A General Mechanism Design Methodology for Social Utility Maximization with Linear Constraints

Abhinav Sinha
EECS Department
University of Michigan, Ann Arbor
Michigan 48109, USA.
absi@umich.edu

Achilleas Anastasopoulos
EECS Department
University of Michigan, Ann Arbor
Michigan 48109, USA.
anastas@umich.edu

ABSTRACT

Social utility maximization refers to the process of allocating resources in a way that maximizes the sum of agents' utilities, under the system constraints. Such allocation arises in several problems in the general area of communications, including unicast (and multicast multi-rate) service on the Internet, as well as in applications with (local) public goods, such as power allocation in wireless networks, spectrum allocation, etc. Mechanisms that implement such allocations in Nash equilibrium have also been studied but either they do not possess the full implementation property, or are given in a case-by-case fashion, thus obscuring fundamental understanding of these problems. In this paper we propose a unified methodology for creating mechanisms that fully implement, in Nash equilibria, social utility maximizing functions arising in various contexts where the constraints are convex. Two additional design goals are the focus of this paper: a) the size of the message space scaling linearly with the number of agents (even if agents' types are entire valuation functions), b) allocation being feasible on and off equilibrium.

Keywords

Mechanism design, Full implementation, Game theory, Nash equilibrium, Proportional allocation, Social utility, Unicast service, Multi-rate multicast, Public goods.

1. INTRODUCTION

In the general area of communications, a number of decentralized resource allocation problems have been studied in the context of mechanism design. Such problems include unicast and multi-rate multicast service on the Internet [18], [11], [10], [9], [15], power allocation in wireless networks [14], spectrum allocation [5] and pricing in a peer-to-peer network [12]. The mechanism design framework is both appropriate and useful in the above problems since these problems are motivated by the designer's desire to allocate resources efficiently in the presence of strategic agents who possess private information about their level of satisfaction from the allocation.

Usually mechanism design solutions define a contract such that the induced game has at least one equilibrium that corresponds to the desired allocation. This is usually obtained with *direct* mechanisms by appealing to the revelation principle [1], [3]. The drawbacks of these direct approaches is that they require agents to quote their types (which may be entire valuation functions) and that the induced game may have other extraneous equilibria that are not efficient. In this paper, the focus is on **full Nash implementation**. Without going into a formal definition, full Nash implementation refers to the design of contracts such that only the designer's most preferred outcome is realized at equilibrium, as opposed to general mechanism design, where other less preferred outcomes are also possible.

Concentrating on those proposed solutions in the literature that guarantee full Nash implementation, one observes a fragmented and case-by-case approach. One may ask how fundamentally different are the various problems to justify a separately designed mechanism for each case. Alternatively, one may ask what are the common features in all these problems that can lead to a unified mechanism design approach. These questions provide the motivation for the work presented in this paper.

In particular, our starting point is to state a class of problems as social utility maximization under linear inequality/equality constraints. We define a general mechanism (by analyzing the dual optimization problem) and show that it results in full Nash implementation when the agents' valuation functions are their private information. Since the application domain of interest is in the area of communications, we give special emphasis on the size of the message space (as a consequence, VCG-type mechanisms[1] are inappropriate since they require quoting of types, which are entire valuation function in this set-up[2]). Further, the message space scales linearly with the number of agents so that the proposed mechanism is scalable.

Finally, a mechanism may be endowed with auxiliary properties off-equilibrium. These are meant to improve the applicability in practical settings. For instance, the NE is interpreted as the convergent outcome when agents in the system "learn" the game by playing it repeatedly, which implies that during this process the messages (and thus allocations and taxes) are off equilibrium. This prevents each player from having to calculate the Nash equilibrium by themselves and the whole system can collectively learn the NE via an offline Learning process (for example by best-responding to each others' average past moves). In our opinion, the property of **feasibility of allocation on and off equilibrium**

[1] see [1, Chapter 5], [3].
[2] See [8, Section 5] for work that adapts VCG for this issue.

is essential whenever the system constraints are hard constraints on resources and cannot be violated by any means. One of the main features of the work in this paper is that the allocation scheme is designed to guarantee feasibility off equilibrium. This is achieved by utilizing *radial allocation*, which will be seen in Section 3.1.

