A Survey on Multicast Cost Sharing

Jonatan Schroeder
Project Report - CPSC532L
University of British Columbia
jonatan@cs.ubc.ca

Abstract

In multi-point Internet communication such as multicast, sometimes there is the need to share the cost of communication between the several agents involved in the connection. This work presents some of the most commonly used incentive-compatible, individual-rational mechanisms that select the participating users and their cost share in a multicast transmission, focusing on efficiency and/or budget balance. The marginal cost and the Shapley value are presented, as well as algorithms to compute the mechanism in a distributed system. Related research is briefly presented, focusing on alternative mechanisms and specific applications of cost sharing.

1 Introduction

With the popularization of the Internet, some important emerging applications, such as multimedia broadcasting, desktop conferencing and collaborative computing, require simultaneous communication between groups of computers, what is called multi-point communication. One or many hosts may be required to send the same message through the network to several hosts, leading to unnecessary redundancy on the network, and increasing the bandwidth usage in the network. To address this issue, new strategies and approaches were developed, such as *Multicast.* In a multicast communication, a message is sent not to a single host or to an entire network, but to a group of hosts, with less redundancy and, thus, less cost of communication.

In multicast transmissions, several hosts, or agents, are interested in receiving the transmission of a flow of data, with each agent valueing the participation differently. Because of the heterogeneity of the network, the participation of an agent in a multicast connection is not only dependent on the individual valuation and cost of transmission associated to this agent in particular, but also to other agents' participation in the connection, which may lead to a share of the associated cost for transmitting the data.

In this project we present some mechanisms that were designed specifically to address this problem. These mechanisms will identify, given the agents' valuations and the cost involved in sending the data to each subset of receivers, the subset of agents that will participate in the multicast communication, and the amount each of these agents will have to pay. This selection has to take into consideration several properties, described in section 2.

This project is organized as follows. Section 2 describes the general models and properties assumed by most mechanisms described in this report. Section 3 addresses two common mechanisms used to address cost sharing. Section 4 discusses some complexity issues, mainly in terms of network usage and computational complexity, and describes distributed algorithms to implement cost sharing mechanisms. Related work on multicast cost sharing mechanism is shown in section 5. Finally, section 6 presents the conclusions.

2 Assumptions and Properties

In this report we will use a model that is based mostly on the model presented in [4], with some changes to focus on the terminology used in [16] and to simplify some aspects that are explained more superficially in this report.

We assume a user (agent) population P, where each node resides in a location in the network. One location may have more than one user. We assume a single source, and that there is a single path $T(i)$ that connects the node in which $i \in P$ resides to the source. We also assume that, for a subset of receivers $R \subseteq P$, the union of all paths of users that belong to R is given by $T(R) = \bigcup_{i \in R} T(i)$[1]. Each link l in the network has a cost $c(l) \geq 0$ associated, known by the nodes on each end of the link. The cost to reach a receiver i is given by $c(T(i)) = \sum_{l \in T(i)} c(l)$. Similarly, the cost to reach a set of receivers R is given by $c(T(R)) = \sum_{l \in T(R)} c(l)$.

Given a multicast transmission, consider that each user i has a valuation $u_i \geq 0$ to participate in the transmission. The user's valuation for not participating in the transmission is 0. Each user will also be assigned a

[1] We assume, for simplicity, that $T(R)$ is a tree. This is a valid assumption if we consider that the path $T(i)$ for any user i is a minimum cost path, and is also assumed by many multicast routing proposals.

payment p_i. Considering a value σ_i that is 1 if the user participates in the transmission, and 0 if the user does not participate, we assume that the user's utility (or welfare, as defined in [4]) is defined by $w_i = \sigma_i u_i - p_i$. Considering a set R of receivers, we define the overall welfare, or net worth, of the set R as $NW(R) = \sum_{i \in R} u_i - c(T(R))$. Note that the overall welfare is not dependent of the payments.

A multicast cost sharing mechanism is defined as a mechanism that, given the model defined above, selects the values for the functions $\sigma_i(u)$ and $x_i(u)$ for each agent i, i.e. it selects the receivers that will participate in the transmission and the payment each receiver is assigned. Some properties are required in these mechanism:

- *Truthfulness*: also known as incentive-compatibility, defines that each agent will have the maximum utility if he declares his own valuation truthfully;
- *Individual Rationality*: for each agent i, $w_i(u) \geq 0$, i.e. only agents that participate in the transmission have payments, and this payments are limited to the agent's valuation;
- *No Positive Transfer*: limits the payments to non-negative values ($p_i(u) \geq 0$), so we only consider situations in which users pay for the right to participate;
- *Consumer Sovereignty*: guarantees the participation of a user if his valuation is high enough;
- *Symmetry*: if two users i and j are at the same node or at different nodes separated by a zero-cost path, and $u_i = u_j$, then $x_i = x_j$ [3].

