network to other consumer entities. A situation where there is a mismatch between supply and demand, and where each fulfilled demand reduces the utility of all, but where all consumer entities are marginally better in the short-term by consuming as much as they can, will potentially lead to the 'tragedy of the commons', or in other words, the collapse of the P2P network. If such a drastic outcome does not come to pass, we are still faced with the issue of entities consuming more of the free services than their 'fair share', or by choosing to restrict their contribution to the network, to avoid contributing their 'fair share' of the cost. Such entities would in fact get a free-ride on the P2P network.

In effective incentive mechanisms, an entity should expect to either give or receive some form of compensation. Producers willing to offer services for free do exist, for example as a promotional give-away or as a loss leader strategy. Similarly, entities consuming a free service may decide to make a donation, in effect buying a service from a producer. Those examples however remain an exception rather than the rule and entities will consume services for free if they can, while producers would prefer a more predictable revenue stream.

An obvious form of compensation is an exchange of money for the service, just as in the case of cash or credit card usage. Payment systems assume the existence of a currency that can be used outside of the network. Participating peers then have a motivation to offer services in order to receive this currency. 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. Such systems assume the existence of a robust framework for payment processing, which in P2P networks presents issues of its own such as, for example, reliable accounting mechanisms (Hausheer et al., 2003). An example of an implementation based on a payment system is outlined by Balazinska et al. (2004).

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. Ideas outlined by Serjantov and Lewis (2003) could be applied to a potentially open and anonymous membership and are based on client puzzles that participants might be asked to solve in exchange for services. On the other hand, an entity's intrinsic need or desire (eagerness) to obtain a service may be leveraged. Methods in this branch can assume that participating entities freely decide on the level of their participation. An entity might be eager to, for example, obtain some files from a file-sharing network or store its data on the network in such manner that it is both secure and highly available.

Another form of compensation can be based on trade, or payment-in-kind. Entities trade if they allow access to each other's storage, bandwidth or computing power. A producer node can be compensated with another service or a promise of a service in the future. In both cases, entities are required to assign a value to their services, especially if these are not heterogeneous. In the first case, entities barter (negotiate) or allocate some portion of their capacity to perform a service for other services, while in the second case a producer entity obtains a voucher redeemable at a later time from the network or the consuming entity. Such vouchers 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. If there is no trading between entities or in some cases trading takes place but is inefficient, a situation allowing extensive free-ride exists, creating an unbalanced network.

Whenever there is potential to receive something tangible, such as money or service, participating entities have a motivation to collaborate. Figure 2 presents key design choices for the motivation element of incentive mechanisms. 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.

3.3 Differentiation

The ability of a network to provide differentiation of services, 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. Without loss of generality then, the following discussion will use quality of service as example.

Figure 2 Motivation

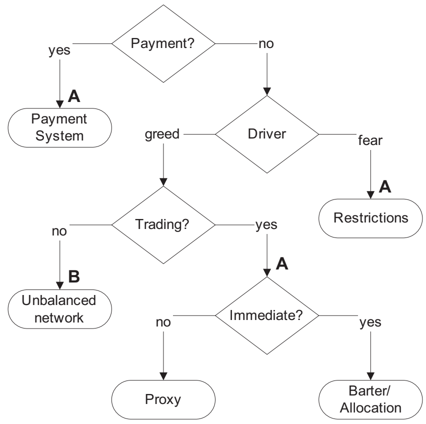

Two main choices exist for differentiation. The first, already familiar to most shoppers, places the burden of describing a service and its cost on producers. Consumers then need to understand what is offered and how to compare offers. The second involves producers calculating what class of service to offer to consumers. These calculations can be based on factors such as credentials supplied by consumer, the result of optimisation of an objective function or a combination of those two methods. For example, in the case of a customer with a bad credit history, a bank might decide to extend a line of credit at a punishing interest rate, if at all.

In the case of the first choice, we can further analyse this element based on which entity describes a particular service. In the simplest form, an entity is responsible for making its own claim (Dewan and Dasgupta, 2004). 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 (Dewan, 2004; Singh and Liu, 2003; Terreen and Hills, 2001; Wang and Vassileva, 2003). 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 choice implements an explicit reward and punishment mechanism. Producers can make decisions based either on their personal goals or on some common group-specific goals. In the latter case, a group of producers ensures similarity of service offered to comparable consumers. Since the consumer does not know a priori what quality of service to expect, its natural behaviour will be to try different producers and keep the best one.

