ISBN 978-83-60810-51-4

Multiagent scheme for voice conference moderation

Adam Połomski
Warsaw Univerity of Technology,
ul. Nowowiejska 15/19, Warsaw, Poland
Email: A.K.Polomski@elka.pw.edu.pl

***Index Terms*—Multiagent systems, Auctions, Social welfare, Moderation**

***Abstract*—Conferencing systems utilizing text messages as the communication medium have been around for many years. Since in highly populated social platforms moderation is more of a necessity than luxury, many different mechanisms exist and function being more or less effective. As the Internet bandwidth is becoming more and more accessible, voice over IP communication is gaining on popularity. Multiuser voice conferencing platforms raise the need for different type of moderation mechanisms. Determining a fair moderation scheme which would in the same time maximize the overall discussion quality for each participant is not a trivial task. We introduce a multiagent model for voice conference moderation which utilizes Vickrey auctions as a resource allocation procedure. By applying a concept of communication channel as a resource with an equally shared ownership, we enforce that rules of a fair discussion are fulfilled.**

I. INTRODUCTION

THE WORLD is driven by information [1] and the ability to exchange knowledge is crucial. In some cases it is necessary to collaborate with multiple people at the same time. Instant sharing and exchanging data with a whole group might be indispensable. On one hand, there is an added value related with group conferences, while on the other, it does require extra effort to manage.

The most general definition of a conference would be: ”a meeting for consultation, exchange of information, or discussion”. However, this requires further clarification. There are various types of conferences depending on conversation style, type of participants and, what is most important, main conference target. An ad-hoc discussion among group of unanimous people will most likely have different progress than a lecture on a specific subject given to a class of students. Still, in both examples there is a flow of information. In more or less chaotic ways, participants share their knowledge and opinions with each other. Managing such a discussion might not be a trivial task, especially if there is a significant group of participants with a high need of expressing themselves and aggressively fighting to do so. Such situation might be quite frequently observed in television debates. If two or more people are determined to dominate the discussion, it becomes very chaotic. Flow of information is then minimal, what results in drastically lowering the value of such conference. A significant number of conversation participants is also an important factor. Exchanging information is an action which involves a group of people, therefore requires coordination. Coordination scheme gains on importance as the group grows in members. Another important factor is disruptive behavior. In conversational spaces where entry barriers are set low, phenomena such as trolling, flooding, flaming or spamming emerge frequently. This calls for the need for some sort of moderation mechanism.

There is a number of papers treating about moderation schemes implemented in popular social platforms like Usenet [2] or Slashdot [3]. Both those models are decentralized, what is enforced by the size of communities. There is a lot of ongoing research on multiagent resource allocation schemes. Two of them [4], [5] in particular treat about multiagent network bandwidth allocation. They introduce the idea of agents bargaining and trading the allocated good among each other.

In the next section we describe the background of the conference moderation problem. First we characterize the environment in which a designed model needs to function. We also clarify the basic concept of moderation and present other possible means of moderating conferences. Section 2 discusses the concept of fairness in generalized multiagent environment. It is later described in the context of conference moderation problem. Since discussion is a process which extends in time, preferences of participants might change dynamically and are highly correlated with the discussion flow up to that moment. Thus, any simple social welfare metric cannot be applied. In Section 3 we introduce a multiagent moderation model for multiuser voice conferencing platform. We conclude the paper with the testbed design. A verification methodology which involves applying the proposed automatic moderation model to a voice conferencing platform is presented.

II. BACKGROUND

The environment with which we are dealing is very specific. It is a huge online audio conference platform connecting people from all parts of the world. It is used by hundreds of thousand users every day. It is very likely that conference participants will not know each other. The discussion topics may vary and cannot be anyhow limited or managed. The only thing that can be assumed or enforced is that all participants in a single discussion share the same language. Also the number of concurrent ongoing conferences may be as high as tens of thousands. The model of a single conference is very simple though. All participants connect to a single device - media server, which is responsible for handling the voice stream. Strictly speaking it broadcasts the voice stream transmitted by a single participant to all others. It also has a steering protocol which allows controlling it to some extent. What is most important, it allows specifying which participants are allowed to transmit voice signal in the given moment of time and which ones are only allowed to receive it.

