Networks on Cake Cutting

Ramsundar Anandanarayanan
IIT Bombay
Mumbai, India
ramsundar@cse.iitb.ac.in

Puranjay Datta
IIT Bombay
Mumbai, India
19d070048@iitb.ac.in

Anish Kulkarni
IIT Bombay
Mumbai, India
21b090002@iitb.ac.in

ABSTRACT

A proposed alternative to the classical notion of fairness for cake cutting, has been to introduce an underlying network structure over the agents, and study local fairness in this graphical setting. In this setting, efficient algorithms, both discrete and continuous moving-knife, have been studied for outputting envy-free allocations on several network structures with special properties, most of them being some variant on trees. We study moving-knife algorithms as they are simpler and tend to require a much fewer number of cuts than their discrete counterparts. In this report, first, we modify an existing algorithm for an envy-free allocation on trees, to incorporate an additional edge at level 1 in the network while preserving fairness. Next, we propose moving-knife algorithms for obtaining envy-free allocations on i) cycle networks (upto C_6), and ii) cliques connected via a bridge. We conclude by showing a few more miscellaneous structures, for which we have envy-free algorithms that we will omit, and briefly discuss possible extensions to our results.

KEYWORDS

Mechanism design, Divisible resource, Cake cutting, Fairness, Envy-free, Graphs, Trees, Cycles

ACM Reference Format:
Ramsundar Anandanarayanan, Puranjay Datta, and Anish Kulkarni. 2023. Networks on Cake Cutting. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION AND RELATED WORK

The classical problem of cake-cutting, is an age-old question in the fair allocation setting. The "cake" represents a divisible good to be allotted to n agents. It is, to this day, still a subject of intense research and has many useful variants. The variant we will be concerned with, however, involves imposing a network structure over our set of agents and then studying the problem of fair cake-cutting in this setting. The classical cake-cutting problem was intensively studied in the 20th century, starting with proportional allocations. However, the question of whether or not there exists an *envy-free* allocation was answered only in the 1990s, and a discrete and bounded algorithm for 4 agents ([3]) and n agents ([2]) was only discovered in 2016. Currently, it is known that the lower bound on run-time complexity for discrete algorithms is $\Omega(n^2)$ and the upper bound is $O(n^{n^{n^{n^n}}})$, leaving a very large range of possibilities for its exact complexity.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

As a result, there was a need to look at things from a different perspective. In most practical scenarios, global fairness is unneeded. When allocating a divisible resource over a large group of people, a single agent typically only compares his share with a subset of people, such as his friends or co-workers. We can represent such models by a network structure imposed over the agents; that is, an undirected graph with agents representing nodes and their neighbours being connected by edges. Envy-freeness ([4], [8], [9]) and proportionality ([5]) have been studied on networks with special properties - the idea being that for certain graphs, this less restrictive version of fairness leads to more efficient algorithms. To the best of our knowledge, however, most work in this region of interest has been done on tree networks/networks related to trees, and no one has extensively studied the structures for which we are proposing algorithms.

Another aspect is the usage of moving-knife algorithms over discrete algorithms - the algorithms we propose are all moving-knife ones. These cannot be implemented in a finite number of discrete steps, and thus can have theoretically infinite run-time, but they also provide much simpler algorithms with a fewer number of cuts/better time complexity w.r.t. number of cuts. Both moving-knife and discrete algorithms are thoroughly studied in this setting, and it is our belief that new discoveries pertaining to moving-knife algorithms are important, both in understanding and accelerating progress in this field.

2 OUR CONTRIBUTIONS

(1) We modify an existing $O(n^2)$ algorithm on binary tree networks ([4]) so that it outputs an envy-free allocation on a tree with an **edge added at level 1**, without changing its complexity.
(2) We propose moving-knife algorithms with an efficient number of cuts that output envy-free allocations on **cycle networks** (upto C_6).
(3) We tackle clique networks connected by bridges, and make headway by proposing an algorithm that outputs an envy-free allocation on **3-cliques connected via a bridge**.

