Information Processing Letters ••• (••••) •••–•••

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

A note on envy-free cake cutting with polynomial valuations

Simina Brânzei

Computer Science Department, Aarhus University, IT-parken, Aabogade 34, DK 8200 Aarhus, Denmark

ARTICLE INFO

Article history:
Received 20 May 2013
Received in revised form 7 July 2014
Accepted 9 July 2014
Available online xxxx
Communicated by B. Doerr

Keywords:
Algorithms
Cake cutting
Query complexity
Envy-freeness

ABSTRACT

The cake cutting problem models the fair allocation of a heterogeneous divisible resource among multiple players. The central fairness criterion is envy-freeness and a major open question in this domain is the design of a bounded protocol that can compute an envy-free allocation of the cake for any number of players. The only existing finite envy-free cake cutting protocol for any number of players, designed by Brams and Taylor [4], has the property that its runtime can be made arbitrarily large by setting up the valuation functions of the players appropriately. Moreover, there is no closed formula that relates the valuation functions to the number of queries required by the protocol.

In this note we show that when the valuations can be represented as polynomial functions, there exists a protocol in the standard query model that is much simpler conceptually and has a runtime bound depending on the maximum degree over all polynomials.

1. Introduction

Cake cutting is a fundamental model in fair division; it represents the problem of allocating a heterogeneous divisible resource – such as land, time, clean water, mineral deposits, and computer memory – among players with different preferences. The cake cutting problem has been studied by Banach, Knaster, and Steinhaus [17] since World War II; since then, a growing body of literature in mathematics, political science, economics, and computer science has been devoted to its study, including two books by Brams and Taylor [5], Robertson and Webb [16], and a line of recent papers [14,8,9,3,2,13,1,12,7,6].

The classical cake cutting literature includes two complementary research directions, namely (*i*) establishing the existence of allocations with desirable properties and (*ii*) designing protocols to compute such allocations (see [15]). All the known discrete cake cutting protocols model the interaction between the center and the players using a standard query model, which was formalized by Robertson and Webb [16] and used in a line of work studying the complexity of cake cutting (Edmonds and Pruhs [10,11], Woeginger and Sgall [19], Procaccia [14], Kurokawa, Lai, and Procaccia [12]).

Arguably the most prominent criterion of fairness in cake cutting is envy-freeness and an outstanding open question in this domain is the design of a bounded, general protocol that can compute an envy-free allocation of the cake for any number of players. The problem of finding an envy-free protocol for any number of players in the Robertson–Webb model was open for almost half a century, until Brams and Taylor [4] announced a solution. However, the runtime of the Brams–Taylor protocol is unbounded: the number of queries required can be made arbitrarily large by setting up the valuation functions of the players appropriately. Moreover, there is no closed formula that relates the valuation functions to the number of steps required by the protocol.

In a recent paper, Kurokawa, Lai, and Procaccia [12] designed an envy-free cake cutting protocol in the Robertson–Webb model for a succinct type of valuations, namely continuous piecewise-linear functions. Their protocol is guaranteed to produce an envy-free allocation within $O(n^6 k \ln k)$ queries on any given instance, where

E-mail address: simina@cs.au.dk.

http://dx.doi.org/10.1016/j.ipl.2014.07.005
0020-0190/© 2014 Elsevier B.V. All rights reserved.

n is the number of players and k is the number of break points in the valuations (break points are discontinuities in the derivatives of the valuation functions). Kurokawa, Lai, and Procaccia [12] also showed that if a protocol can compute envy-free allocations for the class of piecewise uniform valuations (i.e. for which the value density function of any given player can take only one of two values, namely zero or some constant), then it can also solve the envy-free cake cutting problem for general valuations. The result of Kurokawa, Lai, and Procaccia suggests that the main difficulty is detecting the break points in the value density function (and possibly its derivative). Polynomial valuations are interesting from the point of view of envy-free cake cutting because no such discontinuities exist, yet no bounded protocol is known for this class. In this note, we show that polynomial valuations admit a conceptually intuitive protocol for which the query complexity depends on the maximum degree of the polynomials.

