Using agent model to find efficient envy-free allocation of a continuously divisible good between two agents

Sandip Sen

Department of Computer Science
University of Tulsa
e-mail:sandip@utulsa.edu

ABSTRACT

Division of a resource among multiple agents is a critical 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. Envy-free procedures are not efficient, i.e., they are not Pareto optimal in general. We provide a procedure for improving the efficiency of an envy-free division of a continuously divisible good among two agents using models of agents' utility function. We further provide an anytime recursive algorithm that increases, if possible, the efficiency of an initial envy-free allocation. Though efficient envy-free division is not possible for general utility functions, we show that our procedure produces optimal, envy-free divisions for certain classes of utility functions.

Keywords: Envy-free division, efficient outcome, cake-cutting
1998 ACM CCS: I.2.11:Artificial Intelligence-Distributed Artificial Intelligence-Multiagent systems

1 Introduction

A key research issue in agents and multiagent research is to develop negotiation procedures by which agents can effi ciently and effectively negotiate solutions to their conflicts (Jennings, Faratin, Lomuscio, Parsons, Sierra and Wooldridge, 2001; Kraus, 2001; Rosenschein and Zlotkin, 1994). In this paper, we focus on the problem of agents vying for portions of a good which represents a continuously divisible resource. The negotiation process will produce a partition and allocation of the good among the agents (Huhns and Malhotra, 1999; Robertson and 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 and 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 just the largest share in the group. Hence, it will value procedures following which it might get even larger shares without losing the guarantee of envy-freeness.

Envy-free procedures are 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. Though there are results on efficient envy-free divisions of indivisible or atomic objects (Haake, Raith and Su, 2002), our focus is on division of continuously divisible resources.

We assume that the good being divided is continuously divisible, possibly heterogeneous and the preference of an agent for various parts of the good is represented by a utility function[1]. In literature, a continuously divisible good is often represented as a linear piece of cake, i.e., the cake may be cut at any point and can be cut any number of times.

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 this paper, we first present a moving-knife procedure from literature for two agents that produces envy-free divisions. We then present an augmentation of that procedure to improve the efficiency of the resultant division using a model of the utility function of the agents. The division is guaranteed to be envy-free, and is shown to be optimal for certain classes of utility functions of the agents. We will refer to this procedure as *2e-opt*.

Next, we improve the efficiency of the envy-free divisions among two agents by recursively dividing, using *2e-opt*, the initial partitions produced. Though we show that there does not exist a finite procedure for guaranteeing optimal envy-free division in general, the recursive procedure presented here, *2e-optR*, 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, if feasible, over time.

2 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 (Krishna, 2002). 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

[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 on envy-free division.

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.

2.1 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 and 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.

Here are some observations on the above-mentioned metrics:

- If all agents have identical perception or utility function, any partition is efficient. Also, in this case the only proportional solution is also envy-free and equitable.
- For two agents, proportionality implies envy-freeness. This implication, however, does not hold for arbitrary n. For example, an agent can believe it has got more than its fare share, but can still be envious of another agent which it believes has received an even larger share.

There has been significant interest and progress on the complexity of computing n-person envy-free allocations (Stromquist, 2008; Woeginger and Sgall, 2007).

2.2 Insufficiency of envy-freeness

Envy-free procedures produce allocations such that each party considers itself to be better than anyone else. A self-interested agent may, however, want more. 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!

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 extremely attractive. The question then is what information needs to be known about others for an agent to corner a larger than envy-free share.

Though it can be practically possible for a rational agent to risk losing the envy-free guarantee in the hope of gaining substantial benefit most of the time, we will not pursue this line of investigation in the current paper. For the rest of the paper, we concentrate exclusively on dividing a good between two agents. We assume that the good being divided is heterogeneous and continuously divisible and the preference of an agent for various parts of the good is represented by a non-negative utility function.

