Negotiating efficient envy-free divisions

Sandip Sen and Rajatish Mukherjee
Department of Mathematical & Computer Sciences
University of Tulsa
e-mail: rajatish@euler.mcs.utulsa.edu, sandip@kolkata.mcs.utulsa.edu

Abstract

Division of a resource among multiple agents is a frequent problem in multiagent systems and fair, efficient, and decentralized allocation procedures are highly valued. A division of a resource or good is envy-free when every agent believes that its share is not less than anyone else's share by its own estimate. As envy-free procedures are not efficient (in the sense of Pareto optimality) we have previously worked on improving the efficiency of such envy-free division procedures among two agents using models of other agents' utility functions. In this paper, we extend that work by devising an anytime algorithm that increases the efficiency of the envy-free allocation. The procedure also has the desired property of envy-freeness.

Introduction

A key research issue in agents and multiagent research is to develop negotiation procedures by which agents can efficiently and effectively negotiate solutions to their conflicts (Rosenschein & Zlotkin 1994). In this paper, we focus on the problem of agents vying for portions of a good. The negotiation process will produce a partition and allocation of the goods among the agents (Huhns & Malhotra 1999; Robertson & Webb 1998). We are interested in both protocols by which agents interact and appropriate decision procedures to adopt given a particular procedure.

In an *envy-free* division, each agent believes that it received, by its own estimate, at least as much as the share received by any other agent (Brams & Taylor 1996; Stewart 1999). This also implies that an agent has no incentive to trade its share with anyone else. While such guarantees are indeed extremely useful to bring parties to the discussion table, there is no a priori reason why a self-interested agent should be happy with just the share that is most valuable to its estimate in the group when it can possibly get even more. For example, a procedure by which an agent can possibly improve on its share received by an envy-free procedure without losing the guarantee of envy-freeness, would be of much value.

We assume that the good being divided is possibly heterogeneous and the preference of an agent for various parts of the good is represented by a utility function[1]. For historical reasons, we will represent the good as a continuously divisible rectangular piece of cake, i.e., the cake may be cut at any point and can be cut any number of times. We are only interested in the length of the cake.

Envy-free procedures produce allocations that guarantee that each participating agent considers itself a winner. A self-interested agent, however, may not be contended with the most valuable share (according to its own estimate) that it can get in the group. Put simply, a rational agent wants to maximize its utility, and if there is scope for cornering a larger share of the good being divided, even an envy-free procedure may not be satisficing! Envy-free procedures are also not guaranteed to produce efficient, viz., Pareto-optimal, divisions. This means that it is possible to further re-allocate portions of the cake so that the utility of at least one of the agents is improved without decreasing the utilities of the other agents.

To illustrate this scenario, consider the "divide and choose" procedure, in which one agent cuts the cake into two portions and the other agent gets to choose the portion it wants for itself. The strategy of the cutting agent would be to divide the cake into two portions of equal value by its own estimate. The choosing agent should then choose the portion that is of more value by its own estimate. Thus, both agents would believe they did get the most valuable portion of the cake, and would therefore be envy-free. But there may be portions of the cake in their respective allocations that they can still exchange and further improve their valuations. This problem is only exacerbated with larger number of agents.

In our previous work (Sen & Biswas 2000), we presented an approach for two-agent divisions by which agents can use the model of the utility function of the other agent to increase the efficiency of the resultant

[1]In this paper the good being divided represents a resource of interest to participating agents. We refer to the resource as a good for consistency with literature in envy-free division.

division. The division was guaranteed to be envy-free and was proven to be optimal when such an allocation involved the assignment of a contiguous portion of the cake to one agent. We will refer to this procedure as *2e-opt*.

In the present work, we will focus on the problem of improving the efficiency of envy-free divisions among two agents by recursively dividing, using *2e-opt*, the initial partitions produced. While our previous work shows that there does not exist a guarantee optimal envy-free division in general, the procedure presented here can improve the efficiency of divisions that agents can negotiate by using 2e-opt. The procedure also has the desired property of improving the efficiency of the division monotonically with time.

To set the context for this work, we will summarize the requisites of fair division processes, the framework for negotiation, and the algorithm *2e-opt* from our previous work. We will then present the recursive algorithm that is the contribution of this paper.

Division of a good

We concern ourselves with the problem of dividing up a good between multiple individuals. We will assume that the solution procedure is of a decentralized nature. This means that the agents will be required only to abide by a protocol by which the division is to be made. They can freely choose any strategy to use to determine their actions within the accepted protocol.

