Proceedings of the 2013 Federated Conference on
Computer Science and Information Systems pp. 1021–1027

Fair and truthful multiagent resource allocation for conference moderation

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

***Abstract*—Multiuser voice conferencing platforms are more and more popular. Internet bandwidth is becoming very accessible, what makes voice over IP used on an everyday basis. Being able to communicate with multiple people at the same time can be beneficial, but on the other hand increases the need of coordination mechanisms. Determining a moderation scheme which is fair and efficient is not a trivial problem to solve. We define conference moderation as a multiagent resource allocation problem and introduce a process based on Vickrey auctions to solve it. A concept of co-owned communication channel is what stands as a basis of our definition of fairness.**

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

I. INTRODUCTION

EXCHANGE of information is a crucial part of our existence. In some situations we have to elaborate with larger groups of people. It might be very valuable, as every member of such group brings some extra knowledge, potentially useful for the whole society. Yet, it comes at a certain price. It requires an additional effort to manage the flow of information as the number of participants grows. In a most general definition, a conference is ”a meeting for consultation, exchange of information, or discussion”. Yet, the conferences may vary a lot in its specifics. A discussion among a group of friends lives by a different rules than a company meeting. An environment in which a large group of random people struggles to exchange some information has the biggest potential of becoming extremely chaotic, thus bringing the flow of information to a minimal level. Open societies with low entry barriers often face problems of disruptive behavior such as flooding or spamming. How do we deal with that? How do we coordinate the flow of information in a way which is efficient and fair? That is a role for moderation mechanism.

Decentralized moderation schemes for large scale social platforms like Usenet or Slashdot is being discussed in a number of papers [1], [2]. Multiagent resource allocation and a concept of bargaining or trading agents is also a common topic of research, with distribution of network bandwidth being one of the possible applications [3].

This is an extended version of the paper in which we first introduced the presented allocation model [4]. We mostly focused on the verification of the model. However, we also provided a new, refreshed view of the problem definition and social welfare, which may be found in the next two sections. We gave some more detailed insight into features a fair and effective social welfare metric should possess. Section V contains a description of Vickrey auction based allocation method. We have also introduced a resale procedure to deal with a problem of choosing a proper allocation time. Next we move on to an overview of verification methodology in Section VI followed by simulation results in Section VII. A JADE based implementation of conferencing agents 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 is a basic rule that definitely needs to be fulfilled in order to achieve a fair discussion.

No participant will be able 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 a few standard, commonly known moderation mechanisms, 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

978-1-4673-4471-5/$25.00 ©2013, IEEE

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. Multiagent resource allocation

Multiagent resource allocation is a process of distributing a number of items (resources) among a number of agents [5]. This brief definition, however, does not fully describe the problem.

Resources might differ in characteristics. The whole range of resource types is substantial 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.

III. PROBLEM DEFINITION

We formulate conference moderation task as a resource allocation problem. The following definition is general and may find application whenever multiple agents compete over a non divisible resource across multiple time periods.

Let $N = \{a_1, a_2 \ldots a_n\}$ be a finite set of participating agents. Each agent takes part in $T \in \mathbb{N}^+$ consecutive resource allocation runs, each for a separate time period. A resource can only be allocated to one agent at a time, therefore the set of feasible allocations $\Delta = N$. Let δ_t denote the allocation outcome for t allocation run $\forall t \leq T$. There is also a special null allocation δ^- which represents a situation when no agent holds a resource. Each of the participating agents has its preference regarding every allocation represented by the utility function $u_{i,t} : \Delta \to R, \forall t \leq T, a_i \in N$. Let vector $x = (\delta_1, \delta_2 \ldots \delta_T)$ hold all allocations across all time periods. We call $S(x)$ a social welfare function. $S : \Delta^T \to R$. We will discuss it in details in section IV. For now it is only important that it determines an overall happiness of the whole society N for a given allocation vector x. Objective is to find x^*, which will $\max S(x)$.