2. Background

Let $N = \{1, \ldots, n\}$ be a set of players. The cake is modeled as the interval $[0, 1]$. A *piece of cake* X is a finite set of disjoint intervals of $[0, 1]$.

Each player is endowed with an integrable, non-negative value density function $v_i(x)$ that induces a value for each possible piece of cake. Formally, the value of player i for a piece X is given by:

$$V_i(X) = \sum_{I \subset X} \int_I v_i(x)dx.$$

By definition, the valuations of the players are additive, i.e. $V_i(X \cup Y) = V_i(X) + V_i(Y)$ if X and Y are disjoint; and non-atomic, i.e. $V_i([x, x]) = 0$, for all $x \in [0, 1]$. We assume that each player has a value of one for the entire cake: $V_i([0, 1]) = 1$ for all $i \in N$. This assumption is without loss of generality for the purposes of this paper.

An *allocation* $A = (A_1, \ldots, A_n)$ is a partition of the cake among the players, that is, each player i receives the piece A_i, the pieces are disjoint, and $\bigcup_{i \in N} A_i = [0, 1]$.

The central criteria for determining the fairness of an allocation are *proportionality* and *envy-freeness*. An allocation A is *proportional* if $V_i(A_i) \geq \frac{1}{n}$, for all $i \in N$, and it is *envy-free* if $V_i(A_i) \geq V_i(A_j)$, for all $i, j \in N$. Envy-freeness is a strong fairness notion which implies proportionality when the entire cake is allocated.

The standard query model in cake cutting – which captures all the classical discrete cake cutting protocols – was proposed by Robertson and Webb [16]; it models the interaction between the protocol and the players using two types of queries:

1. $Evaluate_i(x, y)$:
   Player i is asked to output α such that $V_i([x, y]) = \alpha$.
2. $Cut_i(x, \alpha)$:
   Player i is asked to output y such that $V_i([x, y]) = \alpha$.

The number of steps of a discrete cake cutting protocol is measured by the number of *Cut* and *Evaluate* queries made during its execution. The protocol has unlimited computational power, and once it has retrieved enough information about the valuations of the players in the given query model, it can simply output an envy-free allocation of the cake that is consistent with the answers of the players.

3. Protocol for polynomial value density functions

In this section we describe the envy-free protocol for polynomial valuations. First note that if the value density functions can be expressed as polynomials, then for each player $i \in N$, we have: $v_i(x) = \sum_{j=0}^{d_i} a_{i,j} \cdot x^j$, where $d_i \in \mathbb{N}$ and $a_{i,j} \in \mathbb{R}$, for all $j \in \{0, \ldots, d_i\}$. Recall that value densities are always non-negative and normalized to give equal weight to all the players. Thus $v_i(x) \geq 0$, for all $x \in [0, 1]$ and $\int_0^1 v_i(x)dx = 1$.

For each player i, let $P_i(x) = V_i([0, x])$, for all $x \in [0, 1]$. Then by the definition of the value densities, we have:

$$\begin{aligned} P_i(x) &= \int_0^x v_i(y)dy \\ &= \int_0^x \left(\sum_{j=0}^{d_i} a_{i,j} y^j \right) dy = \sum_{j=0}^{d_i} a_{i,j} \int_0^x y^j dy \\ &= \sum_{j=0}^{d_i} \left(\frac{a_{i,j}}{j+1} \right) x^{j+1} \end{aligned}$$

Then P_i is a polynomial of degree $d_i + 1$ over $[0, 1]$ with the following properties: $P_i(0) = 0$, $P_i(1) = 1$, and $(P_i(x))' = v_i(x)$. It is a standard calculus fact that each polynomial P_i can be completely recovered given $d_i + 2$ pairs of distinct points with their values: $\langle x_1, P_i(x_1) \rangle, \ldots, \langle x_{d_i+1}, P_i(x_{d_i+1}) \rangle$, where $x_j \in [0, 1]$ for all $j \in \{1, \ldots, d_i + 2\}$.