For example, a protocol used in auction settings may require that every agent submit a sealed bid for the good to an auctioneer. After every bid is collected, the good is divided up among the bidders in proportion to their bids. Our assumption is that once the agents agree to such a protocol, they are free to choose their bids following any strategies they adopt. We require, however, that agents will agree to the division of the good as specified once they have placed their bids.

From a designer's point of view, the choice of a protocol provides a platform for agents to negotiate an agreeable division. The choice of a strategy will be dictated by concerns for arriving at a preferred share of the good being divided. The protocol designer should then provide protocols which can be used by agents to successfully negotiate agreeable divisions with reasonable computational costs. We now list properties of divisions that can make them agreeable to self-interested agents.

Desired characteristics of divisions

We assume that a single divisible good is to be divided among n agents. The following criteria have been espoused as desirable characteristics of decision procedures or outcomes from such procedures (Brams & Taylor 1996):

Proportional: Each agent believes that it received, by its own estimate, at least $\frac{1}{n}$ of the goods being allocated.

Envy-free: Each agent believes that it received, by its own estimate, at least as valuable a share as that received by any other agent. This also implies that an agent has no incentive to trade its share with anyone else.

Equitable: A solution, i.e., a partition of the good among the n agents, is equitable, when the share received by each agent is identical in terms of their individual utility functions.

Efficient: A solution is said to be Pareto optimal or efficient if there is no other partition which will improve the perceived share of at least one agent without decreasing the perceived share of any other agent.

Envy-free divisions with improved efficiency

We assume that the continuously divisible good is possibly heterogeneous and the preference of an agent for various parts of the good is represented by a utility function. For historical reasons, we will represent the good as a rectangular piece of cake. For all practical purposes, however, we are only interested in the length of the cake. Though we present the *2e-opt* negotiation procedure from our previous work (Sen & Biswas 2000) here for completeness, at first glance the reader can simply use the results from the following theorems to facilitate understanding of the discussion about the extended protocol in the later sections of this paper.

The utility to the ith agent of a piece of the cake cut between points a and b (where $a < b$) is given by $\int_a^b U_i(x)dx$, where $U_i(x)$ is the utility function of the ith agent. Without loss of generality, we will assume that the first agent is having a model, $\overline{U_2}$, of the utility function of the second agent. This model need not be accurate, but the better the model, the more benefit the modeling agent stands to gain by using it.

The *2e-opt* procedure we now describe is derived from Austin's moving knife procedure (Brams & Taylor 1996). In the *2e-opt* procedure (Sen & Biswas 2000), the modeling agent A hold two knives parallel to the side edges of the cake and then move them to the right allowing for wrap around. Let the positions of the left and right knives at time t be l_t and r_t respectively. l_0 and r_0 determine the initial region offered to B. The knives stop at time $t = T$ when the right knife reaches the original position of the left knife, and the left knife reaches the original position of the right knife. The agent B then chooses a time $\tau \leq T$, and the portion of the cake in between l_τ and r_τ (with wrap-around if needed) is given to B, with the rest of the cake going to agent A (see Figure 1).

It can be shown that B can always negotiate an envy-free division for itself if it calls "cut" at an appropriate time irrespective of how A moves the knife. B, however, can be resentful as it can presume that the advantage of moving the knife allows A to obtain a super-equitable share for itself, which guarantees a sub-equitable share for B (Sen & Dutta 2001).

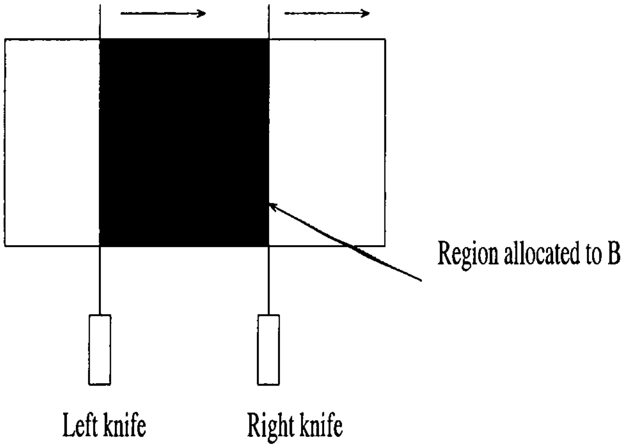

Figure 1: Our augmentation of Austin's procedure with the modeling agent moving two knives.

We now specify the choice of the initial region, the choice of the target region that the modeler wants the other to choose, and a way of moving the knife such that this target region is most likely to be chosen by the other agent.

