10-2015

How to Compute von Neumann-Morgenstern Solutions

Martha Osegueda Escobar
The University of Texas at El Paso, mcoseguedaescobar@miners.utep.edu

Vladik Kreinovich
The University of Texas at El Paso, vladik@utep.edu

Comments:

Technical Report: UTEP-CS-15-76

To appear in *Mathematical Structures and Modeling*, 2016, Vol. 37.

III-я Международная научная конференция ФКН
«Математическое и компьютерное моделирование»,
Россия, Омск, ОмГУ, 12 ноября 2015 г. – С. 1–6. УДК 234.128 : 345.561

Martha Osegueda Escobar and Vladik Kreinovich
Department of Computer Science
University of Texas at El Paso, El Paso, Texas 79968, USA
mcoseguedaescobar@miners.utep.edu, vladik@utep.edu

HOW TO COMPUTE VON NEUMANN-MORGENSTERN SOLUTIONS

Cooperative games and von Neumann-Morgenstern solution. Situations when all participants collaborate with each other is known as a *cooperative game*. One way to describe a cooperative game is to assign, to every subset $S \subseteq N \stackrel{\text{def}}{=} \{1, \ldots, n\}$ of the set of all the participants, the value $v(S)$ that describes what players from S can gain if they collaborate between themselves only. Such subsets S are called *coalitions*.

We consider cooperative situations, so if two disjoint coalitions S and S' collaborate, they should be able to gain not less than they would get on their own, i.e., we should have $v(S \cup S') \geq v(S) + v(S')$.

It always makes sense to consider only gains due to collaboration, so if $v(\{i\}) \neq 0$, then we can take $v'(S) = v(S) - \sum\limits_{i \in S} v(\{i\})$ for which $v'(\{i\}) = 0$.

Definition 1. *Let n be a positive integer. By a* cooperative game, *we mean a function that assigns, to each subset $S \subseteq N \stackrel{\text{def}}{=} \{1, \ldots, n\}$, a non-negative number $v(S)$ so that*

- *$v(\{i\}) = 0$ for all i, and*
- *when $S \cap S' = \emptyset$, then $v(S \cup S') \geq v(S) + v(S')$.*

Since everyone is collaborating, the participants together get the value $v(N)$. The question is: what is a fair way to divide this total amount $v(N)$ between n participants, i.e., how to allocate non-negative amounts $x_1, \ldots, x_n$ for which $\sum\limits_{i=1}^{n} x_i = v(N)$. Such allocations are known as *imputations*.

Definition 2. *By an* imputation, *we mean a tuple $(x_1, \ldots, x_n)$ of non-negative numbers for which $\sum\limits_{i=1}^{n} x_i = v(N)$.*

In their original book [6] that started game theory, John von Neumann and Oscar Morgenstern considered the following notion of *dominance* between imputations $x \succ y$. We say that an imputation x *dominates* an imputation y if there exists a coalition S for which each player from S gets more money in x than in y ($x_i > y_i$) and which is "reachable" for S – i.e., for which $\sum_{i \in S} x_i \le v(S)$.

The condition $\sum_{i \in S} x_i \le v(S)$ means that the coalition S can force a switch from y to x, and the condition $x_i > y_i$ means that this switch is beneficial for all members of the coalition S. So, if both imputations x and y are possible, S will force a switch from x to y.

Definition 3. *We that an imputation x* dominates *an imputation y (and denote it by $x \succ y$) if $x_i > y_i$ for all $i \in S$ and $\sum_{i \in S} x_i \le v(S)$.*

At first glance, it may seem reasonable to consider the set of all non-dominated imputations as a solution; this set is known as a *core* [3, 6]. Alas, often, this set is empty: we may have both $x \prec y$ (because of one coalition S) and $y \prec x$ (because of another coalition S'). In such games, if we allow all possible imputations, we can potentially switch infinitely many times, never reaching an equilibrium.

