How to Cut Cake

An Overview of Fair Online Resource Allocation Problems

Arkaprava Choudhury

University of Toronto

arka@cs.toronto.edu

Canadian Undergraduate Math Conference
July 15, 2022

Presentation Overview

Breaking it down

Fair *Online* *Resource Allocation* *Problems*

Components

1. "Fair"
2. "Online"
3. "Resource allocation"

Allocation problems

Problem statement

You have some quantity, say *m units*, of some *resource*,

Allocation problems

Problem statement

You have some quantity, say *m units*, of some *resource*, and you have *n people* to allocate the resource to.

Allocation problems

Problem statement

You have some quantity, say *m units*, of some *resource*, and you have *n people* to allocate the resource to. You wish to maximize a particular objective through this allocation.

Allocation problems

Problem statement

You have some quantity, say *m units*, of some *resource*, and you have *n people* to allocate the resource to. You wish to maximize a particular objective through this allocation.

Definitions

1. Divisibility whether the resource can be divided, and, if applicable, the finest refinement possible

Allocation problems

Problem statement

You have some quantity, say *m units*, of some *resource*, and you have *n people* to allocate the resource to. You wish to maximize a particular objective through this allocation.

Definitions

1. Divisibility whether the resource can be divided, and, if applicable, the finest refinement possible
2. Homogeneity whether all parts of the resource are worth the same to each person

Allocation problems

Problem statement

You have some quantity, say *m units*, of some *resource*, and you have *n people* to allocate the resource to. You wish to maximize a particular objective through this allocation.

Definitions

1. Divisibility whether the resource can be divided, and, if applicable, the finest refinement possible
2. Homogeneity whether all parts of the resource are worth the same to each person
3. Allocation a *partitioning* of the available resource amongst (a subset of) the *population*

Allocation problems

Problem statement

You have some quantity, say *m units*, of some *resource*, and you have *n people* to allocate the resource to. You wish to maximize a particular objective through this allocation.

Definitions

1. Divisibility whether the resource can be divided, and, if applicable, the finest refinement possible
2. Homogeneity whether all parts of the resource are worth the same to each person
3. Allocation a *partitioning* of the available resource amongst (a subset of) the *population*
4. Objective we shall take this to be the net worth of the allocation, subject to *fairness*

Online algorithms

Intuitive idea

A model of algorithms accepting an input instance given as an unknown sequence of inputs (agents, in this case).

Online algorithms

Intuitive idea

A model of algorithms accepting an input instance given as an unknown sequence of inputs (agents, in this case). After each input agent is presented, the algorithm makes a decision (*irrevocably*, in this case).

(Minimal) online cake-cutting

Defining the problem

Informally...

Congratulations! Today is your birthday so you take a cake into the office to share with your colleagues at tea time. However, as some people have to leave early, you cannot wait for everyone to arrive before you start sharing (allocate) the cake. How do you proceed fairly?
— Toby Walsh (Online Cake Cutting, 2011)

(Minimal) online cake-cutting

Defining the problem

Informally...

Congratulations! Today is your birthday so you take a cake into the office to share with your colleagues at tea time. However, as some people have to leave early, you cannot wait for everyone to arrive before you start sharing (allocate) the cake. How do you proceed fairly?
— Toby Walsh (Online Cake Cutting, 2011)

Simplification

cake $\rightarrow I = [0, 1]$;

Allocation

Cutting

If S is a finite set of closed intervals, then:

1. S is a cutting;
2. $\forall [a,b] \in S$ and $c \in (a,b)$, the set $S \cup \{[a,c],[b,c]\} \setminus [a,b]$ is a cutting.

Allocation

Cutting

If S is a finite set of closed intervals, then:

1. S is a cutting;
2. $\forall [a,b] \in S$ and $c \in (a,b)$, the set $S \cup \{[a,c],[b,c]\} \backslash [a,b]$ is a cutting.

Allocation

An allocation of the cake $I = [0,1]$ among the set of agents $[n]$ is a partition of some cutting of $\{I\}$ into n subsets, $A_1, \ldots, A_n$.