Initial region selection: The initial knife placements should be such that $\int_n^m U_A(x)dx = \frac{1}{2}\int_0^L U_A(x)dx$, where L is the length of the cake. It is also desirable that the region in between does divide the cake in half for agent B. Otherwise B may find it beneficial to choose $t = 0$ (if, according to B's estimate, the region (n, m) is the most valuable) or $t = T$ (if, according to B's estimate, the region (m, n) is the most valuable). But a contiguous region that satisfy both these conditions may not be found in general. So, the initial placement (m, n) should also satisfy the following condition:

$$(m, n) = \arg\min_{y,x} \left| \int_x^y \overline{U_B}(z)dz - \frac{1}{2}\int_0^L \overline{U_B}(z)dz \right|.$$

For ease of exposition, if $y < x$, we use $\int_x^y f(z)dz$ to represent $\int_x^L f(x)dx + \int_0^y f(z)dz$.

Target region selection: The target region must result in an envy-free division. Eliminating beginning and end points, the target region should be selected from the following set: $Z = \{(x,y) : (\int_x^y \overline{U_B}(z)dz > max(\int_{l_0}^{r_0} \overline{U_B}(x)dx, \int_{l_T}^{r_T} \overline{U_B}(x)dx)) \wedge (\int_x^y U_A(z)dz < \frac{1}{2}\int_0^L U_A(z)dz)\}$.

From the regions in Z, we filter out those regions that are of least value to A: $Y = \arg\min_{(x,y)\in Z} \int_x^y U_A(z)dz$. From the regions in Y, A can select those regions which are estimated to be of most value to B: $X = \arg\max_{(x,y)\in Y} \int_x^y \overline{U_B}(z)dz$. This maximizing choice can be viewed either as a cooperative gesture or as an attempt to improve the likelihood of B accepting what A would prefer it to accept. If X is non-empty, then A's goal is to select one of its elements and then move the knifes such that the corresponding region is the most attractive offer received by B. If X is empty, however, A will have to be satisfied with half of the cake.

Moving the knife: While moving between pairs of points as identified in the previous paragraph, the knife locations (u, v) should satisfy the following inequality: $\int_u^v \overline{U_B}(x)dx < \frac{1}{2}\int_0^L \overline{U_B}(x)dx$. For regions where A's utility is high, the spacing between the knives will be reduced to make that region non-envy-free for B.

An instance of the above procedure is shown in Figure 2. We present some properties of this decision procedure (for proof see (Sen & Biswas 2000)):

Theorem 1 *Our scheme dominates Austin's procedure with respect to efficiency of allocations.*

Theorem 2 *If a Pareto-optimal allocation for a problem involves a contiguous region, our proposed scheme will select it when given an accurate agent model.*

Recursive division to improve optimality

The protocol *2e-opt* guarantees an envy-free division among two agents that is more optimal than Austin's moving knife procedure (Austin 1982). The division obtained after applying *2e-opt* will be Pareto-optimal only if there exists some contiguous Pareto-optimal allocation in the good. In a large majority of cases, however, such a contiguous Pareto-optimal allocation may not exist. Under such circumstances, we propose an augmented procedure to increase the efficiency of the allocation without affecting the envy-free guarantee. We now present a time-constrained procedure (with time-limit T) that recursively invokes itself, checking at each stage if further improvements are feasible. The input to this procedure consists of the envy-free division of the cake into two portions (portion X for agent A and portion Y for agent B) obtained after applying *2e-opt* or any other envy-free division procedure. The algorithm is presented in Figure 3.

The above algorithm combines and encapsulates both the specification of a protocol, i.e., when agents communicate and the roles played by each agent in the division process, and the strategy adopted for negotiation, i.e., what portions are offered by one agent and which portion is selected by the other. This negotiation procedure is decentralized as each agent is in control of what to offer and accept. Such a procedure can be used both in a cooperative and competitive scenario. In the competitive scenario, the modeler will be interested only in maximizing its own share of the good. In a cooperative scenario agents can negotiate more efficient and equitable divisions by truthfully revealing their preferences or utilities (incorrect utility estimates can limit the performance of *2e-opt* and *2e-optR* but they will still produce envy-free divisions).