3 An envy-free procedure

Division procedures will, in general, have two components: a rule set to be followed in arriving at an allocation, and strategies to be used by the players to obtain fair shares. For example, in the above-mentioned “divide and choose” scheme, the procedure is that one player divides the cake, and the other gets to choose the part it wants. Within this procedure, a strategy that guarantees envy-freeness is as follows: the divider divides the cake into half by its estimate, and the chooser chooses the part that it considers to be at least as big as the other part.

We first narrate Austin’s moving-knife procedure for arriving at an envy-free division of a cake (Austin, 1982). One of the players, say A, places one knife at the edge of the cake and parallel to it, and a second knife parallel to the first one such that half of the cake, by its estimate, lies in between. Then A moves the knives in parallel over the cake to the right always making sure that exactly half of the cake, by its estimate, lies between the knife. The other player, B, can call *stop* at any time and receives the portion of the cake in between the knives at that point (see Figure 1). If B does not call out in between, the final knife positions will be such that the right knife is at the right edge of the cake and the left knife is at the starting position of the right knife. Since the beginning and the end halves of the allocation presented to B are exactly complimentary, there must be a position in between when the portion of the cake in between the two moving knives is exactly half of the cake by B’s estimate. In this procedure B is guaranteed to get half or more of the cake by its estimate while A is guaranteed to get exactly half of the cake by its estimate if it always presents half of the cake, by its estimate, between the two knives. To be envy-free, A must start, and hence end, with exactly half of the cake between the knives. But A can get more than that by presenting less than half of the cake at intermediate positions, and if B happens to stop the procedure at such a configuration.

Now we analyze this procedure for efficiency. Suppose that agent B calls cut and gets the portion between the knives. It might be the case that there is a portion immediately to the right of the right knife that has no utility to A, but is of value to agent B. This region can be transferred to agent B, thereby increasing its utility. At the same time this transfer does not reduce the utility of agent A. Thus the strategy to guarantee envy-free division in the Austin’s moving-knife procedure does not guarantee efficiency.

We now present a strong result that eliminates the hope of producing optimal divisions in

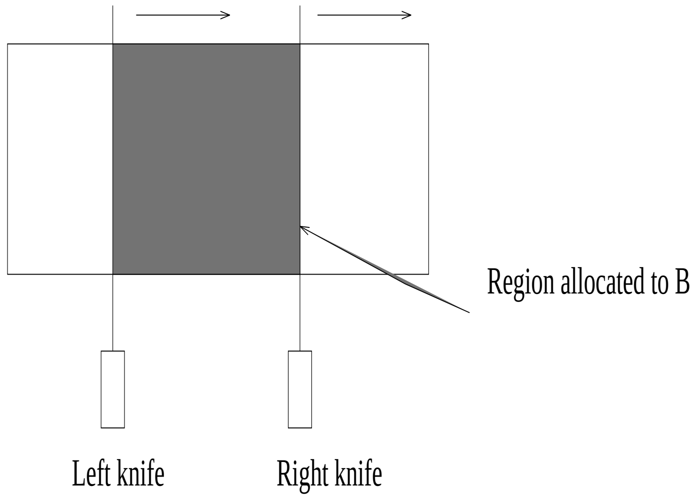

Figure 1: Austin's two-knife procedure. In our case the modeling agent moves the knives.

general for continuously divisible goods.

Theorem 1. There does not exist an algorithm that guarantees an effi cient division of a continuously divisible good between two agents.

Proof: Our proof involves generating a problem scenario for which fi nite decision mechanisms to generate optimal allocations do not exist. Consider the scenario where the continuously divisible good can be divided into infi nitely many pieces such that any two neighboring pieces have the property that the fi rst piece is valued by one agent and is of no value to the other agent and the opposite holds true for the second piece. To produce an optimal allocation in this case requires infi nitely many cuts of the cakes. As such, no fi nite procedure or algorithm can produce an optimal allocation in this situation.