We have obtained algorithms for several other miscellaneous structures as well, but these 3 results are our main focus, as they present the possibility of leading to generalized solutions for some network types.

3 MODEL

3.1 Problem Formulation

In our setting we have a cake depicted as an interval $C = [0, 1]$ and n agents. Each agent possesses a valuation density function, $u_i : [0, 1] \to \mathbb{R}$ which is continuous. A piece of C 'S' is defined to be a union of finitely many disjoint intervals, i.e $S = \cup_j s_j$ where

$s_j \subseteq [0,1]$ is an interval and $\forall i, j,\ i \neq j$, we have $s_j \cap s_i = \phi$. The valuation of agent i for a piece S is defined to be $v_i(S) = \int_S u_i(x)dx$. Our assumptions for each v_i:
Additive -$\forall\, x, y \subseteq [0,\ 1],\ v_i(x \cup y) + v_i(x \cap y) = v_i(x) + v_i(y)$,
Divisible-$\forall\, x \subseteq [0,\ 1], \lambda \in [0, 1], \exists y \subseteq x\, s.t.\, v_i(y) = \lambda v_i(x)$.
Furthermore, we assume each v_i is normalized, i.e. $v_i([0,1]) = 1$. Our goal is to divide the cake into n disjoint pieces $A = (A_1, A_2, ...A_n)$ such that agent i gets the piece A_i and $\cup_i A_i = [0,1]$.

Definition 3.1. (Envy Freeness) For a cake-division instance, an allocation A is said to be envy-free if we have for all agents $i, j \in [n]$, $v_i(A_i) \geq v_i(A_j)$.

Definition 3.2. (Envy Freeness on Networks) Assuming a simple graph $G = (V, E)$ where each vertex denotes agent i and $N(i)$ denotes the set of neighbours of agent i. An allocation A is called envy-free on a network $G = (V, E)$ if for $\forall i$ *and* $\forall j \in N(i), v_i(A_i) \geq v_i(A_j)$. An important observation to make is that an allocation that is envy-free over a graph G, is also envy-free over any subgraph $G' \subseteq G$.

Special Graph Structures: We will be looking at envy freeness over cycle graphs C_n particularly C_4, C_5, C_6. Apart from this we will also focus on cliques of size 3 connected by a bridge and some structures involving trees.

3.2 Subroutines

3.2.1 Austin Cut. The Austin moving-knife procedures [1] are procedures for equitable division of a cake. Define the subroutine for 2 players and m equal division as $AustinCut(i, j, m, S)$-partitions S into m parts, s.t. for every piece $P, v_i(P) = \frac{v_i(S)}{m}$ and $v_j(P) = \frac{v_j(S)}{m}$ both hold.

Algorithm 1 $AustinCut(Player_1, Player_2, m, S)$

```
Require: A = (A_1, A_2, ...A_m)
  i = 0
  Player_1 marks S into m equal pieces (according to v_i).
  while S is unallocated do
      i = i + 1
      if ∃ piece P s.t. v_2(P) = 1/m then
          A_i = P
      else
          Find adjacent pieces A, B s.t v_2(A) < 1/m, v_2(B) > 1/m.
          C = A ∪ B
          l = x                              ▷ x=left endpoint of C
          r = y                    ▷ y=point s.t v_1([l, r]) = 1/m
          while v_2([l, r]) ≠ 1/m do
              move l, r s.t l < r, v_1([l, y]) = 1/m
                               ▷ By IVT, ∃l, r,  v_2([l, r]) = 1/m
          end while
          A_i = [l, r]
          S = S\{A_i}
      end if
  end while
```

This procedure requires atmost $2(m-1)$ cuts - $m-1$ cuts when $Player_1$ marks S into equal pieces, and atmost $m-1$ more cuts (1 cut for every time the **else** block is called), since after $m-1$ pieces are allocated the last piece P will have $v_j(P) = 1/m$.