Due to space limitations, all the proofs are omitted in the exposition here. Readers may refer to the technical report [17] for more discussion and proofs.

2. CENTRALIZED PROBLEM

2.1 Some interesting Resource allocation problems in Communications

Unicast Transmission on the Internet

Consider agents on the Internet from set $\mathcal{N} = \{1, \dots, N\}$, where each agent $i \in \mathcal{N}$ is a pair of source and destination that communicate via a pre-decided route consisting of links in $\mathcal{L}_i$. All the agents together communicate on the network consisting of links in $\mathcal{L} = \cup_{i\in\mathcal{N}}\mathcal{L}_i$. Since each link in the network has a limited capacity, this results in constraints on the information rate allocated to each agent. Considering a scenario where agents have (concave) utility functions/profiles $\{v_i(\cdot)\}_{i\in\mathcal{N}}$ that measure the satisfaction received by agents for various allocated rates, we can write the social utility optimization problem as

$$\max_{x\in\mathbb{R}_+^N} \sum_{i\in\mathcal{N}} v_i(x_i) \qquad (\mathrm{CP_u})$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{N}} \alpha_i^l x_i \le c^l \quad \forall\, l \in \mathcal{L}. \qquad (1)$$

In the above, c^l is the capacity of link l and α_i^l are non-negative weights meant to differentiate between true information rate x_i and its imposition on capacity of the links of the network through coding rate, packet error etc.

The above feasible set is a polytope in the first quadrant of $\mathbb{R}_+^N$ and is created by faces that have outward normal vectors pointing away from the origin. For the details of a full implementation mechanism specifically for the unicast problem readers may refer to [7], [10], [16].

Public Goods and Local Public Goods

In contrast to the private consumption problem above, there are public goods problems where the resources are shared directly between agents instead of sharing via constraints. A well-studied example of this kind is the wireless transmission with interference, in which the power level of each agent affects any other agent through the signal-to-interference-and-noise ratio (SINR).

Here we consider a general local public goods problem which encompasses the public goods problem. With the same basic idea of direct sharing of resources, here the sharing is only among agents locally. So there are local groups of agents for whom the allocation has to be the same, but this common allocation can be different from one group to the next. If we divide the set of agents into disjoint local groups: $\mathcal{N} = \sqcup_{k\in\mathcal{K}}\mathcal{N}_k$, then the centralized problem can be

$$\max_{x\in\mathbb{R}_+^K} \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{N}_k} v_i(x_k) \qquad (\mathrm{CP_{lpb}})$$

$$\text{s.t.} \quad x \in \mathcal{X} \subset \mathbb{R}^N \qquad (2)$$

with $\mathcal{X}$ being a convex subset of $\mathbb{R}^N$. Note that the argument for all utility functions in a local group is the same; since there is one public good being simultaneously used by all the agents in that group (for which x_k marks the usage level). One can rewrite the above with separate allocation for each user with additional constraint of equal allocations within each group being added additionally in the constraints.

Wireless transmission can be seen as a specific local public goods example since each agent affects the SINR only locally (either spatially or in the frequency domain). Implementation for the public goods problem is the most studied of all the examples in this paper, see [4], [6], [2]. Readers may refer to [14] for a specific mechanism for the local public goods problem.

