Can Trust Increase the Efficiency of Cake Cutting Algorithms?

(Extended Abstract)

Roie Zivan,
Industrial Engineering and Management department,
Ben Gurion University of the Negev,
Beer-Sheva, Israel
{zivanr}@bgu.ac.il

ABSTRACT

Fair division methods offer guarantees to agents of the proportional size or quality of their share in a division of a resource (cake). These guarantees come with a price. Standard fair division methods (or "cake cutting" algorithms) do not find efficient allocations (not Pareto optimal). The lack of efficiency of these methods makes them less attractive for solving multi-agent resource and task allocation. Previous attempts to increase the efficiency of cake cutting algorithms for two agents resulted in asymmetric methods that were limited in their ability to find allocations in which both agents receive more than their proportional share.

Trust can be the foundation on which agents exchange information and enable the exploration of allocations that are beneficial for both sides. On the other hand, the willingness of agents to put themselves in a vulnerable position due to their trust in others, results in loss of the fairness guarantees that motivate the design of fair division methods.

In this work we extend the study on fair and efficient cake cutting algorithms by proposing a new notion of *trust-based efficiency*, which formulates a relation between the level of trust between agents and the efficiency of the allocation. Furthermore, we propose a method for finding trust-based efficiency. The proposed method offers a balance between the guarantees that fair division methods offer to agents and the efficiency that can be achieved by exposing themselves to the actions of other agents. When the level of trust is the highest, the allocation produced by the method is globally optimal (social welfare).

Keywords

Game Theory, Social choice theory

1. INTRODUCTION

One of the main challenges in multi-agent systems (MAS) is encouraging self-interested agents to cooperate. Fair division methods offer a possible solution to this challenge for resource and task allocation, by offering guarantees to agents on the quality or size of their share, as long as they are cooperative (follow the instructions of the method's protocol). Moreover, these guarantees hold for an agent, even if other agents choose an uncooperative strategy.

A fair division method guarantees fairness properties but may be inefficient (not Pareto optimal). In other words, there can exist a different allocation that is preferred by both (or preferred by one and is equal in the eyes of the other).

Previous attempts to introduce efficiency into a fair division method offered asymmetric extensions of *Austin's method* [1, 3]. These methods have the following limitations: (1) Only allocations that include up to two cuts of the cake are considered. (2) The method does not consider allocations in which both agents value their share as more than 50%.

The possibility of finding solutions to negotiation problems that *expand the pie*, i.e., the sum of the benefit for the negotiating parties exceeds 100%, was acknowledged by social scientists and triggered studies that investigated the success of different strategies in producing such agreements. Intuitively, integrative strategies that increase the cooperation and information exchange between the negotiating parties increase the chances for efficient agreements.

Trust is a concept that has been intensively studied by social scientists and by the multi agent systems community. The common and accepted definition for trust is the willingness of an agent to put herself in a situation in which she is vulnerable to the actions of another (the party she *trusts*). The relation between trust and efficiency was also acknowledged by multi-agent system studies.

In this paper we extend the research on fair and efficient cake cutting methods by:

1. Proposing a new notion of *trust based efficiency*. It defines the level of efficiency that can be achieved as a function of the level of trust among the agents.
2. Proposing a method for finding *trust based efficiency* that is independent of the role of the agents. The method proposed allows agents to expose themselves with respect to the level of trust and make use of this exposure to increase efficiency while maintaining the guarantees on the fraction of the proportional share that the agents were not willing to risk. When the level of trust is maximal, the allocation found by the method is globally optimal (social welfare).

2. AUSTIN'S METHOD AND ASYMMETRIC EXTENSIONS

Austin's moving knife procedure is famous for being the only method that can find a division of a cake between two agents such that both agents value their share as exactly half of the cake (*exact allocation*) [1, 2].

In Austin's procedure, an infinitely divisible but bounded resource (cake) X is divided between two agents, a and b.

Cite as: Can Trust Increase the Efficiency of Cake Cutting Algorithms? (Extended Abstract), Roie Zivan, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. XXX-XXX.

We assume that the cake has a rectangular shape with length L and width 1. We further assume that all cuts are planar. Each agent has its own utility function, U_a and U_b respectively, which defines the utility she derives from an allocation of any piece of the cake to her. One agent (a) holds two parallel knives. In the initial state, the left knife is placed at the left edge of the cake and the right knife is placed so that the utility she derives from the piece between the knives (we will refer to the piece between the knives as P and to the remainder of the cake as $\bar{P}$) would be $U_a(P) = \frac{1}{2}U_a(X)$ (for simplicity we will assume that $U_a(X) = U_b(X) = 1$). Agent a then moves both knives to the right so that at all times $U_a(P) = \frac{1}{2}$. When $U_b(P) = \frac{1}{2}$ as well, agent b calls "stop" and is allocated P while a gets $\bar{P}$. Thus, the utilities derived by both agents from their share are $U_a(\bar{P}) = U_b(P) = \frac{1}{2}$ (an *exact division* [2]).

If we allow agent b to observe the full process in which agent a moves the knives from the initial position to the final complementary position, and then choose the piece that she values the most and that was between the knives at some point during the process, we can increase the efficiency of the method. However, it is clear that this increment in efficiency is one-sided ($U_b \geq \frac{1}{2}$ while $U_a = \frac{1}{2}$).