The term strategyproofness is sometimes used in the literature to refer to dominant strategy truthfulness [3, 4, 10]. Group strategyproofness is used to denote mechanisms that maximize the agents valuation for a truthful declaration even if a set of users collude.

Beside these properties, some additional properties are not required, but desired:

- *Budget-balance*: the revenue raised from the receivers covers the cost of the transmission exactly ($\sum_i x_i(u) = c(T(R(u)))$);
- *Efficiency*: the mechanism maximizes the overall welfare $NW(R(u))$.

According to a theorem proposed by Green-Laffont in [8], these properties are not achievable simultaneously by any dominant-strategy truthful mechanism. Feigenbaum goes beyond, proving in [3] that approximate efficiency cannot be achieved simultaneously with approximate budget-balance in strategyproof mechanisms. So, the mechanisms described in this report will address one of these properties, but not both. Specifically, the marginal cost mechanism, described in section 3.1, is efficient, but not budget-balanced, while the Shapley value mechanism, described in section 3.2, is budget-balanced, but not efficient.

3 Basic Mechanisms

As stated in the previous section, a multicast cost sharing mechanism is expected to be an incentive-compatible individual-rational mechanism with no negative payments. It is also expected that budget balance and/or efficiency be sought, although it is not possible to achieve both simultaneously. In this section we present two mechanisms that support these properties: the marginal cost mechanism and the Shapley value mechanism. These mechanisms are the basis for most of the recent research on the area. Both are presented in [4].

3.1 Marginal Cost

Green and Laffont state in [7] that any efficient truthful mechanism is based on the Groves mechanism [9]. The most well-known mechanism that satisfies this property is known as VCG (Vickrey-Clarke-Groves) [1, 16]. The *marginal cost mechanism* is an application of the VCG mechanism to the multicast cost sharing problem. This mechanism is both dominant-strategy truthful and efficient, but is not budget-balanced. In fact, this mechanism never results in a positive revenue, and in some cases even leads to all agents paying nothing [12]. This mechanism is defined as follows.

$$
\begin{aligned}
R'(u) &= \operatorname*{argmax}_{R \subseteq P} NW(R) \\
R^*(u) &\in \operatorname*{argmax}_{R \in R'} |R| \\
\sigma_i(u) &= 1 \text{ if } i \in R^*(u),\ 0 \text{ otherwise} \\
x_i(u) &= u_i \sigma_i - \big(NW(R^*(u)) - NW(R^*(u|u_i = 0))\big)
\end{aligned}
$$

This mechanism selects as participants the users that correspond to the largest efficient subset in the network. The payment of each user represents the marginal contribution to the overall welfare that the user produces for having nonzero utility for the transmission.

3.2 Shapley Value

While the marginal cost mechanism provides an incentive-compatible efficient mechanism for multicast cost sharing, it lacks budget balance. For situations where this property is required over efficiency, this mechanism is no longer valid. Another mechanism, called the *Shapley value mechanism*, provides a truthful budget-balanced mechanism for multicast cost sharing [12, 15]. This mechanism is based on coalition theory, and thus imposes not only truthfulness in the individual level, but also that groups of users cannot increase their joint welfare by declaring incorrect utilities.

A budget-balanced mechanism is defined by a payment function $f : 2^P \mapsto \mathbb{R}^{|P|}$, $f_i(R) \geq 0$ with the property that $\sum_i f_i(R) = c(T(R))$ (the sum of the values for the function applied to a subset of users corresponds to the cost to reach these users). Based on this function, the following algorithm is used to find the values for $x_i(u)$ and $\sigma_i(u)$:

1. $R' \leftarrow P$;
2. $R' \leftarrow R$;
3. $x_i \leftarrow f_i(R')$ for each i;
4. $\sigma_i \leftarrow 1$ if $u_i \geq x_i$, 0 otherwise;
5. $R \leftarrow \{i|\sigma_i = 1\}$;
6. Repeat steps 2 to 5 until $R' = R$.