Allocation

Cutting

If S is a finite set of closed intervals, then:

1. S is a cutting;
2. $\forall [a, b] \in S$ and $c \in (a, b)$, the set $S \cup \{[a, c], [b, c]\} \backslash [a, b]$ is a cutting.

Allocation

An allocation of the cake $I = [0, 1]$ among the set of agents $[n]$ is a partition of some cutting of $\{I\}$ into n subsets, $A_1, \ldots, A_n$.

Simple allocation

An allocation using only n disjoint intervals.

Agent preferences

Valuation

For each $j \in [n]$, define the valuation of agent j denoted by $v_j : 2^I \to \mathbb{R}_{\geq 0} = \int f_j$ which is, for all $j \in [n]$:

Agent preferences

Valuation

For each $j \in [n]$, define the valuation of agent j denoted by $v_j : 2^I \to \mathbb{R}_{\geq 0} = \int f_j$ which is, for all $j \in [n]$:

- normalized: $v_j(I) = 1$
- additive: for any two closed disjoint sub-intervals X, Y, $v_j(X \sqcup Y) = v_j(X) + v_j(Y)$

Agent preferences

Valuation

For each $j \in [n]$, define the valuation of agent j denoted by $v_j : 2^I \to \mathbb{R}_{\geq 0} = \int f_j$ which is, for all $j \in [n]$:

- normalized: $v_j(I) = 1$
- additive: for any two closed disjoint sub-intervals X, Y, $v_j(X \sqcup Y) = v_j(X) + v_j(Y)$

Set valuation

For a finite set of intervals S, we define, for all $j \in [n]$, $v_j(S) = \sum_{[a,b] \in S} v_j([a, b])$

Some classic requisites

1. (Strong) proportionality:
   "Each agent *feels* they got a fair share of the cake"

Some classic requisites

1. (Strong) proportionality:
   "Each agent *feels* they got a fair share of the cake"

$$\forall j \in [n], v_i(A_i) \geq 1/n$$

Some classic requisites

1. Proportionality
2. No envy:
   "No agent is envious of some other agent's share"

Some classic requisites

1. Proportionality
2. No envy:
   "No agent is envious of some other agent's share"

$$\forall i, j \in [n], v_i(A_i) \geq v_i(A_j)$$

Some classic requisites

1. Proportionality
2. No envy
3. Equitability:
   "All agents are equally content with their share"

Some classic requisites

1. Proportionality
2. No envy
3. Equitability:
   "All agents are equally content with their share"

$$\forall i, j \in [n], v_i(A_i) = v_j(A_j)$$

Some classic requisites

1. Proportionality
2. No envy
3. Equitability
4. Truthfulness:
   “No agent can profit by falsifying their preferences”

...

Online fairness criteria

Lemma

No envy implies proportionality.

Online fairness criteria

Lemma

No envy implies proportionality.

Lemma

No online cake cutting algorithm is proportional, envy-free, or equitable

Proof.

Suppose agent i leaves before agent n arrives. A_i is then independent of v_n. If $v_n(A_i) = 1$, agent n will not value any allocation outside A_i. So, not proportional. Since no envy implies proportionality, not envy-free either.

Suppose allocation was equitable, so all agents receive some cake. Again, A_i is independent of v_n for the first leaving agent i. □

Online fairness criteria

Online proportionality

Weak proportionality

Each agent j assigns at least r/k of the total value of the cake to their pieces where

1. r is the value of the remaining amount of unallocated cake when agent j arrives;
2. k is the number of agents yet to be allocated cake at this point.

Online fairness criteria

Online no envy

1. Weakly envy-free: agents do not value cake allocated to agents after their arrival more than their own;
2. Immediately envy-free: agents do not value cake allocated to any agent after their arrival and before their departure more than their own

Online fairness criteria

Online no envy

1. Weakly envy-free: agents do not value cake allocated to agents after their arrival more than their own;
2. Immediately envy-free: agents do not value cake allocated to any agent after their arrival and before their departure more than their own

Lemma

No envy implies weakly envy-free. Weakly envy-free implies immediately envy-free.