A. Moderation

For this sort of audio conversation platform to function successfully, a moderation mechanism is required. User experience would suffer otherwise. It is hard to strictly define what would be seen as ”good” or ”bad” conversation by the participants. There are three general rules though, which for sure have huge influence on the overall outcome of such measurement:

- Each participant can express him/herself.
- Each participant will be heard.
- No participant will dominate the discussion.

Moderation is a mean to fulfill those requirements. If the ability to speak is distributed properly among participants over time, it should be possible to maximize welfare of the whole group. There are few standard types of moderation, which can be observed:

- No moderation - For example a group of friends talking at the cafeteria will not require any moderation to get the most of their discussion. It is important that the group is relatively small. The fact that participants know each other well and have no point in dominating the discussion also helps.
- Discussion rules - For example a lecture at the university. Both, the lecturer and the students are aware of discussion rules up front and will respect them.
- Human moderator - A designated person is responsible for leading the discussion by granting/revoking voice to participants. It is also very important that moderator understands and follows the discussion, as it is crucial to pick the right people to speak in a given time slot.
- Queuing - All participants queue up and are the voice is granted in a ”round robin” scheme.

All of the above moderation models except queuing cannot be applied in this specific environment. That is mostly because of size of the system. Queuing is the only model that does not require manual management and puts no trust that users will apply to some rules without forcing them into it. A huge drawback of queuing in this case, is the maximum pessimistic waiting time. It is possible that a participant who needs the voice most will be forced to wait until everyone else uses the granted time slot. That is an area, where introducing a multi agent solution could bring better results.

B. Fairness

When it comes to resource allocation procedures the concept of fairness is a very important issue to solve. Is this specific resource distribution fair for all agents? What criteria are necessary to match in order to obtain the resource distribution, which might be seen as fair? Fairness is a normative issue [6], which means it involves judgments. That is why it is hard to claim a special expertise on this issue and it may be considered as a philosophical debate. Different agents may have different goals, therefore the mechanism which results in unequal resource allocation cannot be considered unfair without more detailed analysis. We can distinguish three main concepts of fairness, which focus on different core attributes:

- Equal opportunity - in this view it is the process that matters, not the result. This concept treats the whole division process as a race. As long as there are rules, which have to be obeyed by every participant in the same way, the division is fair. Of course, some participants might be faster, stronger, more agile and gain more than the weaker ones. This is still considered as fair, because of the same rules applied to everyone
- Egalitarian - this view focuses only on the result, not the process itself. This concept argues that more equality of income is always better then less, and the complete equality is the state which should be the aim.
- Need - need is the prioritized attribute here. To implement this view, we have to be able to determine/measure the need of each agent. This measurement is most practical in small populations, while it is hard to determine the need in large groups. One escape from this problem is to assume, that everyone’s need is the same, but this results with the point of view identical to egalitarianism.

C. Multiagent resource allocation

Multiagent resource allocation is a process of distributing a number of items (resources) among a number of agents [7]. This brief definition, however, does not fully describe the problem.

Resources might differ in characteristics. The whole range of resource types is substiantial and each might require different allocation technique. For instance, we can distinguish divisible goods (like network bandwidth) and indivisible. Also it may, or may not be allowed to share an indivisible resource among a number of agents.

Agents may have preferences over different allocation outcomes. Not only may they have preferences over resource bundles they receive, but also over bundles received by others. Preferences can be represented in a various ways, like utility function or binary relation on alternatives. Moreover, agents may or may not be truthful while reporting their preferences.