This algorithm starts selecting all users for the multicast transmission ("grand coalition"). If any of the users has a negative utility (the cost of sending the transmission to the user is higher than his valuation for the transmission), this user is removed from the coalition, and the Shapley value is recomputed. This is repeated until the subset of participants converges.

Jain proves in [10] that this algorithm provides a budget-balanced, non-positive-transfer, individual-rational, consumer-sovereign, group strategyproof mechanism for any cross-monotonic function f, i.e. if $R' \subset R$, then $f_i(R') > f_i(R)$.

The Shapley value mechanism, specifically, defines the function f using the general definition of the Shapley value [15, 17], that is

$$f_i(R) = \sum_{S \subseteq R \setminus i} \frac{|S|!(|R| - |S| - 1)!}{|R|!} \left(c(T(S \cup i)) - c(T(S)) \right)$$

In this formula, the cost of a link l is shared equally by all receivers who use the link to receive the stream. Although this mechanism is not efficient, the overall welfare produced by this mechanism is strictly higher than any other budget-balanced multicast cost sharing mechanism, as proven in [12]. So, the Shapley value mechanism is the closest to an efficient mechanism that one can obtain if budget-balance is required.

4 Distributed Algorithms

Beside the properties presented in section 2, another aspect that is expected in a multicast cost sharing mechanism is *tractability*, that states that the complexity of the mechanism should be reasonable. We will consider as complexity not only the computational complexity required to run the mechanism, but also what we will call network complexity, that corresponds to the number of messages sent in total and in each link, and to the size of each of these messages [4].

Both mechanisms presented in section 3 assume that the mechanism designer uses as inputs the full set of valuation declarations and all the link costs. The problem is that this information is available only locally at the network locations, and thus has to be sent to the root. If all valuations and costs are sent individually, the complexity either in terms of the number of messages (if values are sent separately) or in terms of the size of a message (if values are concatenated in a message) is linear in terms of the number of users. The messages returning the results of σ and x are also worth considering. Thus, to reduce the number of messages, the mechanisms have to be adapted.

Feigenbaum presents a version of the marginal cost algorithm in [4] that uses only two messages in each link of $T(P)$ to compute and distribute all valuations and costs, one message in each direction. In fact, each message sent from a node α in the network to its corresponding parent (the next node in the path to the source) contains only one value, corresponding to the sum of the utilities of the agents in node α and in its child nodes (nodes that use α to reach the source), minus the sum of the costs of the link connecting α to its parent and of the children links. The message is sent with value 0 if the cost is higher than the sum of utilities, and in this case the node does not expect to be part of the multicast transmission. The root sends back to each child a message corresponding to the welfare contribution of the node, and this value is used by each node to compute the user's participation and payment in the multicast transmission. The marginal cost functionality is maintained, while the number and size of messages in each link is constant.

Feigenbaum also presents in [4] a distributed version of the Shapley value mechanism. This distributed version uses, in the worst case, $O(np)$ messages in total, where n is the number of users in the population ($|P|$) and p is the number of links in the tree that reaches P ($|T(P)|$). Each link will send, in the worst case, $O(n)$ messages, equivalent to two messages in each step of the algorithm (one in each direction). The algorithm works as follows:

1. Each node α sends, using a bottom-up strategy, the number of users in the subtree rooted at α (p_α);
2. The root initiates a top-down traversal, sending a message with $md = 0$ to all its children;
3. Each node, after receiving message md, computes $md' = \frac{c(l)}{p_\alpha} + md$ (where $c(l)$ is the cost of the link that connects α to its parent) and sends it to its children;
4. The value md' computed in each node is attributed as payment to each of its users (x_i);
5. If a user's valuation is lower than the payment ($u_i < x_i$), then this user decides to resign the participation ($\sigma_i = 0$);
6. Repeat the algorithm, now only with users that have not resigned, until the number of users received by the root converges.

This algorithm provides a tool for mechanism designers that, while keeping all the properties of a regular Shapley value mechanism, like truthfulness, budget-balance and individual rationality, distributes the computation to the nodes, and reduces the amount of both processing and information the root has to deal with.

Both these distributed algorithms are based on a model called the *Temper Proof Model* (TPM), which assumes that the nodes may lie about their valuations, but the algorithm being run is not changed. Mitchell and Teague propose a new algorithm for marginal cost in [11] on a different model called the *Autonomous Nodes Model* (ANM). In this model, agents are able to deviate not only their declarations, but also the distributed version of the algorithm that runs in their nodes. Mitchell and Teague propose a mechanism that prevents cheating through the introduction of asymmetric key criptographic primitives, to verify that the user used the correct version of the algorithm. A new version of the Shapley value mechanism algorithm is also presented in [5].