3.2.2 Brams–Taylor–Zwicker procedure. It is a protocol ([6]) for envy-free cake-cutting among 4 partners. The procedure is represented by an algorithm shown below. The run-time of the procedure is, technically, infinite as Austin's procedure cannot be discretized, as it involves two knives moving continuously, resulting in an infinite run-time. However, the number of cuts required for the procedure is limited. Applying the Austin cut between 2 people to obtain 4 pieces in **Step A**, for example, requires six cuts. Additionally, **Step B** requires one cut, and **Step C** requires six more cuts, summing up to 13 cuts. An advanced variant of the Brams-Taylor-Zwicker procedure only requires up to 11 cuts. In Step B, WLOG we let player 3 trim his best piece to create a tie with his second best. We establish a ChooseOrder to ensure that the trimmed piece either goes to player 3 or 4. Step C involves an Austin Cut of the trimming T with player 1 and the player out of 3 and 4 who didn't get the trimmed piece, into 4 more pieces. We then let the player with trimmed piece choose first, since the remaining player 2 will never envy it even if all of T is given to him.

Algorithm 2 $BTZ(1, 2, 3, 4, S)$

```
  AustinCut(1, 2, 4, S)                                  ▷ Step A
  T = Trim(3)      ▷ Step B: Player 3-two way tie for largest piece
  ChooseOrder(4, 3, 2, 1, S\{T})                 ▷ Ensure T → 3 or 4
  if T → 3 then                                          ▷ Step C
      AustinCut(4, 1, 4, T)
      ChooseOrder(3, 2, 4, 1, T)
  else
      AustinCut(3, 1, 4, T)
      ChooseOrder(4, 2, 3, 1, T)
  end if
```

4 STRUCTURES

In this section, we provide cake-cutting procedures using the ideas of Austin cut and BTZ [6] trimming procedure as described before that satisfy network-envy-free property on the following structures: (1) Binary Tree with an additional level 1 edge (2) Cycles (C_4, C_5, C_6) and (3) clique (3 vertices) connected via a bridge. In addition, we also provide the reader with the number of cuts required for each of these structures. Finally, we close this section with a mention of other miscellaneous structures for which we have a network-envy-free procedure.

4.1 Binary Tree with an additional level 1 edge

- Step 1: Perform an Austin cut w.r.t two agents connected at level 1 on the initial cake P to produce n pieces where n is the number of agents. WLOG assume the agents performing the Austin cut are A and B. As a result of the Austin cut, both A and B value each of the above pieces equally according to them.
- Step 2: Let R, the root agent of the tree(parent of A, B), choose its most preferred piece and then let A choose its best $n_A = n - 2$ pieces from the left $n - 1$ pieces where n_A

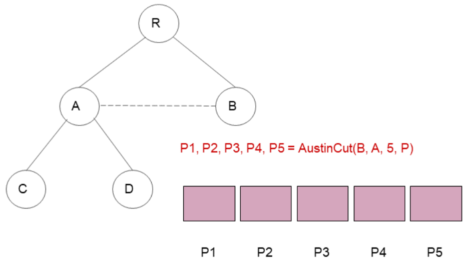

Figure 1: Step1: Binary tree with an additional edge at level 1

represents the number of agents in the subtree of A. Give the leftover piece to B.

- Step 3: Follow the same procedure[4] here onwards for the usual tree structure with agent A as the root and $P = \cup_{i \in n_A} P_i$.