Allocation can be performed with the use of various allocation procedures, which can be either centralized or distributed. Typical examples of centralized procedures are auctions or voting mechanisms, with a central entity empowered to decide on the final allocation. In distributed solutions agents try to come to a common agreement through a sequence of local negotiation steps. In both cases, the objective is to find an allocation which is feasible or optimal. What stands behind the concept of optimal depends on the specific multiagent resource allocation scenario.

D. Social welfare

A typical objective for multiagent resource allocation problem is to find the resource distribution, which maximizes value of a specified metric [7]. This metric somehow depends on the preferences of individual agents. The way that preferences of agents are aggregated defines the most basic rules, on which the society of agents will collaborate while allocating the resources. There are plenty of different social welfare notions and many of them find application in multiagent resource allocation problem. The choice should be dictated by the general aim of the whole society. For most e-commerce applications the aim is to maximize the average profit generated by negotiating agents. Utilitarian social welfare should provide the pretty good assessing system in this case. On the other hand, the aim might be to fairly share the resource, which has been earlier cofounded (in unequal proportions) by all agents. Clearly, optimizing the agents' average utility should not be the goal here. Thus, utilitarian social welfare is not applicable in this case. Generally, before the negotiation starts, the agents should know under what rules, the society operates. Otherwise, it may not be prepared to agree to the distribution which in this society is seen as the most optimal.

If $>_i$ is a preference function over the set of outcomes, a desirable relation of global social preferences $>^*$ should have the following set of characteristics:

- $>^*$ exists for all $>_i$
- $>^*$ exists for every pair of outcomes
- $>^*$ is asymmetric and transitive over the set of outcomes
- $>^*$ is Pareto optimal
- $>^*$ should be independent of irrelevant alternatives
- No agent is a dictator

Unfortunately Arrows has proven that no such mechanism exists [8]. Therefore we are forced to pick from social welfare metrics which none of them is perfect, but may be relevant to use in the specific scenario we are facing.

III. MODEL PRESENTATION

Let $N = \{a_1, a_2 ... a_n\}$ be a finite set of n agents, where $n \geq 2$. Each of n agents is a conference participant. Strictly speaking a conference with 2 participants is not a group conference. Nevertheless, the model is also valid for face to face discussions. The conference lasts for a specified amount of time, which is not known from upfront. Conference is divided into k shorter periods $T = (t_1, t_2, t_3 ... t_k)$, each of an equal length t. Throughout the whole conference time all agents compete for obtaining full rights to the transmission channel, therefore the possible set of allocation outcomes Δ is equal to N. During each time period t_k only a single agent may speak, while all others listen. Let δ_j denote the allocation outcome at the period t_j. There is also a special null allocation δ^-, which represents a situation where no agents holds the transmition rights. Preferences of each participating agent i in time period j are represented by the utility function $u_{ij} : \Delta \rightarrow R$. Utility functions are not known upfront, as u_{ij} depends on the whole history of allocations up to this point in time $(\delta_1, \delta_2, \delta_3 ... \delta_{j-1})$ and what those allocations brought. In other words, in every moment of an ongoing discussion agents may alter their valuations based on what they have heard so far from different conference participants.

Determining if a given allocation is fair in such conditions is not very straight forward. Utilitarian social welfare metric is a common choice for various applications, as it maximizes the overall profit of the whole society.

$$\frac{1}{n} \sum_{a_i \in N, t_j \in T} u_{ij}(\delta_j) \tag{1}$$

Although it would indeed promote allocations which result in greater average happiness, it is not in line with the rules of fair discussion proposed in Section II-A. Let A be an agent and M a very big positive number. A group of $\lceil \frac{n}{2} \rceil$ agents all with valuation functions shaped that $u_j(A) = M, \forall j$ would deny allocation for any agent other then A.