Utility functions are not known upfront. While deciding upon an allocation δ_t, our knowledge consists of:

- current and past utility functions for all participating agents
- allocations that have been performed up to this point

Note that the shapes of current utility functions u_{it} are very likely to depend on the whole history of allocations up to this point in time $(\delta_1, \delta_2, \delta_3 \ldots \delta_{t-1})$ and what those allocations brought. In other words, every allocation decision can affect the way agents shape their utilities in the future. It might also have an impact on the total number of consecutive allocations T.

IV. SOCIAL WELFARE FUNCTION

Social welfare function S ranks every feasible allocation vector x. This ranking represents a welfare of the whole society N if an allocation x took place. Since S determines whether we choose one allocation vector over the other, it also defines what we consider as fair. It is not a straightforward task to rule what is fair in these specific conditions. Moreover, according to Arrows impossibility theorem [6] there exists no reasonably consistent social welfare metric. It needs to be carefully chosen to fit the specific scenario.

For a given application Nash social welfare metric has advantages over Utilitarian or Egalitarian viewpoints [7]. It leverages between fairness by equalizing utility distribution among agents and higher overall utility. However, we cannot define the social welfare metric as in Equation 1.

$$S_N(x) = \prod_{a_i \in N, t \leq T, t \in \mathbb{N}^+} u_{it}(\delta_t) \tag{1}$$

Consider the scenario presented in Table I and two allocation vectors $x_1 = (a_1, a_1, a_2), x_2 = (a_2, a_2, a_1)$. Both vectors are equally valuable in terms of fairness, as each of the two agents gets his share of resources and the division is as equal as possible. After applying Nash social welfare metric we get $S_N(x_1) = 250 > S_N(x_2) = 200$, thus it clearly favors x_1 over x_2. From the utilitarian perspective we have

TABLE I
SAMPLE WELFARE VALUES

t	δ_t	$u_{1,t}$	$u_{2,t}$
1	a_1	5	1
1	a_2	1	10
2	a_1	5	1
2	a_2	1	2
3	a_1	10	1
3	a_2	1	10

$S_U(x_1) = 23 < S_N(x_2) = 25$, what means that overall utility is lower in case of allocation x_1.

An ideal social welfare metric for this environment should in the first place ensure that all rules of fair discussion proposed in Section II-A are fulfilled. Once this is secured, higher overall profit should be promoted. We propose the conference metric given by equation 2, which mixes the concepts of utilitarian social welfare and Nash product together.

$$S_M(x) = \prod_{a_i \in N} \sum_{t \leq T, t \in \mathbb{N}^+} u_{it}(x_t) + \epsilon \tag{2}$$

A. Unanymity

Unanimity principle is the most important concept of welfare economics [8]. It says that the chosen utility vector should not be Pareto inferior to any other feasible utility vector. Metric defined by Equation 2 fulfills the rule of unanimity, as it prefers allocations which bring strongly Pareto optimal utility vectors. Let x_1 be a feasible allocation vector preferred by the metric. Let x_2 be a different feasible allocation dominating x_1 in terms of Pareto domination. Therefore we have:

- $u_{i,t}(x_{1,t}) \leq u_{i,t}(x_{2,t}), \forall a_i \in N, \forall t < T, t \in \mathbb{N}^+$
- $\exists a_i \in N, \exists t < T, t \in \mathbb{N}^+, u_{i,t}(x_{1,t}) < u_{i,t}(x_{2,t})$

It is easy to see, that $S_M(x_1) < S_M(x_2)$. This means that x_2 would be a preferred allocation vector by the metric.

Note that it does not take place if we remove ϵ from the equation. For instance, if one of the agents had zero utility for every possible allocation, all allocation vectors would be seen as equal.