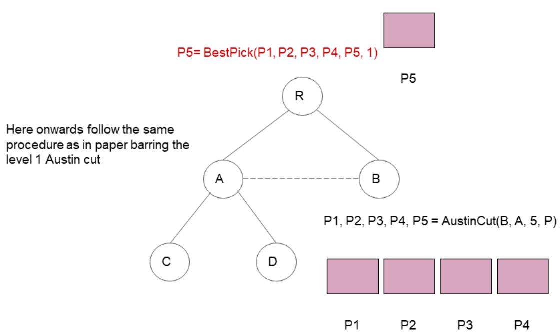

Figure 2: Step2: Binary tree with an additional edge at level 1

4.2 Cycles

4.2.1 *Length four cycle: C_4.*

- Step 1: Perform an Austin cut w.r.t any two agents who are maximally separated (separated via maximum number of edges) on the initial cake P to produce 4 pieces P_1, P_2, P_3, P_4. WLOG assume the agents performing the Austin cut are A and C. As a result of the Austin cut, both A and C value each of the above pieces equally according to them.
- Step 2: Let B choose its most preferred piece and then let D choose its most preferred piece among the remaining.
- Step 3: Finally, the two leftover pieces are arbitrarily allocated to A and C.

4.2.2 *Length five cycle: C_5.*

Phase 1:

- Step 1: Perform an Austin cut w.r.t any two agents who are maximally separated (separated via maximum number of edges) on the initial cake P to produce 5 pieces P_1, P_2, P_3, P_4, P_5. WLOG assume the agents performing the Austin cut are A and C. As a result of the Austin cut, both A and C value each of the above pieces equally according to them.
- Step 2: Let B choose its most preferred piece. WLOG assume B picks P_1.
- Step 3: Given the remaining four pieces, let either D or E trim its most preferred piece to equate its valuation to its second-best piece. This essentially creates a two-way tie for the best piece for the agent making the trim just as in the Selfridge-Conway procedure [7]. Assume that WLOG D performs the trim on its most preferred piece P_2 which gives T_1, T_2, and consequently, according to D the valuation of T_2 is equal to the valuation of P_3(its second best preference). Additionally, note that the trimmed piece T_1 is kept aside for now and will be allocated later in the next phase.
- Step 4: Let E choose its most preferred piece among T_2, P_3, P_4, P_5. If E chooses T_2, then D gets P_3, else, D gets T_2. WLOG assume D gets the trimmed piece.
- Step 5: Finally, the two leftover pieces are arbitrarily allocated to A and C.

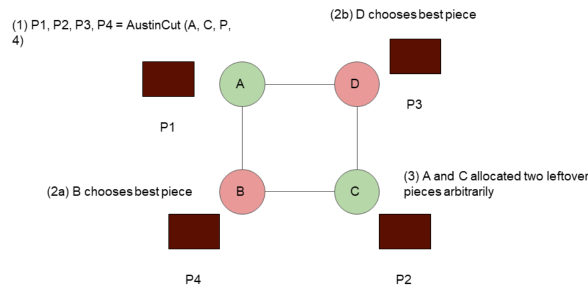

Figure 3: Cake cutting procedure for C_4

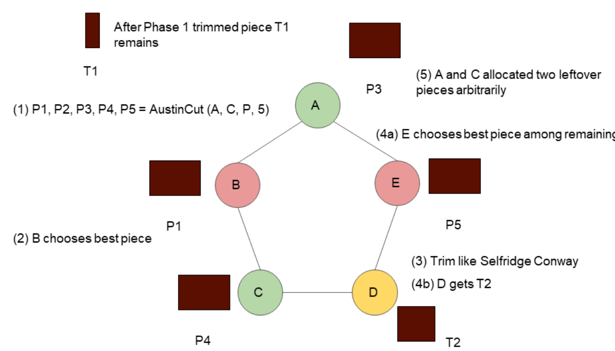

Figure 4: Cake cutting procedure for C_5: Phase 1

After Phase 1, trimmed piece T_1 remains to be allocated. However, note that under the current interim allocation agents are network envy-free of others' pieces just as in the Brams-Taylor-Zwicker procedure [6]. Additionally, we now have agent D who has been allocated the trimmed piece. As a result of this, agent C(D's neighbor and one of the agents who performed the Austin cut in phase 1) will never envy D even if the whole of T_1 were allocated to it. This change in envy structure w.r.t D (we call it the *anchor agent*) after phase 1 helps us stop the further trimming of the trimmed piece T_1 just as in again the Brams-Taylor-Zwicker procedure [6].