By definition of P_i, each tuple $\langle x_j, P_i(x_j) \rangle$ has the property that $V_i([0, x_j]) = P_i(x_j)$. If the center knows the maximum degree, d, it can address each player $i \in N$ the queries $Evaluate_i(0, x_j)$ and use the answers to determine the unique interpolating polynomial P_i; this completely recovers the player's value density function, v_i. An envy-free allocation with $n - 1$ cuts is guaranteed to exist for any cake cutting instance (see [18]), and so once the valuations are recovered, the center can simply output such a contiguous envy free allocation. Finally, the center can determine the maximum degree by iteratively trying increasing integer values of d, starting from $d = 0$. The pseudocode is given in Algorithm 1.

Theorem 1. *There exists a protocol in the Robertson–Webb communication model such that on every n-player cake cutting instance with value density functions given by polynomials, the protocol is guaranteed to terminate with an envy-free allocation using $O(d \cdot n^2)$ queries, where d is the maximum degree of any polynomial in the representation.*

Proof. Algorithm 1 starts by assuming that the players have valuations given by polynomials of degree zero (i.e.

Algorithm 1: Envy-Free Cake Cutting with Polynomial Valuations.

1 **input:** Set of players $N = \{1, \ldots, n\}$
2 **output:** Envy-free allocation $A = (A_1, \ldots, A_n)$
3 $d \leftarrow 0$ *// Running upper bound on the degrees of the polynomials*
4 **while** *(true)* **do**
5   **foreach** $(i \in N)$ **do**
6         $x_{i,d} \leftarrow \frac{1}{d+1}$
7         $y_{i,d} \leftarrow \textbf{Evaluate}_i([0, x_{i,d}])$
8         $P_i(x) \leftarrow \text{Interpolate}((0,0), (x_{i,0}, y_{i,0}), \ldots, (x_{i,d}, y_{i,d}))$
9         $w_i(x) \leftarrow (P_i(x))'$ *// Player i's value density function assuming it's a polynomial of max degree d*
10   *// Compute a contiguous envy-free allocation w.r.t.* $\{w_1, \ldots, w_n\}$
11     $A \leftarrow$ Contiguous envy-free allocation w.r.t. $\{w_1, \ldots, w_n\}$
12   *// Ask the players if A is envy-free*
13   **foreach** $((i, j) \in N^2)$ **do**
14         $W_{i,j} \leftarrow \textbf{Evaluate}_i(A_j)$
15   **if** *(Envy-Free(W))* **then**
16   return A *// Output and exit*
17   **else**
18         $d \leftarrow d + 1$ *// Increase the maximum degree and try again*

constant) and increases the degrees with every iteration. Consider the iteration in which the correct upper bound has been reached: $d = \max\{d_1, \ldots, d_n\}$. Then the answers of player i to the evaluate queries on the intervals:

$$\left\{[0,0], \left[0, \frac{1}{d+1}\right], \left[0, \frac{1}{d}\right], \ldots, [0,1]\right\}$$

can be used to obtain $d + 2$ values for the unique interpolating polynomial, of maximum degree $d + 2$. That is, the protocol has obtained the following values:

$$\left\{P_i(0), P_i\left(\frac{1}{d+1}\right), P_i\left(\frac{1}{d}\right), \ldots, P_i(1)\right\}$$

By taking the derivative of the interpolating polynomial (Line 9), the protocol can find the exact value density function of player i. Since d is an upper bound on the degrees of all the players, it follows that all the value density functions have been guessed correctly, and so the allocation A computed in this iteration (Line 11) is guaranteed to be envy-free.