We illustrate this process by considering two agents A and B. Let us assume that an envy-free division has produced an allocation of portion X to agent A

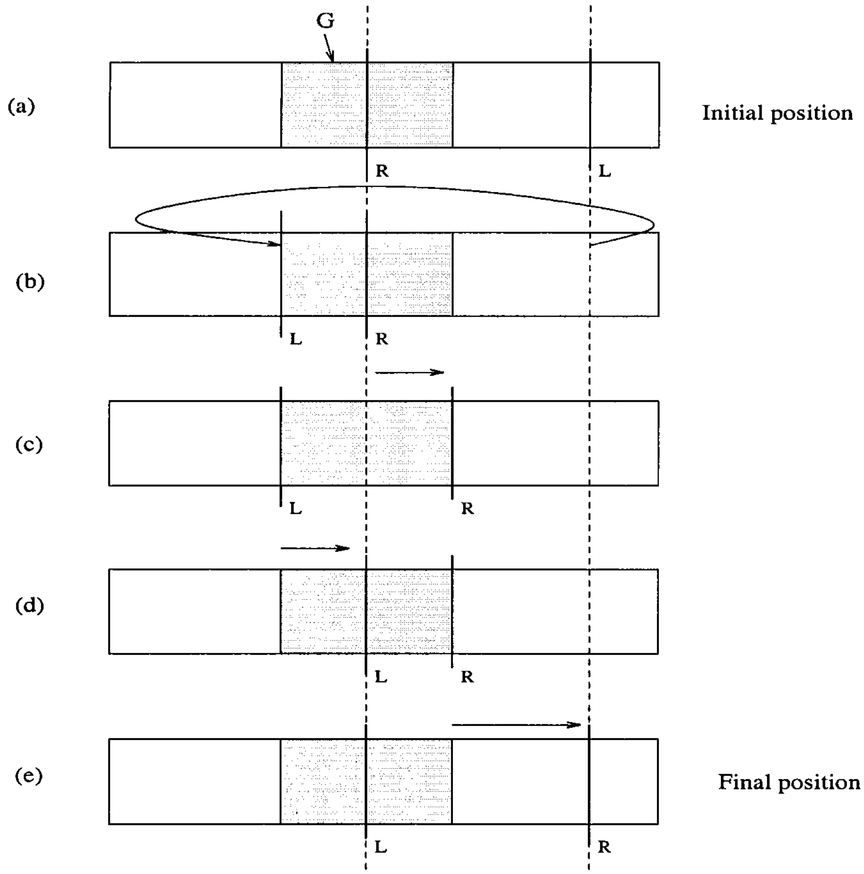

Figure 2: Knife-moving procedure snapshots (G: target region; L,R: left, right knife).

and portion Y to agent B. Now, *2e-opt* is applied on both X and Y dividing it into portions (X_A, Y_B) and (Y_A, Y_B) respectively. This means that if X were to be the only portion to be divided up between A and B, then (X_A, X_B) would be an envy-free division with A receiving X_A and B receiving X_B, i.e., X_A (X_B) is the more valuable portion of X in A's (B's) estimate. Note that (X_A, X_B) is not the optimal envy-free division possible, but the one that is computed by *2e-opt* . Similar reasoning holds for the portion Y. There exists a plausible exchange between A and B if A believes Y_A is more valuable than X_B, and B believes that X_B is more valuable than Y_A. This process is illustrated in Figure 4. We have illustrated only a single step of our algorithm. After the exchange is complete, the algorithm recursively calls *2e-optR* over the new allocations.

The above algorithm will return an envy-free allocation at least as efficient as the *2e-opt* procedure. Also, if A and B have non-conflicting areas of interest within the cake, the above algorithm guarantees a more efficient solution than that achieved by Austin's two person envy-free allocation procedure. Our algorithm *2e-optR* terminates if A and B decide not to exchange any portion at a particular iteration or if the running time has exceeded T.

Figure 4 illustrates a possible iteration of our efficient envy-free division algorithm. Envy-free allocation for A and B (at the start of the algorithm) is shown. Given that A likes Y_A more than X_B and similarly B likes X_B more than Y_A, we can exchange these portions to achieve more efficiency.

Discussion

In our prior work we presented an algorithm that produced optimal envy-free divisions of continuously divisible goods between two agents under certain assumptions of utility functions (Sen & Biswas 2000). In this work, we have introduced a recursive extension (of course, an equivalent iterative version would do just as well) of the previous protocol that results in an anytime algorithm which monotonically increases the optimality of envy-free divisions with available computation time.