Phase 2:

- Step 1: Perform an Austin cut w.r.t agents B and E (Agents who were not part of either the Austin cut or did not get the trimmed piece T_2) on the trimmed piece T_1 to produce 5 pieces Q_1, Q_2, Q_3, Q_4, Q_5. As a result of the Austin cut, both B and E value each of the above pieces equally according to them.
- Step 2: Let A choose its most preferred piece. WLOG assume A picks Q_1.
- Step 3: Given the remaining four pieces and the fact that D is an anchor agent, we can let D choose its best pick and then let agent C choose its best piece among the remaining.
- Step 4: Finally, the two leftover pieces are allocated to B and E arbitrarily.

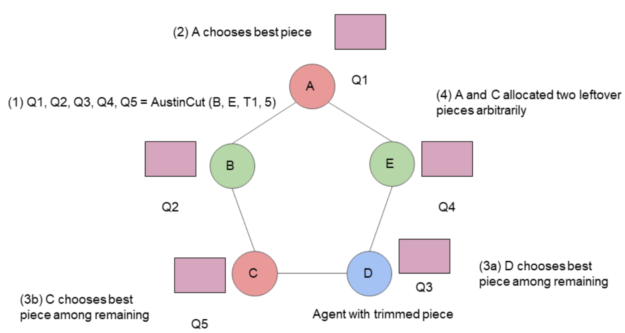

Figure 5: Cake cutting procedure for C_5: Phase 2

4.3 Cliques connected via a bridge

4.3.1 Cliques with three vertices connected via a bridge.

Phase 1:

- Step 1: Perform an Austin cut w.r.t agents A and D (agents part of the bridge) on the initial cake P to produce 6 pieces $P_1, P_2, P_3, P_4, P_5, P_6$. As a result of the Austin cut, both A and D value each of the above pieces equally according to them.
- Step 2: Given the six pieces, let either B or C trim its most preferred piece to equate its valuation to its second-best piece. This essentially creates a two-way tie for the best piece for the agent making the trim just as in the Selfridge-Conway procedure [7]. Assume that WLOG B performs the trim on its most preferred piece P_1 which gives T_1, T_2, and consequently, according to B the valuation of T_2 is equal to the valuation of P_2(its second best preference). Additionally, note that the trimmed piece T_1 is kept aside for now and will be allocated later in the next phase.
- Step 3: Let C choose its most preferred piece among $T_2, P_2, P_3, P_4, P_5, P_6$. If C chooses T_2, then B gets P_2, else, B gets T_2.
- Let's assume WLOG B gets T_2 and C gets P_2. The remaining four pieces P_3, P_4, P_5, P_6 are now to be fully allocated among A, D, E, and F such that the allocation is network-envy-free. We do this on a case basis as explained below.
- Step 4: We consider the following three cases which are exhaustive:

(1) Case 1: Top preferences of agents E and F among the remaining pieces are different. In this case, we directly allocate agents E and F their best choices, and the two leftover pieces are allocated to A and D arbitrarily.
(2) Case 2: Top preferences of agents E and F among the remaining pieces are the same (say P_3), however, their second most preferred piece is different. In this case, we allocate P_3 to agent A, allocate agents E and F their second best choices, and the remaining piece is given to D.
(3) Case 3: Top two preferences of agents E and F among the remaining pieces are the same (say P_3, P_4). In this case, we perform an Austin cut w.r.t agent D and WLOG E (any of E or F) on $P_3 \cup P_4$ to produce two new pieces K_1, K_2. As a result of the Austin cut, both D and E value each of the above two pieces equally according to them. Next, we let agent F choose its best pick with the other going to E. Finally, the two leftover pieces P_5, P_6 are arbitrarily allocated to A and D.