Figure 3 presents key design choices for the differentiation element. By introducing the notion of receiving something tangible in compensation for services in the preceding layer, we automatically require the concept of differentiation in an incentive system. This element therefore describes the manner in which differentiation is defined and made accessible in the network.

3.4 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 neighbours. Basic design choices of peer selection element are presented in Figure 4.

Figure 3 Differentiation

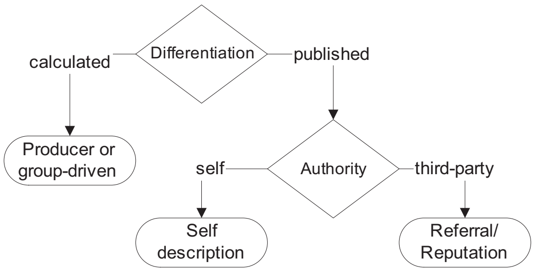

Figure 4 Peer selection

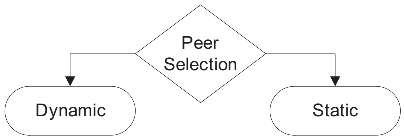

In simple cases where a consumer is presented with a choice of producers, peer selection reduces to selecting a producer that maximises an objective function of the consumer. Similarly, if the differentiation element is defined by a probability function for a producer, the peer selection element reduces to the choice of serving or not serving a particular consumer.

Why is the peer selection element important then? Considering the transient nature of entities in a P2P 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 optimises 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.

3.5 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 Adler et al. (2004), for example, this function is defined in terms of cost and requirements as specified by parallel downloading or streaming of audio/video objects. In Bernstein et al. (2003) on the other hand, this function tries to minimise the time necessary to select the best source for a file download.

Despite the potentially unlimited number of utility functions, the basic design choice of the utility calculation element is quite simple and is illustrated in Figure 5. Each entity will optimise an objective function that results in highest utility as perceived by itself or by a group. In the former case, each entity is independently making rational decisions and hence exemplifies a 'free-for-all' concept. In the latter case, each entity is making decisions that are rational in a collective context and thus tries to optimise 'social welfare' of the group. A group can be the whole or some subset of a P2P network.

This element therefore defines the characteristic of the chosen utility function, whether an entity explicitly acts in its own interest or the interest of some group.

Figure 5 Utility calculation

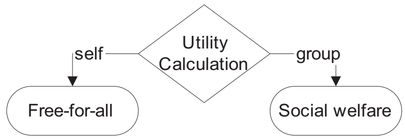

3.6 Metric scope

This element includes the definition of parameters used in the utility calculation element. For example, in the case of file-sharing or network storage applications entities can exchange information about their available upload bandwidth or available disk space. 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. The basic design choice of this element is illustrated in Figure 6.

Figure 6 Metric scope

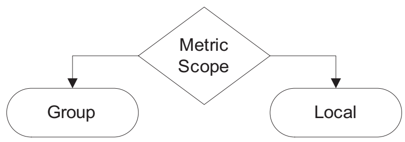

Although many metrics can be used, the basic characteristic of this element is the scope of such metrics within P2P networks. A group in a network can share the same metric or each entity can create and use its own. Referral systems are examples of systems that need to carefully manage the semantics of metrics, while simple machine-learning implementations use metrics that have local meaning only.

If the chosen metric is an abstract concept such as 'reputation' or 'trust', a corresponding model, associated data structures and operations need to be defined. Work presented by Abdul-Rahman and Hailes (2000) assists users in identifying trustworthy entities, and gives them the ability to enhance their understanding of other entities' subjective recommendations. A method to assign each entity in a network a unique global trust value in a distributed and secure manner is presented in Kamvar et al. (2003) and works by aggregating local trust scores of all entities, reflecting their experience with a specific entity.

The main concern of this element is therefore the scope of the metric within the P2P network. In cases where a group of entities share a metric, this element may also be concerned with storing and managing metric information in a decentralised manner.

4 INCA evaluation

Addition of incentive mechanisms to collaborative environments built on a P2P architecture has the potential to increase the robustness of various applications. To compare relative qualities of incentive-enabled systems many QoS parameters can be used, but ultimately, the level of satisfaction of a user will determine if a system is good or not (Rowe and Jain, 2004). The experience of a user while interacting with a particular collaborative environment can be quantified in a single measure, the QoE. This measure is in most cases a function of multiple variables and is quite complex since it takes into account user's perception of interaction.

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 four examples of incentive mechanisms and how their characteristics may impact QoE.