Online fairness criteria

Online equitability

First-come-first-serve

No agent's value of their assigned share can decrease if they arrive earlier in the input sequence and all other agents are left in the same relative positions; formally defined as arrival monotone.

Lemma

Equitability implies arrival monotonicity.

Cut-and-choose algorithm

Each application shall [...] allow two mining operations. The Authority shall designate which part is to be reserved solely for the conduct of activities by the Authority.
— UN Convention on the Law of the Sea

Cut-and-choose algorithm

Algorithm 1 I cut but you choose

```
procedure CUT-AND-CHOOSE
    for j = 1 → n − 1 rounds do
        The earliest arriving agent cuts the cake into two disjoint intervals X, Y such that v_j(X) = v_{j+1}(Y) and X ⊔ Y = I_j.
        The second earliest arriving agent j + 1 chooses whether to take X and leave, or give X to the cutting agent who leaves.
        I_{j+1} ← Y.
    end for
    The last remaining agent takes the leftover cake.
end procedure
```

Fairness of cut-and-choose

Lemma

The online cut-and-choose procedure is weakly proportional and immediately envy free. However, it is not weakly envy free, equitable, or arrival monotonic.

Proof.

Suppose agent i cuts a slice c_i. If allocated the slice, they would want $v_i(c_i) \geq r/k$. But, if not allocated this piece, they would want $v_i(c_i) \leq r/k$. Thus, the best option is to choose $v_i(c_i) = r/k$.

By generalization, this holds for all i, so this is weakly proportional. Also, trivially immediately envy-free.
Consider the following counter-example with 4 agents in the order $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$ for the negative results. □

Fairness of cut-and-choose (contd.)

Proof.

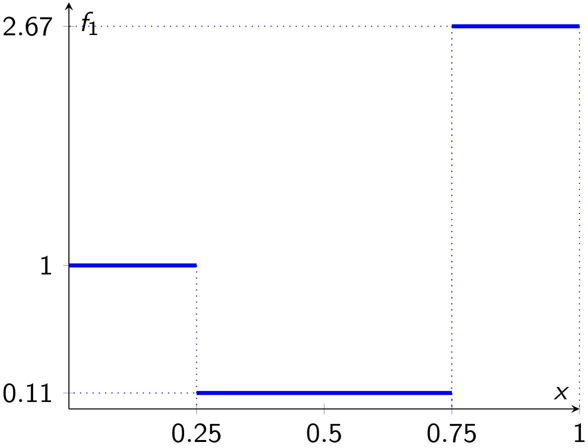

□

Fairness of cut-and-choose (contd.)

Proof.

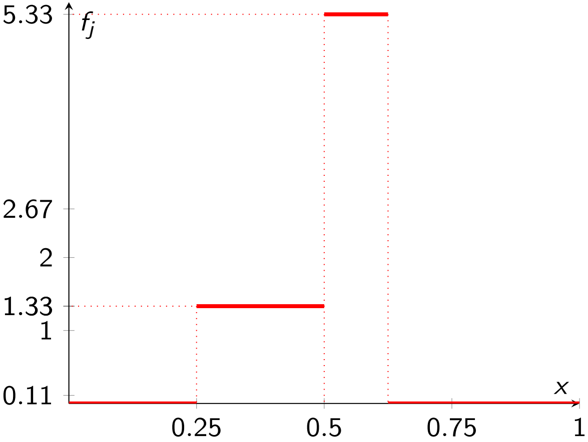

□

Fairness of cut-and-choose (contd.)

Proof.

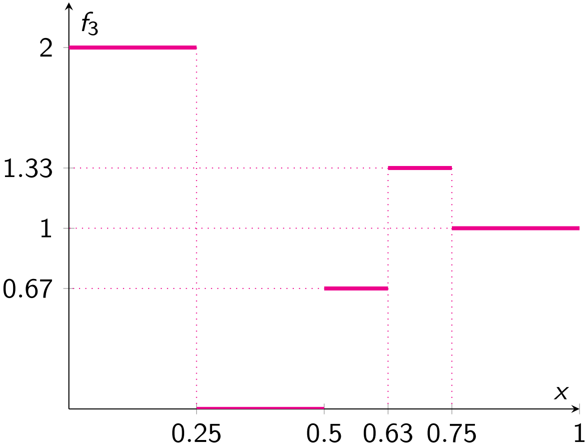

□

Fairness of cut-and-choose (contd.)