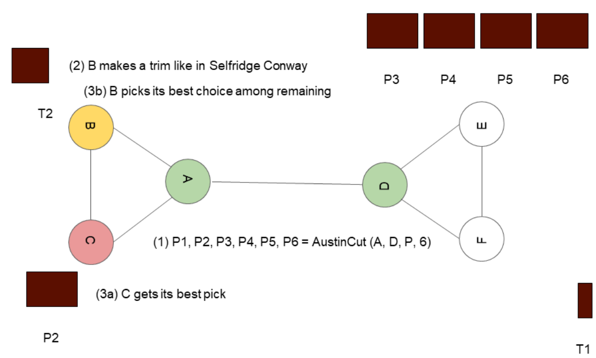

Figure 6: Cake cutting procedure for clique (3 vertices) connected via a bridge: phase 1.

After Phase 1, trimmed piece T_1 remains to be allocated. However, note that under the current interim allocation, agents are network envy-free of others' pieces just as in the Brams-Taylor-Zwicker procedure [6]. Additionally, we now have agent B who has been allocated the trimmed piece. As a result of this, agent A(B's neighbor and one of the agents who performed the Austin cut in phase 1) will never envy B even if the whole of T_1 were allocated to it. This change in envy structure w.r.t B (the *anchor agent*) after phase 1 helps us stop the further trimming of the trimmed piece T_1 just as in again the Brams-Taylor-Zwicker procedure [6].

Phase 2:

- Step 1: Perform an Austin cut w.r.t agents C and D on the trimmed piece T_1 to produce 6 pieces $Q_1, Q_2, Q_3, Q_4, Q_5, Q_6$.

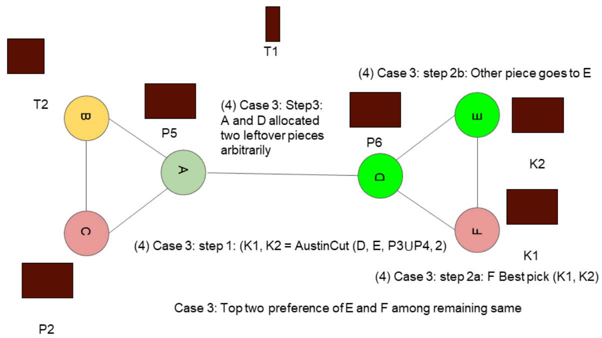

Figure 7: Cake cutting procedure for clique (3 vertices) connected via a bridge: phase 1 case 3.

As a result of the Austin cut, both C and D value each of the above pieces equally according to them.

- Step 2: Given the six pieces and the fact that B is an anchor agent, we can let B choose its best pick and then let agent A choose its best piece among the remaining.
- Let's assume WLOG B gets Q_1 and A gets Q_2. The remaining four pieces Q_3, Q_4, Q_5, Q_6 are now to be fully allocated among C, D, E, and F such that the allocation is network envy-free. We do this on a case basis exactly as done in phase 1.

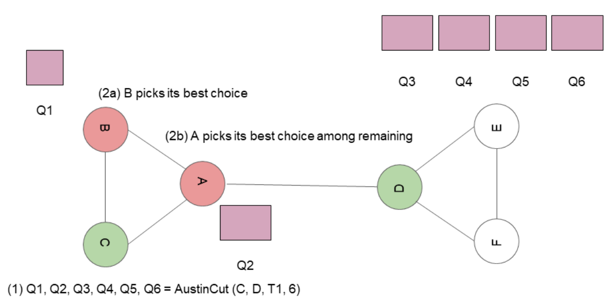

Figure 8: Cake cutting procedure for clique (3 vertices) connected via a bridge: phase 2.