4.1 INCA and game theoretic approach to file sharing

Work presented by Ma et al. (2004) strives to offer service differentiation dependent on the amount of each entity's contribution to the P2P network.

Motivation: in this case, peers trade by means of a proxy, present production for future consumption. Consumers compete for resources provided by the producer. The producer allocates more resources to the consumer whose contribution level is higher related to other competing consumers.

Differentiation: each entity has publicly available credentials supplied by a trusted third-party. Producers use these to calculate how many resources should be offered to a consumer, depending on the group currently being served. Producers' differentiation is not specifically addressed, but since their contribution levels and available bandwidth are known, consumers can presumably devise strategies that take this information into account.

Peer selection: peer selection is not discussed from consumers' perspective: it is ancillary to the differentiation element. Producers tolerate joining and leaving consumers and stop interacting with a consumer that does not provide proof of service. Producers attempt to distribute their resources according to a specific formula among all consumers, but do not make any attempt to interact only with a subset of requesting consumers; as such producers have a static peer selection element: it is subordinate to the utility calculation element.

Utility calculation: consumers' utility calculation is not explicitly discussed. Instead, if consumers follow the given protocol (and have a non-zero contribution level), they can be guaranteed a share of resources as calculated by the producer. The producer computes consumers' utilities and tries to optimise for all connected consumers. By maximising the sum of connected consumers' utilities, a producer guarantees to maximise its contribution level. A producer is in fact maximising its contribution level by maximising the utilities of connected peers.

The producer assumes that each connected consumer has the following utility function: $U_i(x_i) = \log(x_i/b_i + 1)$ where $x_i \in [0, b_i]$ and is the allocated resource to consumer i, and b_i is the requested resource by consumer i. The producer then performs the following constrained optimisation: $\max \sum_{i=1}^{N} C_i U_i(x_i)$ subject to: $\sum_{i=1}^{N} x_i \leq W_A$, where $x_i \in [0, b_i] \forall i$, W_A is the total available resource at the producer and C_i is the contribution level of consumer i. The authors use a competition game model where C_i and W_{Ai} are known to all participants and prove that $b_i^* = W_A C_i / \sum_{j=1}^{N} C_j$ for $i = 1..., N$ is a Nash equilibrium. Similarly, the authors prove that the resource allocation to consumers (x_i) maximises the producer's contribution level.

Metric scope: metric is based on globally known quantities: contribution level of an entity (C_i) and the amount of bandwidth resource provided by a producer (W_{Ai}). The contribution level is calculated based on evidence of rendered service: digitally signed requests for a specific amount of resource from a consumer (b_i).

4.1.1 Analysis

This section highlights some further research areas for the preceding publication based on INCA.

- The same consumer might be offered varying amounts of resources by different producers since at equilibrium $b_i^* = x_i^*$ for $i = 1..., N$ and which depends on the contribution level of all connected consumers to a particular producer offering W_A. It would be in the consumers' best interest to search for the most accommodating producer. This can lead to an impact on the overhead communication between entities.
- Producers do not necessarily need to maximise their contribution level and can simply provide some resources to a consumer for its contribution level to go up. The resulting contribution level might be enough to satisfy an entity that subsequently wishes to consume resources (it is guaranteed not to be resource starved). This is due to the fact that a producer does not know a consumer's true utility function, assuming $U_i(x_i) = \log(x_i/b_i + 1)$ instead. If such a situation

occurs and is persistent, then an entity, even conforming to all of the protocol's rules, cannot be guaranteed to receive its fair share of resources undermining the basic premise of the motivation element.

- If there is no concept of decay of the contribution level, producers can build it up and then free-ride, undermining the basic premise of the motivation element.
- Over the long-term and as contribution levels of entities increase, new participants might have to offer resources for a long time before they can successfully compete with more established consumers. This could be unfair and would undermine the basic premise of the motivation element.
- Producers are interested in maximising their contribution level. Maximising the sum of consumers' utilities is a more complicated way of just maximising the usage of offered service: a producer might find it more expedient to provide resources to only one consumer if $b_i \geq W_A$ rather than calculating precise allocations to multiple consumers. By not following the prescribed mechanism, such producers would undermine the basic premise of the motivation element.
- Two entities can collude by sending evidence of an improbably high allocation of resources. An entity would then have a very high level of contribution since there is no way to verify or exclude that evidence. Such a security loophole would undermine the basic premise of the motivation element.
- Rare files versus popular files: if an entity shares rare files – generating less demand for its resources – its contribution level might be smaller than if it were offering popular files. In order to increase its contribution level, a producer could switch from offering rare files to offering popular files. This can lead to an impact on the diversity offered by the network.