4 Envy-free divisions with improved effi ciency

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 (we assume $U_i(x)$ to be non-negative). Without loss of generality, we will assume that the fi rst 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 benefi t 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 and Taylor, 1996). In the *2e-opt* procedure, 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, where both values are in the range $[0, L]$, where L is the length of the cake. 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.

We allow B to see the entire process before deciding what region to choose as a partial counter-measure to the advantage that A has in being the one who moves the knives. Even then, the procedure remains asymmetric. It does not mean, however, that A can necessarily gain at the expense of B. Rather, the goal here is to find allocations that dominate those obtained with Austin's procedure.

Given the above rules, what should be the strategy of the modeling agent? Its goal is to give B just enough to make it envy-free. Intuitively, A should try to give B half of the cake according to B's estimate while giving up as little as possible of what it values.

On first thought, then, A should start with the left and right knives at the positions for which the intermediate part is least useful for A and the same part is most wanted by B. But this will not guarantee envy-freeness for A. To see why this is the case let m and n be the initial positions of the left and right knives. Now, it may be the case, that B calls cut immediately or never calls cut. In the first case, A gets a part of the cake of utility $\int_n^m U_A(x)dx$; in the second case, A gets a part of the cake of utility $\int_m^n U_A(x)dx$ (for ease of exposition, if $y < x$, we use $\int_x^y f(z)dz$ to represent $\int_x^L f(x)dx + \int_0^y f(z)dz$). Since these two parts are complementary, it follows that to guarantee envy-freeness, the initial knife placements should be such that they divide the cake exactly in half by A's estimate. So, the initial knife placements should be such that

$$\int_n^m U_A(x)dx = \frac{1}{2}\int_0^L U_A(x)dx, \quad (4.1)$$

where L is the length of the cake.

It is also desirable that the initial knife placements be such that the region in between does divide the cake in half for agent B. Otherwise it will call cut immediately or may get more than a proportional share by waiting until the end. In either case, A will not be able to exploit its model of B's utility function and would have to be satisfied with just half of the cake by its estimate. But a contiguous region that satisfy both these conditions may not be found in general.

4.1 Initial region selection

The initial placement (m, n) of the knives should be such that in addition to satisfying Equation 4.1 it should also satisfy the following condition:

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

i.e., the initial allocation is to be as close as possible to half of the cake according to B's estimate. Since A has only an approximate model of B's utility function, the best it can do is to choose the initial placements such that

$$(m, n) = \arg\min_{y,x} \left| \int_x^y \overline{U_B}(z)dz - \frac{1}{2}\int_0^L \overline{U_B}(z)dz \right|. \quad (4.2)$$

Though B can get close to half of the cake by its estimate at the beginning or the end, some intermediate offering may be of more value to it.

4.2 Target region selection

We will now present an analysis by which the modeling agent selects the region of the cake that it will most prefer to offer to B which will still be acceptable to B, i.e., accepting it will give an envy-free allocation to B. Such a region delimited by pair of points m and n must satisfy both of the following inequalities:

$$\int_m^n U_A(x)dx \leq \frac{1}{2}\int_0^L U_A(x)dx,$$

$$\int_m^n \overline{U_B}(x)dx \geq \frac{1}{2}\int_0^L \overline{U_B}(x)dx.$$

The fi rst of the above says that A will be envy-free in giving up the region between m and n, and the second inequality says that B will be envy-free in receiving that region. Put together, these inequalities identify regions which A will prefer to give up assuming that B will not accept a non-envy-free share. A slight enhancement, eliminates the beginning and end regions to return the following set of point pairs (each point pair defi nes a region or possible allocation):

$$\begin{aligned} Z &= \{(x,y) : \int_x^y \overline{U_B}(z)dz > max(\int_{l_0}^{r_0} \overline{U_B}(x)dx, \int_{l_T}^{r_T} \overline{U_B}(x)dx)) \\ &\quad \wedge (\int_x^y U_A(z)dz < \frac{1}{2}\int_0^L U_A(z)dz)\}. \end{aligned}$$

