One-Way Games
(Extended Abstract)

Andres Abeliuk
NICTA and
University of Melbourne
andabek@gmail.com

Gerardo Berbeglia
NICTA and
Melbourne Business School
g.berbeglia@mbs.edu

Pascal Van Hentenryck
NICTA and
Australian National University
pvh@nicta.com.au

Categories and Subject Descriptors

[**Algorithmic game theory and mechanism design**]: Algorithmic mechanism design; [**Distributed artificial intelligence**]: Cooperation and coordination

General Terms

Algorithms, Design, Economics

Keywords

Mechanism Design, Price of Anarchy, Distributed Problem Solving

1. INTRODUCTION

This paper is motivated by optimization applications involving multiple stakeholders. Examples of such applications are found in large-scale restoration of interdependent infrastructures after significant disruptions (e.g.,[1, 2]), humanitarian logistics over multiple states or regions (e.g., [3]), supply chain coordination (e.g., [4]) and integrated logistics involving, say, a port, inland terminals, and railway and truck operators.

Consider, for instance, the restoration of the power system and the telecommunication network after a major disaster. As explained in [3], there are one-way dependencies between the power system and the telecommunication network. This means, for instance, that some power lines must be restored before some part of the telecommunication network can become available. It is possible to use centralized mechanisms for restoring the system as a whole. However, in practice, it is often the case that these restorations are performed by different agencies with independent objectives and selfish behavior may have a strong impact on the social welfare. It is thus important to study whether it is possible to find high-quality outcome to these problems in decentralized settings when the stakeholders proceed independently and do not share complete information about their costs.

This paper aims at taking a first step in this direction. We propose a class of one-way games that abstracts some of the salient features of these applications. In this first approach, we restrict the attention to one-way dependencies.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

These dependencies capture a significant class of applications in infrastructure restoration, supply-chain coordination, and integrated logistics. Our goal is to find ways of incentivizing players to cooperate to achieve a better social welfare outcome.

2. ONE-WAY GAMES

A one-way game features 2 players A and B. Each player $i \in A, B$ has a **public** strategy set $\mathcal{S}_i$ and we write $\mathcal{S} = \mathcal{S}_A \times \mathcal{S}_B$ to denote the set of joint strategies profiles. Each player also has a **private** payoff function $u_i : \mathcal{S} \to \mathbb{R}^+$. The payoff $u_A(s_A, s_B) = u_A(s_A)$ of player A is determined only by its own strategy, whereas the payoff $u_B(s_B, s_A)$ for player B is determined by the strategies selected by both players. Notice that the payoff value $u_B(s_i, s_j)$ for player B is known for any strategy $s_i, s_j \in \mathcal{S}$. Player B however has no information about the strategy that player A will select.

A key feature of one-way games is that the payoff of player A is independent of player B, while B must act according to what player A chooses to do. Thus, player B is at a disadvantage with respect to player A. This fact entails the following lemma.

LEMMA 2.1. *There always exists a pure Nash equilibrium in one-way games. Furthermore, Player A has its optimal payoff in every Nash Equilibrium.*

Our motivating applications aim at optimizing a global welfare function $SW(s_A, s_B) = u_A(s_A) + u_B(s_A, s_B)$. Since player A does not depend on B's action, her strategy may induce a bad outcome for player B even when B has a much greater payoff. We quantify the quality of Nash equilibria with the *price of anarchy* (PoA).

LEMMA 2.2. *The price of anarchy for one-way games is*

$$\frac{\max_{s \in \mathcal{S}} u_B(s)}{\max_{s \in \mathcal{S}} u_A(s)} \leq PoA \leq 1 + \frac{\max_{s \in \mathcal{S}} u_B(s)}{\max_{s \in \mathcal{S}} u_A(s)},$$

The price of anarchy can thus be arbitrarily large. When it is large enough, Lemma 2.2 indicates that in this case, player B has a strong bargaining power to monetary incentivize player A to cooperate. This paper explores this possibility.

3. THE VCG MECHANISM

Since we are interested in solving a welfare maximization problem with private values, it is useful to determine the properties of the VCG mechanism in this context. VCG is a direct revelation mechanism that truthfully implements the welfare maximizing social choice function but may not be weakly budget-balanced. Where weakly budget-balanced means that there are no net transfers into the system. Additionally, the mechanism is ex-post individually rational, meaning that no individual wishes to walk away from a mechanism after all information has been revealed.

To quantify how unbalanced is the VCG mechanism, consider the social cost of the VCG mechanism defined as follows.

Definition 3.1. *The social cost (SC) of the VCG mechanism is the ratio between the optimal social outcome without and with the payments of the mechanism.*

The payments are always negative making the social cost strictly greater than 1. Furthermore, using the explicit value of the payments yields $SC = \frac{\max_{s \in \mathcal{S}} SW(s)}{\max_{s_A \in \mathcal{S}_A} u_A(s_A)}$.

Lemma 3.2. *The social cost of implementing the optimal social outcome can be at least as large as the price of anarchy.*

4. BUDGET-BALANCED MECHANISMS

An interesting starting point is the recognition that, when player B has a better payoff than A, player A may let player B play her optimal strategy in exchange for money. The resulting outcome can be viewed as swapping the roles of both players, i.e., player B chooses her optimal strategy and A plays her best response to B's strategy. This observation leads to the following lemma.