5 Related Work

Dutta and Ray describe in [2] a mechanism used for coalition theory (*n-person transferable-utility cooperative game*). The mechanism is called Egalitarian Mechanism, and seeks to distribute the cost equally among all the receivers. Mutuswami shows in [13] that assuming all users draw their utility values from the same probability distribution, this method maximizes the expected size of the set served.

Jain and Vazirani introduce a budget-balanced mechanism in [10] similar to the Shapley value mechanism, but with a function $f_i(R)$ that is computed based on the available paths from i to other nodes including the source. This mechanism is budget-balanced, but results in a smaller overall welfare than the Shapley value mechanism, as shown in [4].

Penna and Ventre extend the marginal cost mechanism in [14], by providing a function of cost that is not just the sum of the link costs. The focus on this work is multicast transmission on wireless networks, and the cost is defined as a function on the power usage in the communication and the Euclidean distance of the nodes. They prove that some of the properties hold for a VCG-like mechanism, but also prove the impossibility of some other properties.

Garg and Grosu studied the performance of both the marginal cost and the Shapley value mechanisms in [6], and compared them through an experimental evaluation on a real network. They compare both mechanisms in both Tamper Proof and Autonomous Nodes models, described in section 4. The results include a comparison among all four algorithm combinations, comparing the running time in each node, the total payment received by the mechanism, the number of users receiving the multicast transmission and the effect of changing the number of users per node. They also study the effect of cheating in the received payment for the Shapley value mechanism in the Tamper Proof model.

6 Conclusion

In this report we summarize some of the available research on multicast cost sharing. We discussed two of the most well-known mechanisms: the marginal cost mechanism, based on VCG, that focuses on maximizing the overall welfare, and the Shapley value mechanism, that focuses on budget balance and collusion truthfulness. We also showed some of the other research focuses in this area, such as different mechanisms and the application of multicast cost sharing in specific areas.

References

[1] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[2] B. Dutta and D. Ray. A concept of egalitarianism under participation constraints. *Econometrica*, 57:615–635, 1989.

[3] J. Feigenbaum, A. Krishnamurthy, R. Sami, and S. Shenker. Hardness results for multicast cost sharing. Technical Report YALEU/DCS/TR1232, Yale University, June 2002. `http://ftp.cs.yale.edu/pub/TR/tr1232.ps`.

[4] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63(1):21–41, 2001.

[5] N. Garg and D. Grosu. Faithful distributed Shapley mechanisms for sharing the cost of multicast transmissions. Technical report, Dept. of Computer Science, Wayne State University, September 2006.

[6] N. Garg and D. Grosu. Performance evaluation of multicast cost sharing mechanisms. In *Proceedings of the 21st International Conference on Advanced Networking and Applications (AINA 2007)*, pages 901–908, Washington, DC, 2007. IEEE Computer Society.

[7] J. Green and J. Laffont. Characterization of satisfatory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45(2):427–438, 1977.

[8] J. Green and J. J. Laffont. Partial equilibrium approach to the free rider problem. *Journal of Public Economics*, 6:375–394, 1976.

[9] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[10] K. Jain and V. Vazirani. Applications of approximation algorithms to cooperative games. In *Proceedings of the 33*$^{\mathrm{rd}}$ *Symposium on the Theory of Computing*, pages 364–372, New York, 2001. ACM Press.

[11] J. Mitchell and V. Teague. Autonomous nodes and distributed mechanisms. In *Software Security - Theory and Systems. Mext-NSF-JSPS International Symposium (ISSS)*, Lecture Notes in Computer Science, pages 58–83. Springler Berlin / Heidelberg, 2003.

[12] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs:budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.

[13] S. Mutuswami. Strategy proof mechanisms for cost sharing. In *Economics Discussion Papers*, volume 520. University of Essex, Department of Economics, 2000.

[14] P. Penna and C. Ventre. *Sharing the Cost of Multicast Transmissions in Wireless Networks*, volume 3104/2004 of *Lecture Notes in Computer Science*, chapter 23, pages 255–266. Springer Berlin, September 2004.

[15] L. S. Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2:31–40, 1953.

[16] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.* Cambridge University Press, 2008. to be published.

[17] M. Shubik. Incentives, decentralized controls, the assignment of joint costs and internal pricing. *Management Science*, 8:325–343, 1962.