B. Anonymity

Anonymity (symmetry across the agents) is another important social welfare metric feature. It is especially significant while considering fairness, as it indicates if any member of the society is in a privileged position. S_M is anonymous to some extent. For any choice of agents a_1 and a_2, switching their utility functions, so that $u_{1t}^{'} = u_{2t}, \forall t \leq T$ and $u_{2t}^{'} = u_{1t}, \forall t \leq T$ will not change the rating of any allocation vector. This only holds true if we perform so for all values of t. All allocations $\delta_j, \forall j < t$ and utility functions $u_{n,t}, \forall j < t$ act as an allocation history, which has a significant impact while deciding on δ_j. This historical data makes some agents more privileged then others. That is all in line with our understanding of fairness.

V. FAIR ALLOCATION MODEL

The proposed model is designed based on two assumptions:

- The whole resource is cofounded by each of the participating agents, therefore each member of the society owns an equal share of rights.
- Every agent may only grant some utility to owning the resource himself—$u_{it}(a_j) = 0, \forall i \neq j$. A more general allocation model with no such restriction has been proposed in [4], but it is out of scope of this paper.

The resource is indivisible and can only be utilized if fully owned. No one can use just a part of resource, 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_{it}, which might be interpreted as agent's wallet for time period (allocation number) t. Initially $r_{i1} = R, \forall a_i \in N$, where R is a positive constant value to initialize all the wallets. The conference is divided into T 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 period t is following:

1) Each participant a_i issues a bid with the valuation v_{it}. The bid cannot be higher than the actual wealth of agent at that time, therefore $v_{it} = \min(u_{it}(a_i), r_{it})$
2) The winning agent a_k and the price to pay p_t is determined with the use of Vickrey auction.
3) Price to pay is deducted from the winner's wallet $r_{k,t+1} = r_{k,t} - p_t$.
4) All agents which sell their resource rights to a_{kt} are rewarded $r_{i,t+1} = r_{i,t} + \frac{p_t}{|N|-1}$.
5) The whole resource is allocated to $\delta_t = a_k$ for time period t

According to the introduced model, resource allocation time is constant and defined upfront. It is a serious limitation given that demands might vary in terms of allocation length. If allocation time was kept constant and set too long, we would waste the resource due to excess allocation period. On the other hand, setting this time too short would cause allocations which only partially meet the demands. Moreover, there is no guarantee that agents assign any utility at all to allocations which only partially meet their demands.

We deal with this problem by setting an allocation time long and introducing a resale procedure. Whenever an agent owns a resource which he no longer requires while there is still some time left before exclusive rights period ends, he can request for an earlier reallocation. If a demand for the resource exists, he might get a fraction of the price he paid back.

Let $\delta_t = a_i$ and $\delta_{t+1} = a_j$. Agent a_i decided to request a resale after utilizing the resource for time l out of full allocation time L. When a_j wins:

- $r_{i,t+2} = r_{i,t+1} + \frac{l}{L}\min(p_t, p_{t+1}) + \frac{p_t - \frac{l}{L}\min(p_t, p_{t+1})}{|N|-1}$
- $r_{k,t+2} = r_{k,t+1} + \frac{p_t - \frac{l}{L}\min(p_t, p_{t+1})}{|N|-1}, \forall k \neq i, j$

It is important that an agent will never gain from the resale procedure. It is only possible to get a fraction of an allocation price back. Otherwise, it would motivate members of the society to purchase a resource with the hope of future resale at a higher price.

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. It is important to mention here that in the dominant strategy equilibrium is a weak equilibrium for VCG process in case of asymmetrical bid ranges [10]. For the richest agent it is equally reasonable to bid anything from the full value of his wallet to a wallet value of the second richest agent. In this case however, it has no effect on the choice of winner or on the price.

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.

VI. METHODOLOGY

In order to verify the proposed model and how it operates as conference moderation mean, we performed a whole range of simulated discussions. We implemented two types of agents playing a role in the conference—coordinator and participant. Implementation was done with the use of JADE platform [11]. Coordinator is responsible for keeping track and deciding whenever the society should transition to a next allocation phase. It is also up to the coordinator to decide upon the resource allocation based on all data (utilities) collected from participants. Participant's main responsibility is to declare its current utility of holding the resource whenever coordinator announces the bidding phase.