Another metric which might be found useful in this particular scenario is Nash product. It promotes allocations which result in equal valuation distribution.

$$\prod_{a_i \in N, t_j \in T} u_{ij}(\delta_j) \tag{2}$$

Unfortunately, it is also not in compliance with the rules of fair discussion. In contrary to utilitarian social welfare, it can be dominated by a single agent A with a utility function such as

$$\forall j u_j(a) = \begin{cases} \epsilon & \text{if } a \neq A \\ M & \text{if } a = A \end{cases}$$

An ideal social welfare metric for transmission channel allocation should in the first place ensure that all rules of fair discussion proposed in Section II-A are fulfilled. Only after this is secured, higher overall profit should be promoted. We propose the conference metric given by equation 3, which mixes the concepts of utilitarian social welfare and Nash product together.

$$\prod_{a_i \in N} \frac{\sum_{t_j \in T} u_{ij}(\delta_j)}{n} \tag{3}$$

Average welfare over the full set of discussion periods is calculated for each participating agent. Our conference metric promotes allocation schemes, where those averages are distributed equally among agents, what stands for fair discussion. Allocations which give generally higher average welfare for all agents are also preferred.

The following example illustrates all aforementioned metrics. Let $N = \{A, B, C, D\}$. Conference consists of 3 allocation periods (t_1, t_2, t_3). Agents A, B, C have utility function as follows:

$$u_j(a) = \begin{cases} 1 & \text{if } a \neq A \\ 2 & \text{if } a = A \end{cases}, \forall j$$

TABLE I
SAMPLE WELFARE VALUES

Metric	(A, A, A)	(D, A, A)	(D, D, A)	(D, D, D)
Utilitarian (Equation 1)	4,575	4.3	4.025	3.75
Nash (Equation 2)	0.512	1.28	3.2	8
Conference (Equation 3)	0.8	3.4	3.24	2.0

Utility function for agent D:

$$u_j(a) = \begin{cases} 0.1 & \text{if } a \neq D \\ 2 & \text{if } a = D \end{cases}, \forall j$$

Table I presents sample welfare values for four different allocation schemes calculated with the use of each metric. Utilitarian social welfare promotes agent A. Nash product does the same for agent D. Out conference metric gives highest value to allocation (D, A, A), which seems to be the most fair of all. It grants a single discussion time period to agent D what ensures fairness. The remaining allocations are performed with respect of preferences of majority.

IV. BASIC MODEL

The proposed model is based on the assumption, that the whole conference is cofounded by each of the participating agents. Therefore each participant owns an equal share of rights to the transmitting channel. Of course, the channel is an indivisible resource and can only be utilized if fully owned. We assume though, that giving any value to δ^- should be interpreted as disruptive, as it leads to canceling any communication whatsoever. We do not want to promote such behavior, therefore $u_i(\delta^-) = 0, \forall i$.

No one can use just a part of the channel, yet agents can negotiate over the price and purchase or sell it to each other. For this purpose each agent a_i is associated with r_{ij}, which might be interpreted as agent's wallet for time period t_j. At the beginning of the discussion $r_{i1} = R, \forall i$, where R is a positive constant value to initialize all the wallets. The conference T is divided into k shorter periods, at the beginning of each agents may express their desire to obtain full rights to transmission channel. Resource allocation is then performed with the use of Vickrey auction mechanism [9]. The allocation pattern for each discussion period t_j is following:

1) Each conference participant a_i issues a bid with the valuation v_{ij}. The bid cannot be higher then the actual wealth of agent at that time, therefore

$$v_i j = \min(u_{ij}(a_i), r_{ij}) \tag{4}$$

2) The winning agent a_{kj} and the price to pay p_j is determined with the use of Vickrey auction.
3) Price to pay is deducted from the winner's wallet $r_{k,j+1} = r_{kj} - p_j$.
4) All agents which sell their transmission rights to a_{kj} are rewarded $r_{i,j+1} = r_{ij} + \frac{p_j}{n-1}$.
5) The whole transmission channel for time period t_j is allocated to $\delta_j = a_{kj}$

This procedure is then repeated for each discussion period. Conference ends when none of the participating parties are willing to compete for transmission channel rights - $v_{ij} = 0, \forall i$.