To avoid such situations, von Neumann and Morgenstern suggested that we adopt some *social norms* that would limit the set of possible imputations in such a way that no two imputations within this norm dominate each other. The social norm has to be *enforceable* meaning that if someone proposes an imputation which is outside this norm, there should be a coalition that forces a switch to a solution within the norm. The resulting definition is known as a *von Neumann-Morgenstern solution* (or *NM-solution*, for short).

Definition 4. *A set C of imputations is called a* von Neumann-Morgenstern solution *if it satisfies the following two properties:*

- *if $x, y \in C$, then $x \not\succ y$;*
- *if $y \notin C$, then there exists an $x \in C$ for which $x \succ y$.*

In this case, our decision making consists of two stages:

- first, all participants agree on an appropriate "social norm", i.e., on an appropriate set of imputations C within which they search for an imputation;

- then, once the set C is selected, the participants select an imputation x from this set.

The two conditions from Definition 4 guarantee that:

- if someone tries to violate an agreement and propose an imputation y outside the set C corresponding to the social norm, then we can force it back into C;
- second, that once an imputation x is selected, no coalition is interested in switching to a different socially acceptable imputation y.

Our goal is to compute the list of all possible social norms C (or at least compute one social norm C).

Computing von Neumann-Morgenstern solution is a challenge. Originally, von Neumann and Morgenstern proposed their solution as the main solution concept for cooperative games. Unfortunately, the two major challenges emerged. A minor challenge is that there are games which do not have this solution at all. A major challenge is that computing this solution is not easy. As of now, it is not even clear whether there exists an algorithm that can compute such a solution [2, 5].

In the discrete case, it is known that the problem of checking the existence of an NM-solution is NP-hard. Indeed, we can represent the conflict situation as a graph, with outcomes as vertices and $x \succ y$ if and only if there is an edge from x to y. In graph terms, an NM-solution is a *minimum independent dominating set*, or a *kernel*. The problem of checking the existence of such a kernel is NP-complete. This result was first proven in [1]; see also [4] (Problem 9.5.10).

Comment. For a general overview of complexity of different conflict resolution notions, see, e.g., [5] and references therein.

Our idea. A real-life division of the overall sum may involve not just division of money, but rather a division of objects whose price is also only approximately known. Thus, the actual value x_i allocated to each person is also only approximately known. Let us introduce the following definition.

Definition 5. *Let* $\varepsilon > 0$.

- *We say that tuples* x *and* x' *are* ε-close *if* $|x_i - x'_i| \le \varepsilon$ *for all* i *and* $\left|\sum\limits_{i\in S} x_i - \sum\limits_{i\in S} x'_i\right| \le \varepsilon$ *for all coalitions* S.

- *We say that two sets of tuples X and X' are* ε-close *if every $x \in X$ is ε-close to some tuple from X', and every $x' \in X'$ is ε-close to some tuple from X.*

When we only know ε-approximations $\widetilde{x}$ and $\widetilde{y}$ to numbers x and y, then we cannot check whether $x > y$. To be more precise, if $\widetilde{x} > \widetilde{y} + 2\varepsilon$, then $x \geq \widetilde{x} - \varepsilon > \widetilde{y} + \varepsilon \geq y$ hence $x > y$. Similarly, if $\widetilde{x} \leq \widetilde{y} - 2\varepsilon$, then we are sure that $x \leq y$ and thus, that $x \not> y$. However, if $\widetilde{y} - 2\varepsilon < \widetilde{x} \leq \widetilde{y} + 2\varepsilon$, then we can have both $x > y$ and $x \not> y$. So, we can distinguish between "necessarily larger" relation $\widetilde{x} > \widetilde{y} + 2\varepsilon$ and "possibly larger" relation $\widetilde{x} > \widetilde{y} - 2\varepsilon$.

When we request that no two imputations from the set C dominate each other, this cannot mean "possibly larger", since even for $x = y$, the value y_i is possible larger than x_i. Thus, we need to require "necessarily larger" condition. On the other hand, when we require that every imputation not from C can be forced into C, we cannot use the "necessarily larger" condition, since this way we may eliminate some "forcings". Thus, we arrive at the following definitions.

Definition 8.