A. Protocol

Naturally, interaction among agents requires a specified communication language. FIPA standard defines a huge set of protocols, which come implemented out of the box with JADE. We have chosen some of as a language for our agents.

”Subscribe” protocol—defined by the FIPA standard. Should be used, whenever ”Initiator” agent wishes to monitor the state of an object owned by the ”Receiver”. For the purpose of this simulation, subscribe protocol is used by the participants to monitor state changes of the communication channel. This information is broadcasted by the coordinator after each performed allocation.

”Allocation” protocol—FIPA standard does not provide a Vickrey auction protocol. We designed the protocol of our own as shown in Figure 1. Every auction is initiated by the coordinator agent, who collects bids from all participants and performs the allocation according to the implemented model.

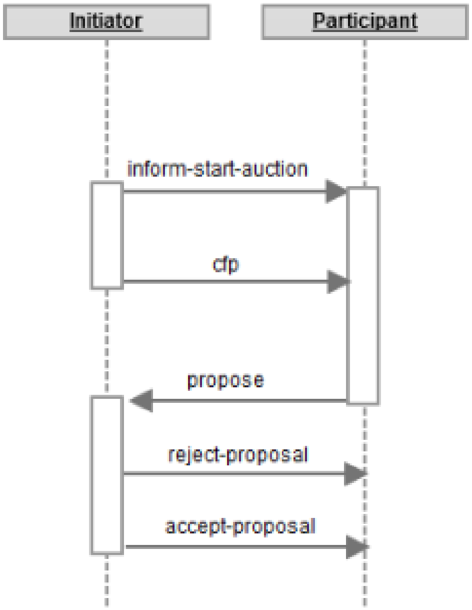

Fig. 1. Allocation protocol

”Request” protocol—defined by the FIPE standard. Used whenever an ”Initiator” agent asks the ”Receiver” to perform some action. In the discussion simulation, this protocol is used by a participant agent to request for an earlier reallocation procedure.

B. Behaviours

The whole process begins with a subscription phase when coordinator awaits for all participants to submit their conference subscription requests. After a certain time limit the discussion moves on to the main phase, which consists of two alternate behaviors:

- Auction—Collecting bids (valuations) from all conference participants. Making an allocation decision based on collected bids. Informing participants about the auction result.
- Allocation—Granting the resource to the new owner. Collecting and distributing the payment. Informing participants about the resource ownership change.

The whole auction process along with a decision upon the next allocation change is performed in advance (auction speedup time) to the actual change, in the background of the previous allocation time. This procedure is designed to eliminate allocation time fluctuations caused by the auction procedure, as it might involve passing a substantial number of messages over the network or performing timely calculations.

C. Demands

The main task of participant agent is to compete over a resource (communication channel) according to his preferences. Determining current utility is performed based on the agents set of demands. A single demand object is shown in the Figure 2. It is characterized by the following features

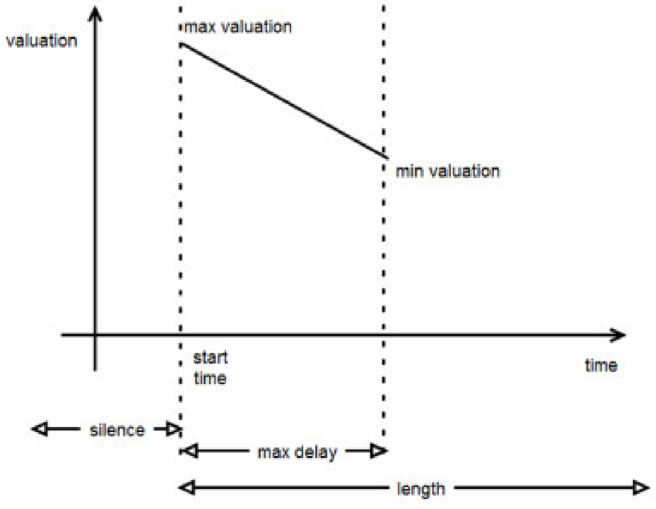

Fig. 2. Demand object

- Start time—the time since when the demand is active and has influence on agent's utility
- Maximal delay—maximal time in which the demand is still active
- Length—allocation time required to satisfy this demand
- Maximal valuation—valuation for an allocation at the demand start time
- Minimal valuation—valuation for an allocation with a maximal delay
- Silence time—time interval from the previous demand

In order to determine the current resource valuation, an agent scans through the whole set of active demands and picks the one with highest active valuation. If a demand is not satisfied right when it arises (start time), valuation deteriorates linearly up till maximal delay.