It is immediate that the protocol terminates after at most $d + 1$ iterations, and the number of *Evaluate* queries asked in each iteration is $n^2 + n$. Thus the total number of queries required to output an envy-free allocation when the maximum degree is d is bounded by $(d + 1)(n^2 + n)$. □

An interesting open question is whether there exists a bounded algorithm for envy-free cake cutting with polynomial valuations, where the runtime of the protocol is only a function of the number of players. A negative result for this class would also answer the existence question for general valuations.

Acknowledgements

I would like to thank Ariel D. Procaccia and the anonymous reviewers for helpful feedback. This work was supported by the Sino-Danish Center for the Theory of Interactive Computation, funded by the Danish National Research Foundation and the National Science Foundation of China (under the grant 61061130540), and by the Center for research in the Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council.

References

[1] Y. Aumann, Y. Dombb, A. Hassidim, Computing socially-efficient cake divisions, in: Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2013, pp. 343–350.
[2] X. Bei, N. Chen, X. Hua, B. Tao, E. Yang, Optimal proportional cake cutting with connected pieces, in: Proceedings of the 26th Conference on Artificial Intelligence (AAAI), 2012, pp. 1263–1269.
[3] S.J. Brams, M. Feldman, J. Morgenstern, J.K. Lai, A.D. Procaccia, On maxsum fair cake divisions, in: Proceedings of the 26th Conference on Artificial Intelligence (AAAI), 2012, pp. 1285–1291.
[4] S. Brams, A. Taylor, An envy-free cake division protocol, Am. Math. Mon. 102 (1) (1995) 9–18.
[5] S.J. Brams, A.D. Taylor, Fair Division: From Cake-Cutting to Dispute Resolution, Cambridge University Press, 1996.
[6] S. Brânzei, P.B. Miltersen, Equilibrium analysis in cake cutting, in: Proceedings of the 2th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2013, pp. 327–334.
[7] S. Brânzei, A.D. Procaccia, J. Zhang, Externalities in cake cutting, in: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI), 2013, pp. 55–61.
[8] I. Caragiannis, J.K. Lai, A.D. Procaccia, Towards more expressive cake cutting, in: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI), 2011, pp. 127–132.
[9] Y.J. Cohler, J.K. Lai, D.C. Parkes, A.D. Procaccia, Optimal envy-free cake cutting, in: Proceedings of the 25th Conference on Artificial Intelligence (AAAI), 2011, pp. 626–631.
[10] J. Edmonds, K. Pruhs, Balanced allocations of cake, in: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2006, pp. 623–634.
[11] J. Edmonds, K. Pruhs, Cake cutting really is not a piece of cake, in: Proceedings of 17th ACM-SIAM Symposium on Discrete Algorithms (SODA), 2006, pp. 271–278.
[12] D. Kurokawa, J.K. Lai, A.D. Procaccia, How to cut a cake before the party ends, in: Proceedings of 27th Conference on Artificial Intelligence (AAAI), 2013, pp. 555–561.
[13] A. Maya, N. Nisan, Incentive compatible two player cake cutting, in: Proceedings of the 8th International Conference on Internet and Network Economics (WINE), 2012, pp. 170–183.
[14] A.D. Procaccia, Thou shalt covet thy neighbor's cake, in: Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI), 2009, pp. 239–244.
[15] A.D. Procaccia, Cake cutting: not just child's play, Commun. ACM 56 (7) (2013) 78–87.
[16] J. Robertson, W. Webb, Cake-Cutting Algorithms: Be Fair if You Can, A K Peters/CRC Press, 1998.
[17] H. Steinhaus, The problem of fair division, Econometrica 16 (1) (1948) 101–104.
[18] W. Stromquist, How to cut a cake fairly, Am. Math. Mon. 87 (8) (1980) 640–644.
[19] G.J. Woeginger, J. Sgall, On the complexity of cake cutting, Discrete Optim. 4 (2) (2007) 213–220.

Please cite this article in press as: S. Brânzei, A note on envy-free cake cutting with polynomial valuations, Information Processing Letters (2014), http://dx.doi.org/10.1016/j.ipl.2014.07.005