From the regions in Z, we fi rst choose the set of Pareto-optimal allocations:

$$Z_{PO} = \{(a,b) \in Z | \forall (m,n) \in Z, \int_m^n U_A(z)dz > \int_a^b U_A(z)dz \vee \int_a^b \overline{U_B}(z)dz > \int_m^n \overline{U_B}(z)dz.$$

From Z_{PO}, we then choose regions with the least value to A:

$$Y = \{(a,b) \in Z_{PO} | \forall (m,n) \in Z_{PO}, \int_m^n U_A(z)dz \geq \int_a^b U_A(z)dz.$$

From the regions in Y, A can select those regions which are estimated to be of most value to B:

$$X = \{(a,b) \in Y | \forall (m,n) \in Y, \int_a^b \overline{U_B}(z)dz \geq \int_m^n \overline{U_B}(z)dz.$$

The above is a maximin procedure for selecting regions that A desires to offer to B. 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 such a procedure for moving the knives is followed, B will choose this region over the initial or fi nal offerings, and A will have succeeded in making use of its model of B's utility function to obtain a larger than equal share. If X is empty, however, A will have to be satisfi ed with half of the cake.

4.3 Moving the knives

While moving between pairs of points as identifi ed in the previous paragraph, the knife locations (u, v) should satisfy the following inequality:

$$\int_u^v \overline{U_B}(x)dx < \frac{1}{2}\int_0^L \overline{U_B}(x)dx.$$

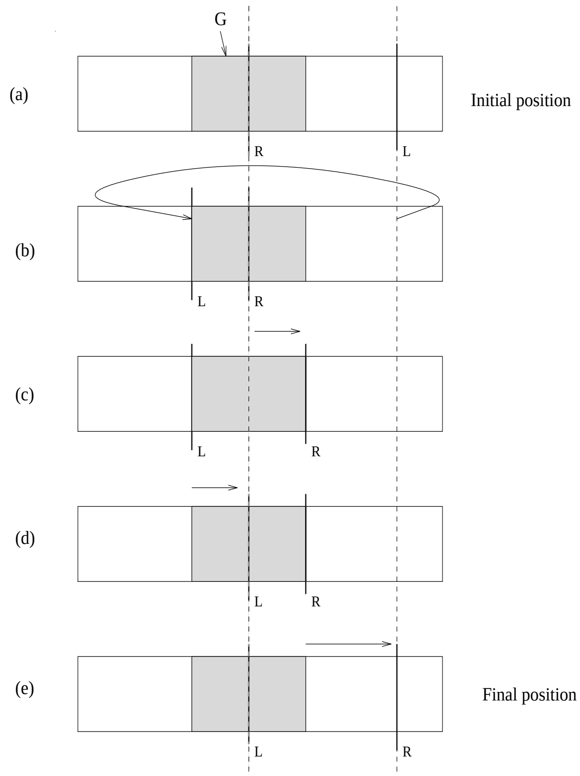

Figure 2: A procedure for knife movement to take advantage of a model of the other agent's utility function.

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.

We now analyze the effects of an inaccurate model of B's utility function as used by A. By enforcing the constraint that at no time during the procedure the region between the two knives is greater than half the cake's worth in A's estimate, we have eliminated the possibility of a division which is non-envy-free for A. The worst case scenario is that A receives exactly half of the cake in its estimate. Interestingly, this would also be the result if both agents have identical utility functions irrespective of whether A has an accurate model of B's utility function.

At this point, we have completed a specifi cation of a procedure to be followed by our modeling agent to exploit its model of the other agent. Another way of looking at the above description is a characterization of a class of strategies that guarantees envy-freeness but can provide a signifi cantly larger share of the cake than that obtained by Austin's moving-knife procedure.

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 and Dutta, 2001).