2.2 General Centralized Problem

Combining the features of problems such as above, we define the resource allocation problem for a system with agents indexed in the set $\mathcal{N} = \{1, 2, \dots, N\}$, who have utility functions $\{v_i(\cdot)\}_{i\in\mathcal{N}}$. The objective is to find the optimum allocation of a single infinitely divisible good that maximizes the sum of utilities subject to constraints on the system. The allocation made to agents will be denoted by the vector $x \in \mathbb{R}_+^N$, with x_i being the allocation to agent i. The centralized optimization problem that we consider is

$$\max_x \sum_{i\in\mathcal{N}} v_i(x_i) \qquad (\mathrm{CP})$$

$$\text{s.t.} \quad x \in \mathbb{R}_+^N \qquad (\mathrm{C_1})$$

$$\text{s.t.} \quad A_l^\top x \le c_l \ \ \forall\, l \in \mathcal{L} \ \text{ where } \ A_l \in \mathbb{R}^N, c_l \ge 0. \qquad (\mathrm{C_2})$$

The set $\mathcal{L} = \{1, 2, \dots, L\}$ indexes all the constraints and A_l, c_l are all parameters of the optimization problem. It is easy to see that the above set-up covers equality constraints (such as from the public goods example), since we can always write $x_1 = \ldots = x_N$ as $x_1 \ge x_2 \ge \ldots \ge x_N \ge x_1$.

Denote by $\mathcal{C} \subset \mathbb{R}_+^N$ the above feasible set. Note that $\mathcal{C}$ is a polytope in the first quadrant of $\mathbb{R}^N$, possibly of a lower dimension than N (due to equality constraints).

2.3 Assumptions

(A1) For any $i \in \mathcal{N}$, $v_i(\cdot)$ is a strictly concave and continuously double differentiable function $\mathbb{R}_+ \to \mathbb{R}$.

The purpose of strict concavity is to have a (CP) whose solution can be described sufficiently by the KKT conditions (note that monotonicity is not assumed).

(A2) The optimal solution $x^\star$ is bounded such that $x^\star \in \times_{i=1}^N (d_i, D)$ for some $0 < d_i < D$, with $d = (d_i)_{i=1}^N \in \mathcal{C}$ arbitrarily close to $\underline{0}$ and D being large enough.

This assumption is used to eliminate corner cases of (CP), since they usually require special treatment and make the exposition unnecessarily convoluted. Because of the choice $c_l \ge 0$ in ($\mathrm{C_2}$), we have $\underline{0} \in \mathcal{C}$. Thus one can always select a point $d \in \mathcal{C}$ arbitrarily close to $\underline{0}$. Since we are considering problems where all the variables x_i have a physical interpretation, it is natural to consider a constraint set whereby every agent getting 0 allocation is feasible. Since d can be chosen to be very close to $\underline{0}$, this assumption doesn't pose significant restrictions on the considered problem (CP).

(A3) For any constraint $l \in \mathcal{L}$ in (C_2) there are at least two distinct $i, j \in \mathcal{N}$ such that $A_{li}, A_{lj} \neq 0$.

This ensures that there is indeed competition for all the constraints that could possibly be active at optimum. Again this is used to avoid special treatment of corner cases.

Denote by $\mathcal{V}_0$ the set of all possible functions $\{v_i(\cdot)\}_{i\in\mathcal{N}}$ that satisfy the above assumptions. Then $\mathcal{V}_0$ will be the environment for our mechanism design problem. We also make an assumption on the overall utility of agents.

(A5) There is a linear taxation component that affects agents' utilities. So overall utility of an agent i is

$$u_i(x, t) = v_i(x_i) - t_i. \tag{3}$$

3. MECHANISM

A mechanism in the Hurwicz-Reiter framework consists of an environment, an outcome space, a (centralized) correspondence between the two, a message space and a contract from the message space to the outcome space. In our case the environment is the set $\mathcal{V}_0$, the outcome space is the set of all possible allocations and taxes, i.e. $\mathcal{C} \times \mathbb{R}^N$. The correspondence between $\mathcal{V}_0$ and $\mathcal{C}$ is provided implicitly by the centralized problem (CP), where for each $\{v_i(\cdot)\}_{i\in\mathcal{N}}$ we get an optimal allocation $x^\star$ by solving (CP). The designer thus has the task of designing the message space and the contract.