It is quite likely that such an extension would follow the law of diminishing returns. That implies that

```
Procedure 2e-optR (X,Y,T)
// T is the maximum running time; X, Y are initial allocations
{
      if (T=0) // if time has elapsed
         return(X,Y); // return new allocations
      else
         Divide portion X into portion {X_A,X_B} using 2e-opt;
         Divide portion Y into portion {Y_A,Y_B} using 2e-opt;
         if (Utility_A(X_B) <= Utility_A(Y_A) and
             Utility_B(Y_A) <= Utility_B(X_B))
         {// if envy-free exchange possible, exchange X_B with Y_A
            X <- X_A U Y_A;
            Y <- X_B U Y_B;
            return 2e-optR(X,Y,(T-C));// Time taken for this call = C
         }
         else // no more exchange possible
           return(X,Y);
}
```

Figure 3: The 2e-optR negotiation procedure.

the earlier exchanges are likely to produce greater increase in efficiency compared to latter exchanges. The computation costs of each negotiation stage should be roughly the same. This raises the issue of a tradeoff between computation cost and efficiency improvement. If the negotiation is done off-line or with a fixed time limit as we have assumed, the termination criteria should be as used here. If the agents are interested in factoring in computation costs, however, negotiation can possibly be terminated when the expected gain from the next negotiation stage is less than the computational cost of participating in that stage.

We are working on extending our envy-free division method to more than two agents. In the general case of n agents, no envy-free division procedure exists. Hence, we would be interested in augmenting approximately envy-free division procedures, where allocations are envy-free within some pre-specified error-bounds (Brams & Taylor 1996). The *2e-optR* algorithm, can be utilized to improve the optimality of an already existing envy-free division among n agents. For example, for $n = 3$, an envy-free allocation divides the cake into three portions. We can apply the *2e-optR* algorithm for two persons for every pair of agents. We allow exchange between two agents if both agree and the third agent does not object, i.e., the resultant allocation after the exchange does not make the third agent envious. The problem is similar to the multi-agent contract problem which consists of exchanges between all agents (Andersson & Sandholm 1999). However, our problem is more difficult because multi-agent contracting is concerned only with individual rationality (no agent will enter into a contract that leaves it worse-off), whereas in this case we have to additionally ensure that all agents remain envy-free after a possible exchange between any two agents.

We have presented a procedure that improves the efficiency of n agent envy-free divisions (Nuchia & Sen 2001). The difference of this work from the approach above is that we assume an initial envy-free allocation and try to improve on that allocation in efficiency without sacrificing the envy-free property.

Acknowledgments This work has been supported in part by an NSF CAREER award IIS-9702672.

References

Andersson, M. R., and Sandholm, T. W. 1999. Time-quality tradeoffs in reallocative negotiation with combinatorial contract types. In *Sixteenth National Conference on Artificial Intelligence*, 3–10. Menlo Park, CA: AAAI Press/MIT Press.

Austin, A. 1982. Sharing a cake. *Mathematical Gazette* 66(437):212–15.

Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From cake-cutting to dispute resolution.* Cambridge University Press, Cambridge: UK.

Huhns, M. N., and Malhotra, A. K. 1999. Negotiating for goods and services. *IEEE Internet Computing* 3(4).

Nuchia, S. W., and Sen, S. 2001. Improving optimality of n agent envy-free divisions. In *Pre-Proceedings of the Eighth International Workshop on AGENT THEORIES, ARCHITECTURE, AND LANGUAGES*, 208–217.

Robertson, J., and Webb, W. 1998. *Cake Cutting Algorithms.* Natick, MA: A. K. Peters.

Rosenschein, J. S., and Zlotkin, G. 1994. Designing conventions for automated negotiation. *AI Magazine* 29–46.

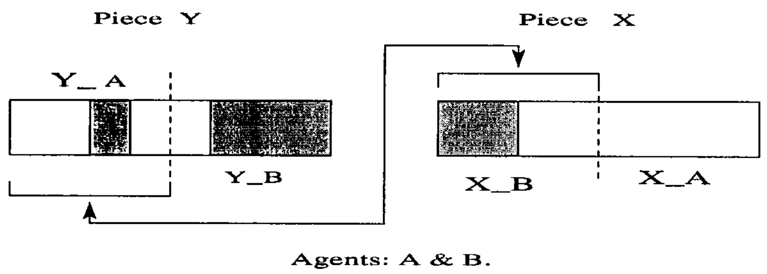

Figure 4: A plausible envy-free exchange between agents A and B.

Sen, S., and Biswas, A. 2000. More than envy-free. In *Proceedings of Fourth International Conference on MultiAgent Systems*, 433–434. Los Alamitos, CA: IEEE Computer Society.

Sen, S., and Dutta, P. S. 2001. Envy and spite free divisions. In *Working Notes of the Fourth Workshop on Deception, Fraud and Trust in Agent Societies*, 71–90.

Stewart, I. 1999. Mathematical recreations: Division without envy. *Scientific American*.