By leveraging Vickrey auction, we gain all characteristics of this auction mechanism. Performing allocation is quick and does not require a lot of overhead network traffic. This is a very important feature, as it is crucial to finish negotiations before the beginning of conference period affected by this allocation. Agents are also highly encouraged to bid their true valuations, as it is in line with the dominant auctioning strategy. Unfortunately Vickrey auction has a couple of drawbacks. It is vulnerable to bidder collusion agreements. A group of agents with the highest valuations may settle not to vote their true valuations in order to lower the resulting price. The model is also exposed to lying auctioneers. Agents may bid shill votes in order to inflate the price and increase the income.

Theorem 1: The proposed basic model ensures that all rules of fair discussion described in Section II-A are fulfilled.

Proof: Let t_l be a time period of an ongoing conversation and a_k an agent willing to express himself. The highest bid a_k can then submit is equal to the state of his wallet r_{kl}. The set of agents capable of overbidding a_k in an auction for time period t_l is N_l. The summaric state of all wallets for agents in N_l is $D_l = \sum_{a_i \in N_l} r_{il}$. If none of the N_l agents decides to bid higher than r_{kl} the transmission channel for period t_l is allocated to a_k. Otherwise a_k looses the auction in favor of some agent from N_l. The lowest possible allocation price is r_{kl}, since that is the bid from a_k. After the payments are transferred from the winning agent to all other participants $r_{kl+1} = r_{kl} + \frac{r_{kl}}{n-1}$ and $D_{l+1} = D_l - r_{kl} + r_{kl}\frac{|N_l|-1}{n-1}$. Since $a_k \notin N_l, |N_l| < |N| = n$, therefore $r_{kl+1} > r_{kl}$ and $D_{l+1} < D_l$. After a finite set of time periods g we will reach to a point where $r_{kl+g} > D_{l+g} \Rightarrow N_{l+g} = \emptyset$. Agent a_k is then guaranteed to win the allocation for t_{l+g}. This proves that regardless of participants' utility functions and state of wallets each agent will have a chance to express himself and no one will ever dominate the discussion. ■

Even though the discussion leveraging proposed model as a moderation mechanism would be fair, there is no guarantee whatsoever, that it would lead to optimal allocations in terms of welfare metric given by Equation 3. The basic limitation is introduced for agents who highly value listening to other participants. This model only allows to express the valuation for getting the ownership of transmission channel.

$$u_{ij}(a_k) = 0, \forall i \forall j \forall k \neq i \tag{5}$$

Even if that reflected the actual situation, agents would not be allowed to express their real valuations in bids due to restriction enforced by Equation 4. Whether or not this model is close to optimal allocations that still requires verification.

TABLE II
SAMPLE BIDS

	a_1	a_2	a_3	a_4	a_5
v_{ikj}	0,75	3	2,75	2	0
Δr_{ij}	0,75	-1,5	-1,25	-0,5	1,5

V. GENERALIZED MODEL

The model proposed in Section (IV) bases on the assumption that each participant does not have any positive valuation associated with an allocation other than to himself. This might not fully represent the real situation in all environments. It is very likely, that some discussion participants actually do value listening to some of the other people in the conference. Generalized model relaxes the previously made assumption characterized by Equation (5). Agents may assign positive utility values to multiple different allocation outcomes.

$$u_{ij}(a_k) \geq 0, \forall i \forall j \forall k$$

The core foundations are just alike in the basic model described in Section (IV). The whole conference channel is yet again cofunded by all participating agents each of whom holds an equal share. The concept of agent's wallet r_{ij} is also leveraged to enforce discussion fairness. Modifications are only related to bidding, auctioning and rewarding schemes. The allocation pattern for each discussion period t_j is now as follows:

1) Each conference participant a_i issues a number of bids with valuations v_{ilj}. One bid for every discussion member a_l for which $u_{ij}(a_l) > 0$. None of the bids can be higher thn the actual agent's wealth at that time, therefore
$$v_{ilj} = \min(u_{ij}(a_l), r_{ij}), \forall l$$
2) All valuations bid for the same allocation option are accumulated.
$$v_{lj} = \sum_{a_i \in N} v_{ilj}$$
3) The winning agent a_{kj} and the price to pay p_j is determined with the use of Vickrey auction operating on accumulated valuations.
4) Each agent is rewarded or charged based on his bid for the winning allocation $r_{ij+1} = r_{ij} + \Delta r_{ij}$, where
$$\Delta r_{ij} = \frac{p_j}{n} - v_{ikj} \tag{6}$$
5) The whole transmission channel for time period t_j is allocated to $\delta_j = a_{kj}$

A reward calculation mechanism for generalized model is illustrated in Figure (1). It shows sample reward calculation for bids as shown in Table (II). Mean valuation for the winning allocation is $\frac{p_j}{n} = 1, 5$.

Theorem 2: The proposed generalized model ensures that all rules of fair discussion described in Section II-A are fulfilled.

Fig. 1. Sample reward calculation

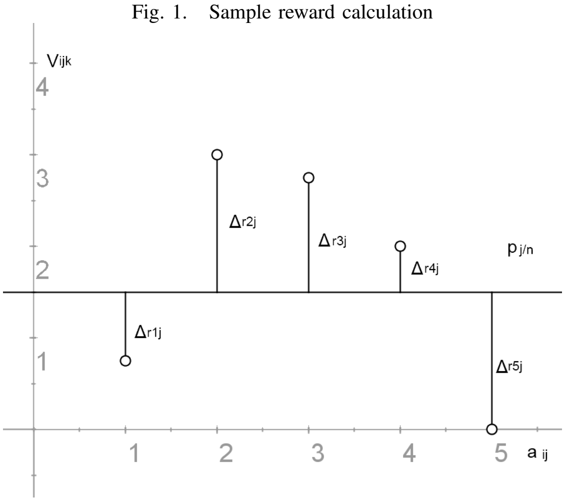

Proof: The concept of proof is very similar as for Theorem (1). If t_l is a time period of an ongoing conversation and a_k an agent willing to express himself, the most optimal set of bids a_k can then submit consists of a single bid, where $v_{kkl} = r_{kl}$. The summaric state of all wallets for agents in $N \backslash \{a_k\}$ is $D_l = \sum_{a_i \in N \backslash \{a_k\}} r_{il}$. If a_k loses the auction to another agent a_m, he is rewarded with Δr_{kj}, as in Equation (6). The reward is positive $\Delta r_{kj} > 0$, since $v_{kml} = 0$. Therefore $r_{k,l+1} > r_{kl}$, while $D_{l+1} < D_l$. This proves that after a finite set of time periods g we will reach to a point where $r_{kl+g} > D_{l+g}$, and a_k is then guaranteed to win the allocation for t_{l+g}. ■

By allowing participating agent to express all of their preferences instead of just a valuation of self-allocation $u_{ij}(a_i)$, the model should result with allocations that are more optimal in terms of welfare metric given by Equation (3). The extended model promotes participants whose statements are interesting to other conference members. This should bring some added value to an overall discussion quality. Furthermore, not all participants might want to take active part in the conversation. This also gives them the mean to express their preferences in who they want to listen to.

The major drawback introduced by this model extension is possibility of manipulations. Apart from collusion agreements and shill bids, aggregating valuations prior to performing Vickrey auction encourages agents to bid lower than their true valuation. There might be such a situation, when $\delta_j = a_k$ no matter if v_{ikj} is reported as in Equation (4) or lower. Therefore, lowering v_{ikj} would then raise Δr_{ij} which profitable for a_i.