4.2 *INCA and incentives in media streaming*

Work presented by Habib and Chuang (2004) describes an incentive system for P2P 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 such as Feldman et al. (2004) or Kamvar et al. (2003). 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 maximise 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 optimisation 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 $C(x_i) = (c_L + c_T)B_{\text{out}}$ 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 $Q_{\text{BestEffort}}$. 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: $Q_i(R_i) = Q_{\text{MAX}}$ when $R_i \geq a$ and $R_i/a(Q_{\text{MAX}} - Q_{\text{BestEffort}}) + Q_{\text{BestEffort}}$ 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 honour a request. Furthermore, scores can be subjected to decay over time.

4.2.1 *Analysis*

This section highlights some further research areas for the preceding publication based on INCA.

- An entity will cooperate 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 cooperate. As other entities cooperate, their score and thus their ranking increase, lowering the relative ranking of the non cooperating entity. By increasing its cooperation, an entity can then maximise its utility. In cases where other entities are also cooperating, 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 cooperate, is better off by not cooperating even when its quality function is close to $Q_{\text{BestEffort}}$. This situation of 'score inflation' could be caused by malicious nodes and is illustrated in Figure 7.
- Since each entity is responsible for maximising 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 would not 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 cooperation, collecting scores appears then to be redundant.

- In order to compute the utility function, an entity has to compute its ranking within the network. This computation requires $O(\sigma^2)$ samples, where σ is the standard deviation of the population. This σ is a priori not known and thus the number of samples taken can introduce variable errors into the calculation of the utility function.

Figure 7 Score inflation

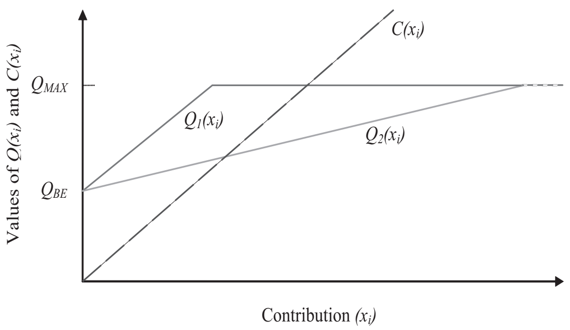

4.3 *INCA and incentives in distributed storage*

An architecture for fair sharing of resources is presented by Ngan et al. (2003). It nominally addresses storage resources, but can be extended to bandwidth as well. An incentive system ensures that participating entities can consume only as much resources as they themselves provide in return to the network. A regime of audits is instituted in order to prevent entities from gaining more than their share of resources.

Motivation: participating entities can take advantage of resources on the network, provided that they make available an equal amount of resources available to other entities. All well-behaved entities will receive the same service and in contrast to other systems where a malicious (or free-riding) entity will have its quality of service lowered or curtailed, a non-cooperating entity will be expelled from the network. The motivation thus comes from the fear of expulsion.

Differentiation: each entity publishes a *usage file*, which contains information about its offered service and currently held contracts with other entities. Consumers simply need to find a producer that offers an amount of storage equal to or greater than their storage requirements: producers thus differentiate themselves by self-description. Producers on the other hand need to perform a simple calculation dependent on consumer's request and usage file. All producers perform the same calculation in order to reject consumers that are over their allowed storage quota: consumer differentiation is thus a group-driven calculation.

Peer selection: peer selection is static for both consumers and producers.

Utility calculation: utility calculation is not explicitly discussed by the authors, but is straight-forward. Consumers are only interested in being able to use network's resources whenever they need to and that is a function of how much resources they offer in return. Producers on the other hand will only provide resources to well-behaved consumers and as such are concerned with social welfare.

Metric scope: as mentioned previously, entities use information contained in the usage file, which is global in scope.

4.3.1 *Analysis*

This section highlights some further research areas for the preceding publication based on INCA.

- Assuming that identities are not easy to obtain, a malicious peer's identity needs to be propagated throughout the whole network in order to ensure that it cannot enter into a contract with other peers. The main issue however is to ensure that a malicious peer cannot rapidly return with another identity.

- Assume that during a regular audit (where a consumer verifies that its content is truly stored by a given producer), the consumer discovers that the producer is malicious. The producer however passes audits by other consumers. Should the producer be punished (expelled from the network)?