The *message space* for our mechanism is $\mathcal{S} = \times_{i\in\mathcal{N}}\mathcal{S}_i$ with $\mathcal{S}_i = (d_i, +\infty) \times \mathbb{R}_+^{L_i}$ where messages from agents are of the form $s_i = (y_i, p_i)$ with $p_i = (p_i^l)_{l\in\mathcal{L}_i}$ and the total message is denoted by $s = (s_i)_{i\in\mathcal{N}} = (y, P)$ with $y = (y_i)_{i\in\mathcal{N}}$ and $P = (p_i)_{i\in\mathcal{N}}$. The message $s_i = (y_i, p_i)$ is to be interpreted as follows: y_i is the level of demand from agent i and p_i is the vector of prices (per unit) that he believes everyone else should pay for the respective constraints. The *contract* $h : \mathcal{S} \to \mathbb{R}_+^N \times \mathbb{R}^N$ will specify allocation and taxes for all agents based on the message s, i.e., $h(s) = (h_{x,i}(s), h_{t,i}(s))_{i\in\mathcal{N}}$. This contract along with agents' utilities will give rise to a one-shot game $\mathfrak{G} = (\mathcal{N}, \times_{i\in\mathcal{N}}\mathcal{S}_i, \{\widehat{u}_i\}_{i\in\mathcal{N}})$ between agents in $\mathcal{N}$ where action sets are $\mathcal{S}_i$ and utilities are

$$\widehat{u}_i(s) = u_i(x, t) = v_i(x_i) - t_i = v_i(h_{x,i}(s)) - h_{t,i}(s). \tag{4}$$

Information assumptions: We assume that the designer doesn't know $\{v_i(\cdot)\}_{i\in\mathcal{N}}$ but knows the set $\mathcal{V}_0$. Also the constraints (C_1) and (C_2) in (CP) are common knowledge.

We say that the mechanism fully implements the centralized problem (CP) if **all** Nash equilibria of the induced game $\mathfrak{G}$ correspond to the unique allocation $x^\star$ - solution of (CP), and additionally individual rationality is satisfied i.e. agents are weakly better-off participating in the contract at equilibrium than not participating at all.

3.1 Allocation

To make the exposition clearer, from here onwards we only consider the case where the constraints in (C_2) do not have any effective degeneracy i.e. equality constraints. Kindly refer to the technical report [17] for analysis with equality constraints. This distinction is based on whether the feasible set $\mathcal{C}$ has a proper interior. Clearly in absence of equality constraints, the constraint set will have a proper interior.

For the allocation, we first choose a point θ in the interior of the feasible set such that

$$\theta \in \text{int}(\mathcal{C}) \cap \times_{i=1}^N (0, d_i) \tag{5}$$

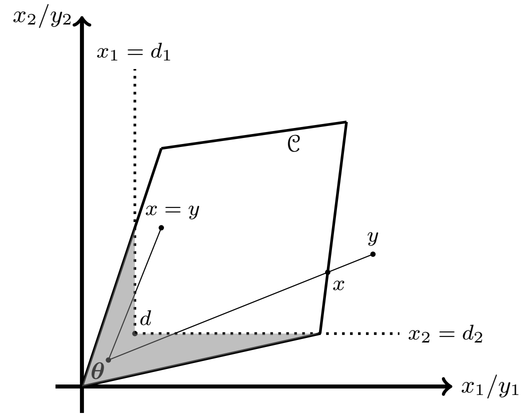

Figure 1: An illustration of the allocation for $N = 2$

Note that we can guarantee the existence of θ since $\underline{0} \in \mathcal{C}$ and clearly $\mathcal{C}$ being the intersection of half-planes, is a convex set. Since d can be made arbitrarily close to $\underline{0}$, the same holds for θ as well.