D. Dynamic demands

Aiming to increase similarity to real live conversational environment, we implemented semi intelligent agents capable of dynamically shaping their utility functions based on discussion history. Human conversation is a process driven by certain rules [12]. Among others, these rules describe how people choose and change the discussion topic:

- RULE 3: In introducing a new topic of conversation, the topic should be chosen so that both speakers have some knowledge and interest in its discussion.
- RULE 6: The topic of conversation may drift to a subject where the conversational participants share a great amount of knowledge.
- RULE 13: Each participant in the conversation has the conversational goal of saying things that are important to the other participant.

Having that in mind, we have developed an agent capable of reacting (adjusting his utility function) to the discussion flow. Such agent is supplied with a set of topic preferences and discussion memory to record the occurrence frequencies of all topics. Given all that as an input data, whenever an agent takes active part in a discussion, it will first choose a theme possibly interesting for all parties and shape it's utility function accordingly. Every topic from the list might be chosen with the probability directly proportional to agent's preferences and the frequency of occurance in the discussion so far.

E. Allocation models

We performed simulations for four different allocation models acting as a moderation procedure. Apart from the "Fair allocation model" proposed in Section V, we use three other allocation schemes for the purpose comparison:

- Queuing—Participants reporting need for the resource (positive valuation) are put into FIFO queue. Resource is always allocated to the first agent in the queue. There is no risk of dominating the discussion, yet it does not give any preference to allocations which increase the overall social welfare.
- Choosing maximal valuation—Resource is allocated to an agent reporting highest valuation at the given time. This model does not encourage agents to bid their true valuations nor does it put any preference to fair allocations. The dominating strategy is to bid just a tiny bit above the highest bidding participant.
- Maximizing social welfare metric—Chooses the allocation which maximizes social welfare metric defined by Equation 2. Prefers allocations which are both fair and increase the overall social welfare. However, participants might try to increase their profits or dominate the discussion by faking their bids. Under the assumption that allocation decision does not have any impact on the future shape of agents' utility functions, this model guarantees to maximize the social welfare.

VII. EVALUATIONS

We came up with three discussion participant profiles which vary in characteristics of their demands set. All demand attributes were generated from uniform distribution on a given range.

- Regular:
  - Silence time: 3 to 15 seconds
  - Maximal delay: 3 to 7 seconds
  - Length: 3 to 10 seconds
  - Maximal valuation: 2 to 10
  - Minimal valuation: 0 to half of maximal valuation
- Aggressive:
  - Silence time: 3 to 3,5 seconds
  - Maximal delay: 3 to 7 seconds
  - Length: 3 to 10 seconds
  - Maximal valuation: 10
  - Minimal valuation: 9 to 10
- Dynamic: (as described in Section VI-D):
  - Silence time: 3 to 15 seconds
  - Maximal delay: 3 to 7 seconds