VI. VERIFICATION

Both models described in Sections (IV, V) require verification. This will be performed by implementing a live multiuser voice conferencing application and exposing it to a real traffic. From a functional perspective the application is an online platform for setting up ad-hoc conferences with random participants. The randomization makes it less likely that discussion members know each other up front, therefore the importance of moderation is intensified. The verification focuses mainly on the following aspects:

- What is the allocation quality in terms of metric described by Equation (3)?
- Are models immune to disruptive behavior?
- What is the usability of the conferences running on top of both allocation models?

A. Metodology

Each of the ad-hoc created conferences has an equal chance of leveraging one of the three different moderation schemes:

- Round robin queue.
- Basic model.
- Generalized model.

Moderation scheme is randomly chosen at the beginning of the conference, each of the three schemes being equally probable. User interface is capable of collecting full preference model for every discussion member. Before the beginning of every conference period t_j all agents submit their full set of actual valuations over all possible allocation outcomes Δ. Even though not all of the collected data might be required to determine the allocation outcome δ_j, having full preference model for all participating agents over the duration of conference is required for later offline analysis.

VII. CONCLUSIONS

In this paper we proposed a social welfare metric to determine the quality of a conference taking fairness aspects into consideration. We also introduced two multiagent moderation models, which enforce fair discussion in terms of the aforementioned metric. Both presented models can potentially find application in case of ad-hoc discussions among multiple unanimous participants. However, the potential set of applications for both presented models is not limited to conference moderation only. Every scenario which involves fair allocation of an indivisible resource over multiple time periods might be a good area for the proposed solution. Finally, we also provided a brief description of verification metodology for multiagent moderation scheme.

Future work includes verifying quality of both models in a real life scenario, by deploying it to broad public and collecting results. Applying multicomodity auction model [10] to minimize the danger of potential manipulation is an area of further research. Another potential direction of study is converting the fair discussion problem into bicriteria optimization.

ACKNOWLEDGMENT

The research was partially supported by the Polish National Budget Funds 2010–2013 for science under the grant N N514 044438.

REFERENCES

[1] S. Cordes, “Broad horizons: The role of multimodal literacy in 21st century library instruction,” 2009.

[2] J. A. Konstan, B. N. Miller, D. Maltz, J. L. Herlocker, L. R. Gordon, J. Riedl, and H. Volume, “Grouplens: Applying collaborative filtering to usenet news,” *Communications of the ACM*, vol. 40, pp. 77–87, 1997.

[3] C. Lampe and P. Resnick, “Slash(dot) and burn: distributed moderation in a large online conversation space,” in *Proceedings of the SIGCHI conference on Human factors in computing systems*, ser. CHI ’04. New York, NY, USA: ACM, 2004, p. 543550. [Online]. Available: http://doi.acm.org/10.1145/985692.985761

[4] T. Hasselrot, “Fair bandwidth allocation in internet access gateways - using agent-based electronic markets,” SICS, Tech. Rep., 2003.

[5] A. Polomski, “Multiagent means of bandwidth allocation for telecommunication market,” in *Proceedings of the 4th KES international conference on Agent and multi-agent systems: technologies and applications, Part II*, ser. KES-AMSTA’10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 400–408.

[6] R. Schenk, “What is fair?” 1997. [Online]. Available: http://ingrimayne.com/econ/resouceProblems/Fairness.html

[7] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lematre, N. Maudet, J. Padget, S. Phelps, J. A. Rodrguez-aguilar, and P. Sousa, “Issues in multiagent resource allocation,” *Informatica*, vol. 30, p. 2006, 2006.

[8] J. M. Vidal, “Fundamentals of multiagent systems,” 2006. [Online]. Available: http://www.multiagent.com/fmas

[9] Y. Narahari, D. Garg, R. Narayanam, and H. Prakash, *Game Theoretic Problems in Network Economics and Mechanism Design Solutions*, 1st ed. Springer Publishing Company, Incorporated, 2009.

[10] K. Koltys, P. Palka, E. Toczylowski, and I. Zoltowska, “Paper multicommodity auction model for indivisible network resource allocation.”