Before formally defining the allocation, we define it informally with the help of Fig. 1. For any demand $y \in \mathcal{S}_y \triangleq \times_{i=1}^N (d_i, +\infty)$, the allocation x will be equal to y if y is inside the feasibility region $\mathcal{C}$. Otherwise the allocation will be the intersection point between the boundary of the feasibility region $\mathcal{C}$ and the line joining θ with y the figure shows two different possible y's and their corresponding allocation x. The shaded region represents that part of $\mathcal{C}$ that can never be allocated. Note that since $\|d\|, \|\theta\| \approx 0$ this is a very small region and thus it doesn't significantly affect the generality of the results presented in this paper. Radial allocation is the reason why assumption (A2) was needed.

Formally, for a demand $y \in \mathcal{S}_y$, the allocation x is

$$x = \begin{cases} y & \text{if } y \in \mathcal{C} \\ y_0 & \text{if } y \notin \mathcal{C}, \end{cases} \tag{6}$$

where y_0 is the projected point on the boundary of $\mathcal{C}$. Explicitly, if the above intersection happens on the hyperplane $\mathcal{F}_l = \{A_l^\top x = c_l\}$ then

$$y_0 - \theta = \alpha_0(y - \theta) \quad \text{with} \quad \alpha_0 = \frac{c_l - A_l^\top \theta}{A_l^\top (y - \theta)} \tag{7}$$

The above allocation mapping is an extension of the generalized proportional allocation idea (see [15] and [16]), but modified in accordance with the generality of the problem (CP) and also so that points in the interior of $\mathcal{C}$ are covered as well. It is easy to verify that the allocation is continuous.

3.2 Taxes

For any agent i, we define his tax t_i as

$$t_i = \sum_{l\in\mathcal{L}_i} t_i^l \tag{8a}$$

$$t_i^l = A_{li} x_i \bar{p}_{-i}^l + \left(p_i^l - \bar{p}_{-i}^l\right)^2 + \eta\, \bar{p}_{-i}^l p_i^l \left(c_l - A_l^\top x\right)^2 \tag{8b}$$

$$\bar{p}_{-i}^l \triangleq \frac{1}{N^l - 1} \sum_{\substack{j\in\mathcal{N}^l \\ j\neq i}} p_j^l. \tag{8c}$$

with $\eta > 0$ being a positive constant.

4. EQUILIBRIUM RESULTS

For lack of space, we avoid stating intermediate lemmas here. More detailed technical analysis can be found in the technical report [17]. The main result in this section is the desired full Nash implementation property for the mechanism defined above. This would require us to prove that all pure strategy NE of game $\mathfrak{G}$ result in allocation $x^\star$ - the unique solution of (CP), and also to show individual rationality. The method for proving this result is as follows: firstly we show that for any pure strategy NE the corresponding allocation and quoted prices must satisfy the KKT conditions. Since by assumptions KKT conditions are both necessary and sufficient, this would mean that if pure strategy NE exist (unique or multiple), the corresponding allocation would have to be the solution of (CP) with quoted prices as the optimal Lagrange multipliers. Then we show existence by showing that for any type in the considered type space, there exists at least one message in the message space such that the corresponding allocation is the optimal $x^\star$ and it is also a Nash equilibrium.

For the intermediate lemma that states that NE must necessarily satisfy stationarity condition of KKT, we have to make assumptions on the feasible set $\mathcal{C}$.

(A6)
$$A_l \in \mathbb{R}_+^N \quad \forall\, l \in \mathcal{L} \tag{9}$$

5. GENERALIZATIONS

Two quite interesting generalizations arise immediately from the set-up and analysis in this paper. The first is a case where agents have utilities based on a vector allocation rather than a scalar allocation i.e. the multiple goods scenario. Note that the assumption of strict concavity can still be made. In a communications scenario, such an example can arise if the Internet agents have utility based on throughput as well as delay or packet error rate. This will also be particularly required if one wants to model a wireless network where the SINR of each agent is affected by the power allocated to neighbouring agents.