Lemma 4.1. *Consider $s' = \arg\max_{s \in \mathcal{S}} \max(u_A(s), u_B(s))$. In the one-way game, strategy s' has a price of anarchy of 2.*

Lemma 4.1 gives us hope for the design of an budget-balanced mechanism that has a constant price of anarchy. Indeed, a simple and distributed implementation of Lemma 4.1 would ask each player to reveal their maximal payoff value and then choose the best strategy to be implemented. If the strategy proposed by player B should be implemented, then player A must receive a monetary compensation for deviating from her maximal strategy.

Lemma 4.2. *There is no ex-post individual rational, budget-balanced mechanism that implements strategy s' from Lemma 4.1.*

This impossibility result comes from the fact that both players have a positive outside option. We have found a counter-example showing that it is not possible to implement strategy s' without violating the individual rational constraint.

5. SINGLE-OFFER MECHANISM

We now consider a bargaining game under a Bayesian setting where each player has private utilities and a belief about the other player utilities. We assume that the default strategy $s_A^* \in \arg\max_{s \in \mathcal{S}_A} u_A(s)$ of player A (i.e. the action A would choose if no monetary incentive is given), is publicly known. This single-offer mechanism is defined as follows:

1. Player B determines $s'_A \in \mathcal{S}_A$ and $s'_B(s'_A) \in \mathcal{S}_B$ such that $s'_B(s'_A) \in \arg\max_{s \in \mathcal{S}} u_B(s)$, i.e., the strategy that maximizes her utility.
2. Player B computes $s_B^*(s_A^*) \in \arg\max_{s_B \in \mathcal{S}_B} u_B(s_A^*, s_B)$, i.e., the best response to player A's default strategy.
3. Player B proposes a monetary value of $\gamma \cdot \Delta_B$ with $\Delta_B = u_B(s'_B(s'_A)) - u_B(s_B^*(s_A^*))$ and $\gamma \in \mathbb{R}_{[0,1]}$ to player A if she accepts to play strategy s'_A rather than her default strategy s_A^*.
4. Player A decides whether to accept the offer.
5. If player A accepts the offer, the game is played with strategy (s'_A, s'_B); Otherwise the outcome of the game is $(s_A^*, s_B^*(s_A^*))$.

It worth to observe that a broker is required in this mechanism to ensure that the strategy $(s_A^*, s_B^*(s_A^*))$ is implemented, if player A rejects the unique offer.

Proposition 5.1. *Player A accepts the offer whenever $\gamma \cdot \Delta_B \geq \Delta_A$, where $\Delta_A = u_A(s_A^*) - u_A(s'_A)$.*

Let $F(\cdot)$: $F(x) \in [0,1]$ is the probability that $\Delta_A \leq x$ and thus $F(\gamma \cdot \Delta_B)$ is the probability that player A accepts the offer $\gamma \cdot \Delta_B$. $F(x)$ is assume to be public.

We have designed a mechanism that satisfies individual rationality by construction. Player B never offers more than Δ_B and its payoff is never worse than the default strategy $s_B^*(s_A^*)$.

Lemma 5.2. *Player's B expected utility is maximized when she offers $\gamma^* \cdot \Delta_B$, where $\gamma^* = \arg\max_\gamma F(\gamma \cdot \Delta_B) \cdot (1 - \gamma)$.*

We now derive the induced *price of anarchy* for the single-offer mechanism.

Lemma 5.3. *Let $PoA^A(\gamma)$ and $PoA^R(\gamma)$ denote the induced price of anarchy if player A accepts and rejects respectively. Then, $PoA^A(\gamma) = 1 + \gamma$ and $PoA^R(\gamma) = 1 + \frac{1}{\gamma}$.*

We are ready to show the main property of this mechanism.

Theorem 5.4. *The Bayes-Nash price of anarchy is*

$$\frac{\gamma^* + 1}{\gamma^*} \left(1 - F(\gamma^* \cdot \Delta_B)(1 - \gamma^*)\right).$$

For example, if $F = U(0, \Delta_B)$, then $\gamma^* = \frac{1}{2}$ and thus the expected price of anarchy is 2.25.

6. REFERENCES

[1] B. Cavdaroglu, E. Hammel, J. E. Mitchell, T. C. Sharkey, and W. A. Wallace. Integrating restoration and scheduling decisions for disrupted interdependent infrastructure systems. *Annals of Operations Research*, pages 1–16, 2013.

[2] C. Coffrin, P. Van Hentenryck, and R. Bent. Last-mile restoration for multiple interdependent infrastructures. In *AAAI*, 2012.

[3] P. Van Hentenryck, R. Bent, and C. Coffrin. Strategic planning for disaster recovery with stochastic last mile distribution. In *Integration of AI and OR techniques in constraint programming for combinatorial optimization problems*, pages 318–333. Springer, 2010.

[4] G. Voigt. *Supply Chain Coordination in Case of Asymmetric Information*. Springer, 2011.