5 A procedure for improving effi ciency of envy-free divisions

Now we present a particular instance of these class of strategies assuming that the set X is non-empty. The particular procedure we are going to describe is also presented pictorially in

Figure 2. We first choose a member of this set that defines a region G that A wants B to take. The distinct steps of the procedure are as follows:

1. A should begin by putting its right knife somewhere inside the region G and the left knife to its right (allowing for wrap-over, which means that in some cases, the left knife would be situated to the left of the right knife to start with) such that the region between these knives satisfies the criteria given in equations 4.1 and 4.2 respectively (see Figure 2(a)).

2. The left knife is moved right, allowing for wrap-around, until it coincides with the left edge of region G (see Figure 2(b)).

3. The right knife is moved until it coincides with the right edge of region G (see Figure 2(c)).

4. The left knife is moved until it coincides with the initial position of the right knife (see Figure 2(d)).

5. The right knife is moved until it coincides with the initial position of the left knife (see Figure 2(e)).

6 Properties of the proposed division procedure

We now present some characterizations of the properties of the decision procedure described above. All of the following theorems will assume that the model being used of the other agent is an accurate one, i.e., $\overline{U_B}(x) \equiv U_B(x)$. We will also assume that B is acting rationally. That is, it chooses the region which maximizes its utility from the set of regions offered. If B chooses irrationally, A benefits by getting a larger allocation than expected.

Theorem 2. If a Pareto-optimal, envy-free, allocation for a problem involves a contiguous region, the *2e-opt* scheme will select it.

Proof: We will prove the theorem by contradiction. We start by assuming that there exists a contiguous region which when allocated to B produces a Pareto-optimal allocation. Let our division procedure allocate the region bounded by points (u, v) to agent B. Suppose there exists another contiguous region, bounded by points (p, q) which when allocated to B produces an envy-free allocation that dominates the allocation selected by our procedure. This means that $\int_p^q U_A(x)dx \leq \int_u^v U_A(x)dx$ and $\int_p^q U_B(x)dx \geq \int_u^v U_B(x)dx$ with at least one of the inequalities being satisfied in the strict sense. But since maximin strategy chooses (u, v) over (p, q) then one of the following must be true:

1. $(p, q) \notin Z$, i.e., the allocation is not envy-free for both agents,
2. $(p, q) \in Z \wedge (p, q) \notin Y$, i.e., $\int_p^q U_A(x)dx > \int_u^v U_A(x)dx$,
3. $(p, q) \in Y \wedge (p, q) \notin X$, i.e., $\int_p^q U_A(x)dx = \int_u^v U_A(x)dx \wedge \int_p^q U_B(x)dx < \int_u^v U_B(x)dx$.

This means the allocation to B of the region bounded by the points (p, q) is either non-envy-free or it does not dominate the allocation of the region bounded by the points (u, v). With this contradiction, we conclude that the contiguous region selected is the Pareto-optimal contiguous allocation.

Theorem 3. The *2e-opt* scheme dominates Austin's procedure with respect to effi ciency of allocations.

Proof: In Austin's procedure using a knife, division of the cake is a contiguous one and from theorem 2 we know that any allocation given by Austin's procedure cannot be more effi cient. However, Austin's procedure cannot make contiguous allocations such that a middle part of the cake goes to one agent and the front and end goes to the other agent. Our procedure can produce that because of we allow the knives to wrap-around. For situations where allocations of a contiguous region from the middle of the cake is necessary, the allocation produced by our procedure improves on the allocation produced by Austin's procedure. Thus our procedure dominates Austin's procedure with respect to effi ciency of allocations.

Theorem 4. *2e-opt* produces Pareto-optimal allocations for rational agents with monotonic utility function.

Proof: When utility functions are monotonic for both agents, the desired region that A would like to give to B will be a contiguous region at one of the two ends of the cake. From theorem 2 it follows that such a contiguous region will be selected by our scheme.

Theorem 5. *2e-opt* scheme produces Pareto-optimal allocations for rational agents with utility functions with a single optimum.