- Following on the preceding point, what happens if a consumer claims – maliciously – that a producer is not storing files it claims to store?

- In order to motivate entities to be truthful, regular and random audits need to be performed. Whereas it is understandable that a consumer would take an interest that its files are properly stored – even at a cost to the consumer, performing random audits is less so. In fact, an entity could free-ride and let other entities perform random audits. The benefit of performing a random audit accrues to all participants, but the cost is borne only by the auditing parties.

- A consumer, without being malicious, could still obtain more than its fair share of network resources by following this simple scenario: before requesting resources, increase its own offered resources. After allocating its files to producers for storage, decrease its own offered resources. Auditing, both regular and random, would not be able to detect this anomaly.

4.4 INCA and incentives in content distribution

In order to improve the performance of data-sharing networks, Condie et al. (2004) present a protocol for creation of self-organising topologies based on peers' assessment of which direct connections are most likely to result in satisfactory file download. This results in formation of clusters of well-behaved entities, which have a broader view of the network than non-cooperative entities.

Motivation: it can be argued that motivation in this case is a combination of fear and greed (carrot and stick approach). Knowledge of entities' actions (choice of cooperation or non-cooperation) is accumulated throughout the network. After a sufficiently long time, each entity can be judged on the difference between the sum of cooperation instances and the sum of non-cooperation instances.

Differentiation: entities select *good* peers to connect to based on history of past interactions. Differentiation therefore is based on the result of a calculation by an entity.

Peer selection: peer selection is dynamic. Entities constantly update their connections with peers that maximise their utility function.

Utility calculation: all entities are interested in increasing only their own utility. The network is defined as an undirected graph with $G = (P, E)$ where P is the set of entities and E is the set of connections such that $c(i, j) = 1$ when there is a connection, $c(i, j) = 0$ when there is no connection and $c(i, i) = 0$ for $i, j \in P$. The objective function is defined as $Q_i = \sum_{j=1}^{v} c(i, j)s_{ij}$ subject to $\sum_{j=0} c(i, j) \leq \tau \ \forall\, i$. The number of satisfactory interactions s_{ij} is defined as $s_{ij} = \text{satisfactory}(i, j) - \text{unsatisfactory}(i, j)$.

Metric scope: metric is a combination of local trust, connection trust and void downloads. Entities assess other entities' cooperation from their own experience, but do not share this information with other entities; scope of the metric is thus local.

4.4.1 Analysis

This section highlights some further research areas for the preceding publication based on INCA.

- In order to establish connection with well behaved peers, an entity must first learn about the entities it interacts with. This is a function of the size of the network, cycle time – the time it takes to make a decision on the success of an interaction, for example, downloading a movie is slower than downloading a song, availability of entities – time they spent connected to the network and the number of connections each peer can have. All of these have an effect on convergence to a state that is stable and/or acceptable to a peer.
- Assuming a stable state is found, entities that join the network will have no history of previous interaction and thus it will be difficult for them to establish links with well behaved peers.

5 Conclusion

This paper presents INCA, 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. Functional and effective incentive mechanism will have all of those elements, except possibly the peer selection element. Each element presents some design choices that will define characteristics of an incentive mechanism.

This reference framework can be used in assessing the impact on the QoE of applications in the context of collaborative environments built on top of P2P architectures, whenever incentive mechanisms are present.

References

Abdul-Rahman, A. and Hailes, S. (2000) 'Supporting trust in virtual communities', *Hawaii International Conference on System Sciences*.

Adar, E. and Huberman, B. (2000) 'Free riding on gnutella', *First Monday*.

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

Antoniadis, P., Courcoubetis, C. and Mason, R. (2004) 'Comparing economic incentives in peer-to-peer networks', Vol. 46, pp.133–146.

Balazinska, M., Balakrishnan, H. and Stonebraker, M. (2004) 'Contract-based load management in federated distributed systems', *Symposium on Networked Systems Design and Implementation*, San Francisco, USA.

Bamberg, G. and Spremann, K. (1989) 'Agency Theory, *Information and Incentives*, Springer.

Bernstein, D., Feng, Z., Levine, B. and Zilberstein, S. (2003) 'Adaptive peer selection', *International Workshop on Peer-to-Peer Systems*, Berkeley, USA.

Bhagwan, R., Savage, S. and Voelker, G.M. (2003) 'Understanding availability', *International Workshop on Peer-to-Peer Systems*, Berkeley, USA.

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

Cohen, B. (2003) 'Incentives build robustness in bittorrent', *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, USA.