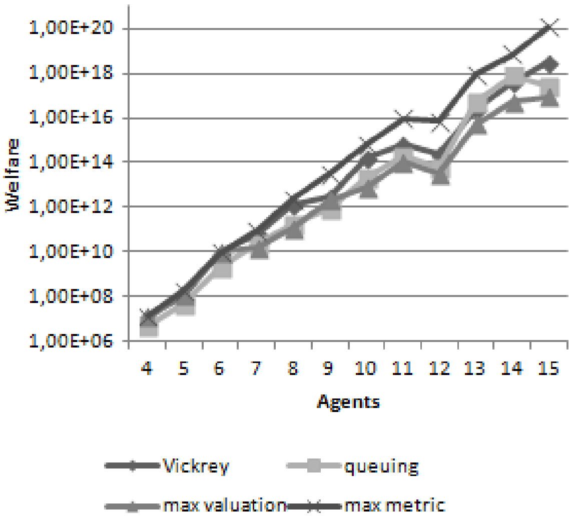

Fig. 3. Regular demands

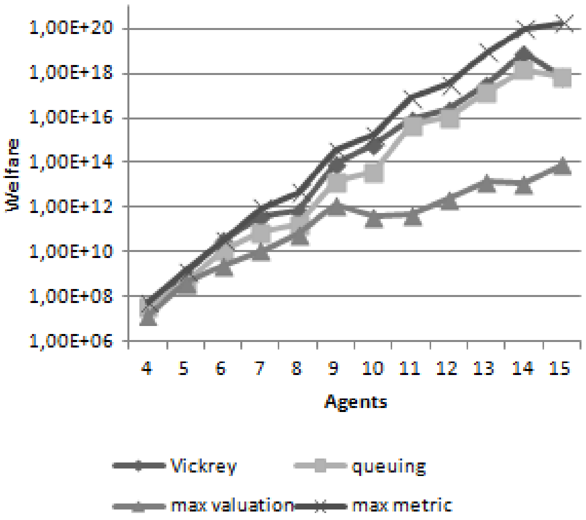

Fig. 4. Regular demands with 10 percent aggressive

  - Length: 3 to 10 seconds
  - Maximal valuation: 10 * topic preference
  - Minimal valuation: 0 to half of maximal valuation
  - Topics of interest: 5 randomly chosen topics out of total 10
  - Topic preference (for each topic of interest): 0 to 1

We performed model evaluations by simulating a number of discussions with the following parameters:

- Discussion time—5 minutes
- Allocation time—10 seconds
- Auction speedup time—3 seconds
- Number of participants—3 to 15

For a chosen distribution of demand profiles among participants, we performed 10 simulations for every number of participants from the range and every moderation model. Welfare has been calculated according to metric defined by Equation 2 and averaged over those 10 runs. Each agent had his set of demands generated before every single simulation.

Figure 3 contains results for the most ideal scenario, where every agent has a ”regular” demands profile. All allocation methods show decent behaviour, as there is was no special need to bother about fairness or overall welfare. ”Maximize metric” model outperforms all others, as expected.

In the next scenario, we had chosen the ceil of 10% of all agents and set their profiles to ”aggressive”. Results are shown in Figure 4. ”Choosing maximal valuation” clearly prefers allocation vectors which bring lower social welfare. It promotes aggressive agents over others and allows them to dominate the whole process.

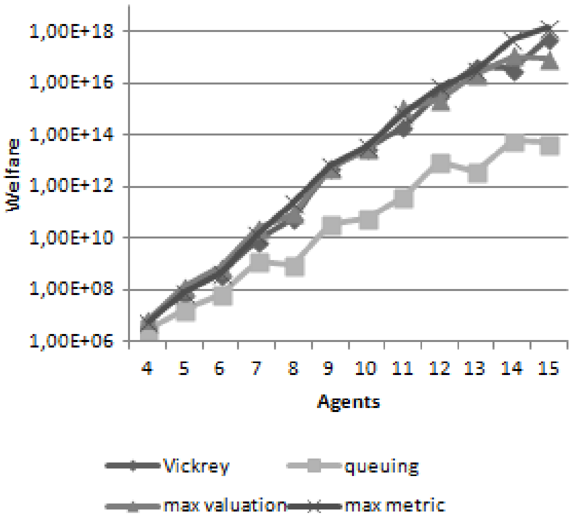

Fig. 5. Dynamic demands

Figure 5 shows results for a scenario where all participants represent ”dynamic” demands profile. In such environment ”Queuing” moderation model performed very poorly. It is due to the way it provides fairness. Allocating a resource to the first agent from the queue might very often prevent participants from finding a common set of preferred topics. ”Maximizing welfare metric” is no longer strictly superior to other allocation model, as the choice of allocation might have an impact on participants’ future utility functions.