Proof.

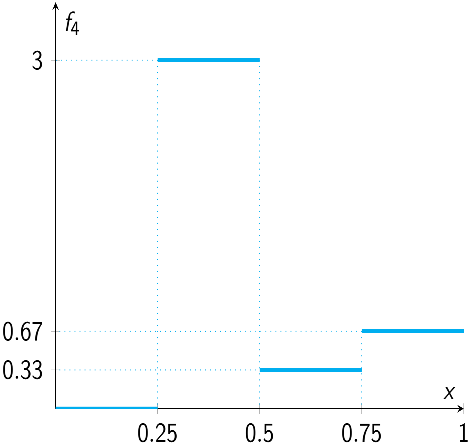

□

Online moving knife

1. several rounds of cutting ($n-1$ rounds for minimal cutting)
2. in each round, the algorithm moves a knife from the left to the right, and only stops when some agent declares it to stop
3. at that point, the algorithm cuts the cake and that agent leaves with their share of cake, i.e., the part to the left of the cut.

Dubins-Spanier procedure

Briefly...

Dubins-Spanier procedure

Briefly...

1. Given $k < n$, start a moving knife procedure with the first k agents.
2. At the end of the procedure, if the last agent is yet to come, then wait for the next agent and restart the procedure with k agents again.
3. If there are no more agents to come, restart with $k - 1$ agents. Repeat until only one agent remains. Allocate the remainder of the cake to that agent.

Dubins-Spanier procedure

Briefly...

1. Given $k < n$, start a moving knife procedure with the first k agents.
2. At the end of the procedure, if the last agent is yet to come, then wait for the next agent and restart the procedure with k agents again.
3. If there are no more agents to come, restart with $k-1$ agents. Repeat until only one agent remains. Allocate the remainder of the cake to that agent.

Lemma

The online moving knife procedure is weakly proportional and immediately envy free. However, it is not (weakly) envy free or arrival monotonic.

Dubins-Spanier procedure

If you are curious...

Theorem

Consider a set S and n agents, and let $\mathbb{U}$ be a σ-algebra on S. Suppose each agent j has a countably-additive and nonatomic value measure $v_i : \mathbb{U} \to \mathbb{R}$. Let K be a k-partition of S. Then, the set of all $n \times k$ matrices $[M]_{ij}$ is a compact and convex set in the space of all real-valued $n \times k$ matrices.

Truthfulness

Existing work

There exist deterministic non-minimal cutting algorithms which guarantee truthfulness. There also exist randomized **minimal** cutting algorithms guaranteeing truthfulness.

Open question

With what restrictions can we sacrifice randomness without losing minimalism?

Other query models

Robertson-Webb

Two oracles for each $j \in [n]$ as follows

1. $Eval_j(x, y)$
2. $Cut_j(x, \alpha)$

Other query models

Robertson-Webb

Two oracles for each $j \in [n]$ as follows

1. $Eval_j(x, y)$
2. $Cut_j(x, \alpha)$

Simultaneous encoding

All agents succinctly report their discretized value allocations on arrival.

More variations of resource allocation

1. Multi-cake
2. Homogeneous goods
3. Indivisible goods
4. Combinatorial auctions

References

T. Walsh (2011)

Online Cake Cutting

Algorithmic Decision Theory. ADT 2011.

L. E. Dubins, E. H. Spanier (1961)

How to Cut a Cake Fairly

American Mathematical Monthly. Jan 1961.

Acknowledgements

UofT

- Prof. Allan Borodin
- Prof. Nathan Wiebe

\+ my friends @ UofT

Funding

- UofT Dept of Comp Sci
- CQIQC
- The MatterLab group

The End

Questions? Comments?