Proof: The Pareto-optimal allocation will be a contiguous region which will be selected by our scheme.

7 Recursive division to improve optimality

The *2e-opt* procedure 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 effi ciency 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, and the time remaining for negotiating an outcome. The algorithm is presented in Figure 3.

The above algorithm combines and encapsulates both the specifi cation 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

```
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.

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 effi cient 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 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 effi cient 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 effi cient solution than that achieved by Austin's two person envy-free allo-

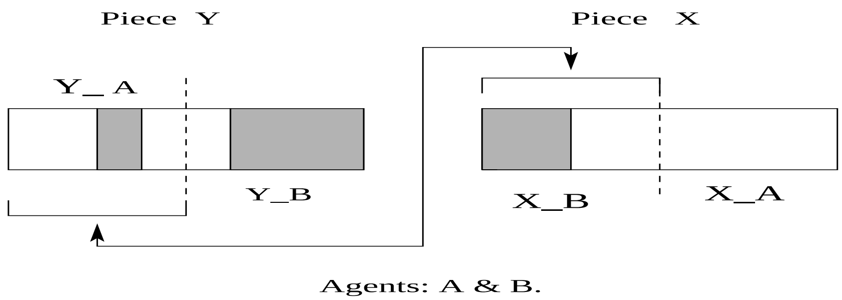

Figure 4: A plausible envy-free exchange between agents A and B.

cation 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 effi cient 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 effi ciency.

8 Discussion

Recent work in multiagent systems literature on envy-free-divisions focus on a fi nite set of rational agents negotiating the allocation of a set of indivisible resources in a distributed setting (Chevaleyre, Endriss, Estivie and Maudet, 2007). In contrast, our work involves negotiation between two agents on a continuously divisible good.

The *2e-opt* protocol presented here produces optimal envy-free divisions of continuously divisible goods between two agents for a restricted class of utility functions. *2e-optR*, the recursive extension (of course, an equivalent iterative version would do just as well) of this 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, i.e., successive recursive calls will produce smaller effi ciency improvements. That implies that the earlier exchanges are likely to produce greater increase in effi ciency compared to latter exchanges. The computation costs of each negotiation stage should be roughly the same. This raises the issue of a trade-off between computation cost and effi ciency improvement. If the negotiation is done off-line or with a fi xed 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 and 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 and 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.

In this paper, the agent moving the knife is willing to consider inequitable shares, e.g., it can offer another agent half of the cake which may be worth more than half to the other player. A more greedy knife-mover may choose only to offer those elements of Z (see Section 4.2) that it values the least. While from a rational, utility-maximizing agent's point of view, that would be the preferred approach, Pareto-optimality guarantees presented in Section 6 would be lost.

References

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

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.

Chevaleyre, Y., Endriss, U., Estivie, S. and Maudet, N. 2007. Reaching envy-free states in distributed negotiation settings, *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pp. 1239–1244.

Haake, C.-J., Raith, M. G. and Su, F. E. 2002. Bidding for envy-freeness: A procedural approach to n-player fair-division problems, *Social Choice and Welfare* **19**: 723–749.

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

Jennings, N., Faratin, P., Lomuscio, A. R., Parsons, S., Sierra, C. and Wooldridge, M. 2001. Automated negotiation: prospects, methods and challenges, *International Journal of Group Decision and Negotiation* **10**(2): 199–215.

Kraus, S. 2001. *Strategic Negotiation in Multi-Agent Environments*, MIT Press.

Krishna, V. 2002. *Auction Theory*, Academic Press.

Robertson, J. and Webb, W. 1998. *Cake Cutting Algorithms: Be Fair If You Can*, A. K. Peters, Natick, MA.

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

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

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

Stromquist, W. 2008. Envy-free cake divisions cannot be found by finite protocols, *The Electronic Journal of Combinatorics* **15**(R#11).

Woeginger, G. and Sgall, J. 2007. On the complexity of cake cutting, *Discrete Optimization* **4**: 213–220.