The second generalization is with problems which can be equivalently formulated in the form of (CP), but perhaps with the help of auxiliary variables. A canonical example of this is the multi-rate multicast problem. Readers may refer to [9], [15] for a full implementing mechanism specifically for the multi-rate multicast problem. The incorporation of this model into the unified mechanism design methodology is a research topic the authors are currently working on.

The aforementioned generalizations together can lead to fully implementing mechanisms with minimal message space that can solve an even larger class of problems of interest.

6. REFERENCES

[1] T. Börgers. *An Introduction to the Theory of Mechanism Design*. 2013. Retrieved February 11, 2014 from `http://www-personal.umich.edu/~tborgers/TheoryOfMechanismDesign.pdf`

[2] Y. Chen. A family of supermodular Nash mechanisms implementing Lindahl allocations. *Economic Theory*, 19(4):773–790, 2002.

[3] D. Garg, Y. Narahari, and S. Gujar. Foundations of mechanism design: A tutorial part 1-key concepts and classical results. In *Sadhana (Academy Proceedings in Engineering Sciences)*, volume 33, pages 83–130. Indian Academy of Sciences, 2008.

[4] T. Groves and J. Ledyard. Optimal allocation of public goods: A solution to the "free rider" problem. *Econometrica: Journal of the Econometric Society*, pages 783–809, 1977.

[5] J. Huang, R. A. Berry, and M. L. Honig. Auction-based spectrum sharing. *Mobile Networks and Applications*, 11(3):405–418, 2006.

[6] L. Hurwicz. Outcome functions yielding Walrasian and Lindahl allocations at Nash equilibrium points. *The Review of Economic Studies*, 46(2):217–225, 1979.

[7] R. Jain and J. Walrand. An efficient Nash-implementation mechanism for network resource allocation. *Automatica*, 46(8):1276–1283, 2010.

[8] R. Johari and J. N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. *Operations Research*, 57(4):823–839, 2009.

[9] A. Kakhbod and D. Teneketzis. Correction to "An efficient game form for multi-rate multicast service provisioning". IEEE Journal on *Selected Areas in Communications*, 31(7):1355–1356, 2013.

[10] A. Kakhbod and D. Teneketzis. Correction to "An efficient game form for unicast service provisioning". 2013.

[11] R. T. Maheswaran and T. Basar. Social welfare of selfish agents: motivating efficiency for divisible resources. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 2, pages 1550–1555. IEEE, 2004.

[12] M. J. Neely. Optimal pricing in a free market wireless network. *Wireless Networks*, 15(7):901–915, 2009.

[13] S. Reichelstein and S. Reiter. Game forms with minimal message spaces. *Econometrica: Journal of the Econometric Society*, pages 661–692, 1988.

[14] S. Sharma and D. Teneketzis. Local public good provisioning in networks: A Nash implementation mechanism. IEEE Journal on *Selected Areas in Communications*, 30(11):2105–2116, 2012.

[15] A. Sinha and A. Anastasopoulos. Generalized proportional allocation mechanism design for multi-rate multicast service on the internet. Technical report, 2013. Retrieved February 11, 2014 from `http://arxiv.org/abs/1307.2569`

[16] A. Sinha and A. Anastasopoulos. Mechanism design for unicast service on the internet. Technical report, University of Michigan, 2013. Retrieved February 11, 2014 from `http://www.eecs.umich.edu/techreports/systems/cspl/cspl-414.pdf`

[17] A. Sinha and A. Anastasopoulos. A general mechanism design methodology for social utility maximization with linear constraints. Technical report, University of Michigan, 2014. Retrieved April 7, 2014 from `http://web.eecs.umich.edu/~anastas/docs/netecontechreport.pdf`

[18] S. Yang and B. Hajek. Revenue and stability of a mechanism for efficient allocation of a divisible good. *preprint*, 2005. Retrieved June 8, 2014 from `http://www.ifp.illinois.edu/~hajek/Papers/YangHajek06.pdf`