Condie, T., Kamvar, S.D. and Garcia-Molina, H. (2004) 'Adaptive peer-to-peer topologies', *IEEE International Conference on Peer-to-Peer Computing*, Switzerland.

Dewan, P. (2004) 'Peer-to-peer reputations', *Twelveth International Workshop on Parallel and Distributed Real-Time Systems*, Santa Fe, USA.

Dewan, P. and Dasgupta, P. (2004) 'Securing reputation data in peer-to-peer networks', *International Conference on Parallel and Distributed Computing and Systems*, MIT Cambridge, USA.

Douceur, J.R. (2002) 'The sybil attack', *International Workshop on Peer-to-Peer Systems*, Cambridge, USA.

Feigenbaum, J. and Shenker, S. (2002) 'Distributed algorithmic mechanism design: Recent results and future directions', *Sixth International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, New York: ACM Press, pp.1–13.

Feldman, M., Lai, K., Stoica, I. and Chuang, J. (2004) 'Robust incentive techniques for peer-to-peer networks', *Fifth ACM conference on Electronic Commerce*, New York, NY, USA, pp.102–111.

Friedman, E. and Resnick, P. (2001) 'The social cost of cheap pseudonyms', Vol. 10, pp.173–199.

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

Habib, A. and Chuang, J. (2004) 'Incentive mechanism for peer-to-peer media streaming', *Twelfth IEEE International Workshop on Quality of Service*, Montreal, Canada.

Hausheer, D., Liebau, N.C., Mauthe, A., Steinmetz, R. and Stiller, B. (2003) 'Token-based accounting and distributed pricing to introduce market mechanisms in a p2p file sharing scenario', *IEEE International Conference on Peer-to-Peer Computing*, Linköping, Sweden.

Hummel, T., Strømme, Ø. and Salle, R.L. (2003) 'Earning a living among peers – the quest for viable p2p revenue models', *Hawaii International Conference on System Sciences*.

Kamvar, S.D., Schlosser, M.T. and Garcia-Molina, H. (2003) 'The eigentrust algorithm for reputation management in p2p networks', *International World Wide Web Conference*, Budapest, Hungary.

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

Lai, K., Feldman, M., Stoica, I. and Chuang, J. (2003) 'Incentives for cooperation in peer-to-peer networks', *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, USA.

Leibowitz, N., Ripeanu, M. and Wierzbicki, A. (2003) 'Deconstructing the kazaa network', *Third IEEE Workshop on Internet Applications*, San Jose, USA.

Ma, R.T.B., Lee, S.C.M., Lui, J.C.S. and Yau, D.K.Y. (2004) 'A game theoretic approach to provide incentive and service differentiation in p2p networks', *SIGMETRICS*, New York, USA, 2004.

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

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

Ngan, T-W., Wallach, D. and Druschel, P. (2003) 'Enforcing fair sharing of peer-to-peer resources', *International Workshop on Peer-to-Peer Systems*, Cambridge, USA.

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

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

Ross, D. (2004) 'Game theory', E.N. Zalta (Ed) *The Stanford Encyclopedia of Philosophy*, Winter.

Rowe, L.A. and Jain, R. (2004) 'Acm sigmm retreat report on future directions in multimedia research'.

Senior, M. and Deters, R. (2002) 'Market structures in peer computation sharing', *IEEE International Conference on Peer-to-Peer Computing*, Linköping, Sweden.

Serjantov, A. and Lewis, A. (2003) 'Puzzles in p2p systems', *Eighth CaberNet Radicals Workshop*, Corsica.

Shneidman, J. and Parkes, D. (2003) 'Rationality and self-interest in peer to peer networks', *International Workshop on Peer-to-Peer Systems*, Berkeley, USA.

Singh, A. and Liu, L. (2003) 'Trustme: Anonymous management of trust relationships in decentralized p2p systems', *IEEE International Conference on Peer-to-Peer Computing*, Linköping, Sweden.

Strulo, B., Smith, A. and Farr, J. (2003) 'An architecture for peer-to-peer economies', *IEEE International Conference on Peer-to-Peer Computing*, Linköping, Sweden.

Terveen, L. and Hill, W. (2001) 'Beyond recommender systems: Helping people help each other', J. Carroll (Ed). *HCI in The New Millennium*, Addison-Wesley.

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

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

Wang, Y. and Vassileva, J. (2003) 'Bayesian network trust model in peer-to-peer networks', *International Workshop on Agents and Peer-to-Peer Computing*, Melbourne, Australia.