The last scenario we simulated is a mixture of the previous two. Among the participants with ”dynamic” demands the chosen ceil of 10% of agents have their profiles set to ”aggressive”. This environment highlights the weaknesses of both ”queuing” and ”choosing maximal valuation” which perform visibly worse than the remaining two.

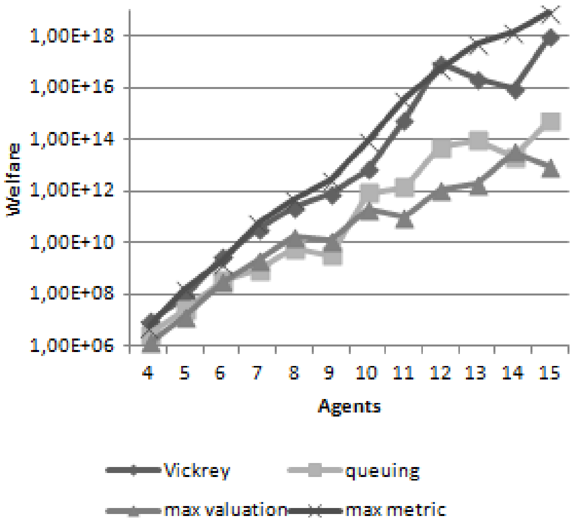

Fig. 6. Dynamic demands with 10 percent aggressive

VIII. CONCLUSIONS

In this paper we proposed a social welfare metric to determine the quality of a conference taking fairness aspects into consideration. Further on, we introduced a multiagent resource allocation scheme which embraces the described concept of fairness while getting as much out of overall social welfare as possible. We have performed a whole series of simulated discussions in order to verify the model's quality. Empirical results show that this resource allocation procedure is very much in line with the social welfare metric defined by Equation 2. It generates an allocation which is fair and highly valued by the whole society in a hostile environment with agents trying to dominate. We have also observed good results with semi-intelligent conversation aware agents.

Future work includes testing the model in a real life scenario by deploying to a broad public. We would also like to lay a more solid theoretical foundation to our definition of fair allocation. Analyzing whether our concept is in line with fair dominance [13] should be a good starting point. Another area of theoretical research are budget bound auction mechanisms [14] and investigate the impact on the introduced Vickrey auction based process.

REFERENCES

[1] 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.

[2] 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.

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

[4] A. Polomski, "Multiagent scheme for voice conference moderation," in *FedCSIS*, 2012, pp. 1215–1220.

[5] 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.

[6] K. J. Arrow, *Social Choice and Individual Values, Second edition (Cowles Foundation Monographs Series)*, 2nd ed. Yale University Press, Sep. 1970.

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

[8] H. Moulin, *Axioms of Cooperative Decision Making (Econometric Society Monographs)*. Cambridge University Press, Jul. 1991.

[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] M. H. Rothkopf, "Thirteen reasons why the vickrey-clarke-groves process is not practical," *Oper. Res.*, vol. 55, no. 2, pp. 191–197, Mar. 2007.

[11] F. Bellifemine, G. Caire, A. Poggi, and G. Rimassa, "JADE: A White Paper," *EXP in search of innovation*, vol. 3, no. 3, pp. 6–19, 2003.

[12] J. G. Carbonell, "Intentionality and human conversations," in *Proceedings of the 1978 workshop on Theoretical issues in natural language processing*, ser. TINLAP '78. Stroudsburg, PA, USA: Association for Computational Linguistics, 1978, pp. 141–148.

[13] W. Ogryczak, "Bicriteria models for fair and efficient resource allocation," in *SocInfo*, ser. LNCS, vol. 6430, 2010, pp. 140–159.

[14] H. Varian, "Position auctions," *International Journal of Industrial Organization*, vol. 25, no. 6, pp. 1163–1178, Dec. 2007.