A different extension to Austin's method, which increases its efficiency, was proposed by Sen and Biswas [3]. Their method reaches a similar result by allowing the cutting agent (a) to hold a model of the other agent preferences. This allows her to manipulate the selection of b and be left with the most beneficial allocation among the allocations that leave agent b with a satisfactory consecutive share.

The two methods described above are both asymmetric, i.e. give an advantage to one of the agents over the other. Both methods do not consider allocations that increase the benefit for both agents beyond their proportional share.

3. TRUST BASED EFFICIENCY

An allocation A will be constructed of two disjoint sets of pieces, X_a and X_b. If we will put together all the pieces in X_a and X_b we will get the entire cake (X). We will use the notation $U_j(x)$ for the utility agent j derives from the allocation of piece x to her. The utility agent j derives from an allocation A will be denoted $U_j(A)$ and will be equal to $U_j(X_j)$, the utility the agent derives from her allocated set of pieces in A, X_j. Once again, for simplicity we will assume that agents' utility functions are normalized, i.e., $U_j(X) = 1$. We propose the following two innovative notions:

1. given $0 \leq l \leq 1$, the symmetric level of trust between agents a and b, an incentive participation constraint for agent $j \in \{a, b\}$ is that for any possible resulting allocation A, $U_j(A) \geq \frac{1-l}{2}$.

2. An allocation A is l trust efficient if there is no piece x held by agent $j \in \{a, b\}$ in A and piece x' held in A by agent $i \in \{a, b\}, i \neq j$ for which: (a) $U_i(x) \geq \frac{1-l}{2}$. (b) $U_j(x') \geq \frac{1-l}{2}$. (c) $U_j(x') > U_j(x)$. (d) $U_i(x) \geq U_i(x')$.

The following method finds l-trust-efficient (LTE) allocations of a cake between two agents:

1. At the initial state, agent a places the left knife on the left edge of the cake and the right knife so that $U_a(P) = \frac{1-l}{2}$ (recall that P is the piece between the knives).

2. Agent a moves the knives to the right, keeping the value of P at $\frac{1-l}{2}$ until at the final state, the right knife reaches the right edge of the cake.

3. Agent b decides which part of the cake to allocate to agent a and which part to herself, cuts the cake and makes the allocation accordingly.

To complete the description of the mechanism, it remains to describe the protocol that agent b follows in the third step. Notice that like in Austin's procedure, while the value of P for agent a remains the same while the knives are moving, its value for agent b may be changing. The value of the piece P for agent b as a function of the location of the left knife (moving to the right between the left edge of the cake and its location in the final state) is observed and analyzed by her in order to produce the allocation.

Agent b selects a set of disjoint pieces X_a to allocate to agent a so that the following conditions are satisfied: (1) $x \in X_a \Rightarrow x$ was equal to P at some time through the movement of the knives. (2) $x \in X_a \Rightarrow U_b(x) \leq \frac{1-l}{2}$, i.e. b values x less than agent a values it. (3) $x \in X_b \Rightarrow U_b(x) > \frac{1-l}{2}$. (4) $U_b(X_b) \geq \frac{1-l}{2}$. (5) $X_a \neq \emptyset$.

If these conditions cannot be satisfied (for example if $U_a = U_b$ the third condition cannot be satisfied), then agent b selects a piece P', which was between the knives at some point during the process and has a lower value in her eyes than any other such piece P, and allocates P' to a, leaving the rest of the cake for herself.

A number of properties can be established for the method presented above. Among them the two properties that the method was designed to achieve, that it finds an l-trust-efficient allocation and that the guarantees for agents are maintained, i.e., for any allocation A found by the method, $U_a(A) \geq \frac{1-l}{2}$ and $U_b(A) \geq \frac{1-l}{2}$. In addition its equivalence to the asymmetric version of Austin's method when the level of trust is minimal and its convergence to a globally optimal social welfare allocation when the level of trust is maximal can be established as well (proofs for these properties were omitted for lack of space).

4. CONCLUSION

We proposed the use of *trust in cake cutting algorithms*. We defined the level of trust between agents as the proportional quantity of their fair share that they are willing to expose to the actions of other agents, and risk losing. We further defined a new concept, l-trust-efficiency, which determines the level of efficiency of an allocation based on the level of trust between the agents.

We proposed a method for finding l-trust-efficient allocations. The method allows agents to increase the efficiency of the allocation with respect to the level of trust between them, but at the same time, guarantees the allocation of the quantity that they were not willing to risk. The method allows the agents to divide the cake between them with respect to the utility they derive from allocations of the different parts of the cake and, as a result, increase not only the efficiency but also the social welfare value of the allocation.

5. REFERENCES

[1] A. K. Austin. Sharing a cake. *Math. Gazett*, 66:212–215, 1982.

[2] Y. J. Robertson and W. Webb. *Cake-Cutting Algorithms, Be Fair If You Can.* A K Peters, Ltd, 1998.

[3] S. Sen and A. Biswas. More than envy-free. In *ICMAS '00: Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS-2000)*, page 433, Washington, DC, USA, 2000. IEEE Computer Society.