- *We say that a tuple x* necessarily ε-dominates *y, and denote it by $x \succ^{\varepsilon} y$ if $x_i > y_i + 2\varepsilon$ for all $i \in S$ and $\sum\limits_{i \in S} x_i \leq v(S) - \varepsilon$.*
- *We say that a tuple x* possibly ε-dominates *y, and denote it by $x \succ_{\varepsilon} y$ if $x_i > y_i - 2\varepsilon$ for all $i \in S$ and $\sum\limits_{i \in S} x_i \leq v(S) + \varepsilon$.*
- *Let X be a set of tuples. We say that a set C is an ε-NM solution for the set X if the following two conditions are satisfied:*
  - *if $x, y \in C$, then $x \not\succ^{\varepsilon} y$;*
  - *if $y \in X - C$, then there exists an $x \in C$ for which $x \succ_{\varepsilon} y$.*

This modification of the original NM-definition takes into account that we only know the values $v(S)$ and x_i only approximately. Such "approximate" NM-solutions can be algorithmically computed in the following sense:

Theorem. *There exists an algorithm that, given a rational-valued game v and positive rational numbers $\varepsilon > \delta > 0$, returns a finite list of finite sets $C_1, \ldots, C_M$ such that:*

- *each of the sets C_i is an $(\varepsilon+\delta)$-NM solution, and*
- *each $(\varepsilon-\delta)$-NM solution C is δ-close to one of the sets C_i.*

Proof. Let us use a grid with step δ to form a finite δ-approximation X' to the set X of all imputations. For any subset C of this finite set X', we can algorithmically check whether this set is an ε-NM solution for the set X'. Let us show that the list of all such solution is the desired list $C_1, \ldots$

Indeed, one can easily check that if $x \succ^\alpha y$, x' is δ-close to x, and y' is δ-close to y, then $x' \succ^{\alpha-\delta} y'$. Thus, if $x' \not\succ^{\alpha-\delta} y'$, then $x \not\succ^\alpha y$. In particular, if $x' \not\succ^\alpha y'$, then $x \not\succ^{\alpha+\delta} y$.

Similarly, if $x \succ_\alpha y$, x' is δ-close to x, and y' is δ-close to y, then $y' \succ_{\alpha+\delta} x'$.

Thus, if C is an α-NM solution for X, C' is δ-close to C, and X' is δ-close to X, then C' is $(\alpha+\delta)$-NM solution for X'.

Since each set C_i is an ε-NM solution for X', and X' is δ-close to X, we conclude that C_i is an $(\varepsilon+\delta)$-NM solution for the set X.

Vice versa, let C be an $(\varepsilon-\delta)$-NM solution for the set X. By construction of the set X', each point from the set X, in particular, each point from C, is δ-close to some point from X'. Let C' denote the set of all the points from X' which are δ-close to some point from C. By definition, this set C' is δ-close to C. Since X' is δ-close to X, we thus conclude that $C' \subseteq X'$ is an ε-NM solution for X' and is, thus, one of the sets C_i. The theorem is proven.

Acknowledgments. This work was supported in part by the National Science Foundation grants HRD-0734825, HRD-1242122, and DUE-0926721.

Литература.

1. *Chvatal, V.* On the computational complexity of finding a kernel // Center de Recherches Mathématiques, Université de Montréal, Technical Report CRM-300, 1973.

2. *Deng, X., and Papadimitriou, C.H.* On the complexity of cooperative solution concepts // Mathematics of Operations Research, 1994, V. 19, No. 2, P. 257–266.

3. *Luce, R.D., and Raiffa, H.* Games and Decisions: Introduction and Critical Survey – New York: Dover, 1989.

4. *Papadimitriou, C.H.* Computational Complexity – San Diego: Addison Wesley, 1994.

5. *Papadimitriou, C.H.* Algorithms, games, and the Internet // Proceedings of ACM Symposium on Theory of Computing STOC'01, Hersonissos, Crete, Greece, July 6–8, 2001.

6. *von Neumann, J., and Morgenstern, O.* Theory of Games and Economic Behavior – Princeton, New Jersey: Princeton University Press, 1944.