4.4 Comparison of number of cuts required

Structures	Number of cuts required in worst case
C_4	**6**
C_5	**17**
C_6	**108** (**54** for each trim)
Cliques connected via bridge	**25**
Tree with additional edge at level1	$O(n^2)$

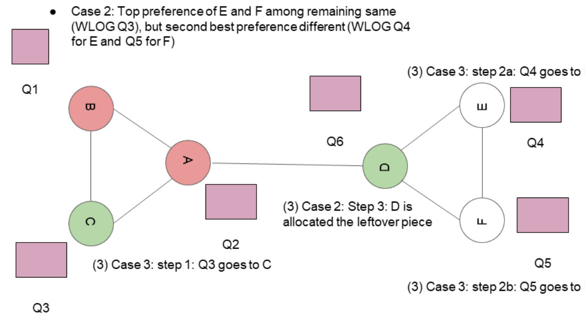

Figure 9: Cake cutting procedure for clique (3 vertices) connected via a bridge: phase 2 case 2.

4.5 Other Miscellaneous structures

In addition to the above-described structures, we have procedures that satisfy network-envy-freeness for the following structures as well. However, we skip these due to space limitation.

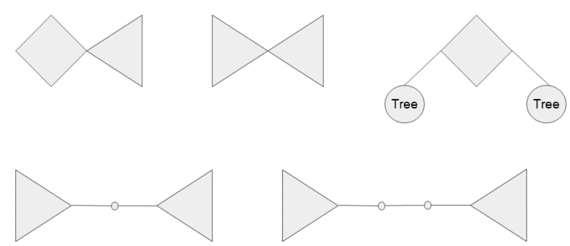

Figure 10: Other Miscellaneous structures.

5 CONCLUSIONS

General algorithm for C_n - For cycle networks, we have algorithms upto C_6.

Extensions to [4] - Our extension only added 1 edge. We currently have an algorithm for adding leaf edges, but it involves wastage.

Other fairness notions - Proportionality is another popular notion of fairness that is studied along with envy-freeness. It is weaker and hence easier to obtain, but is harder to generalize over several networks. Nevertheless, we can try to study and improve upon current proportional and efficient algorithms. We can also try to define and study other notions of networked fairness.

REFERENCES

[1] A.K. Austin. 1982. Sharing a Cake. *The Mathematical Gazette* 66, 437 (1982), 212–215. https://doi.org/10.2307/3616548

[2] Haris Aziz and Simon Mackenzie. 2016. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 416–427.

[3] Haris Aziz and Simon Mackenzie. 2016. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 454–464.

[4] Xiaohui Bei, Youming Qiao, and Shengyu Zhang. 2017. Networked fairness in cake cutting. *arXiv preprint arXiv:1707.02033* (2017).
[5] Xiaohui Bei, Xiaoming Sun, Hao Wu, Jialin Zhang, Zhijie Zhang, and Wei Zi. 2020. Cake cutting on graphs: a discrete and bounded proportional protocol. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2114–2123.
[6] Steven Brams, Alan Taylor, and William Zwicker. 1997. A moving-knife solution to the four-person envy-free cake-division problem. *Proceedings of the american mathematical society* 125, 2 (1997), 547–554.
[7] Steven J Brams and Alan D Taylor. 1995. An envy-free cake division protocol. *The American Mathematical Monthly* 102, 1 (1995), 9–18.
[8] Ganesh Ghalme, Xin Huang, Yuka Machino, and Nidhi Rathi. 2022. A Discrete and Bounded Locally Envy-Free Cake Cutting Protocol on Trees. *arXiv preprint arXiv:2211.06458* (2022).
[9] Ganesh Ghalme, Xin Huang, and Nidhi Rathi. 2022. Envy-Free Cake Cutting with Graph Constraints. *arXiv preprint arXiv:2205.12559* (2022).

6 APPENDIX

6.1 Cycles

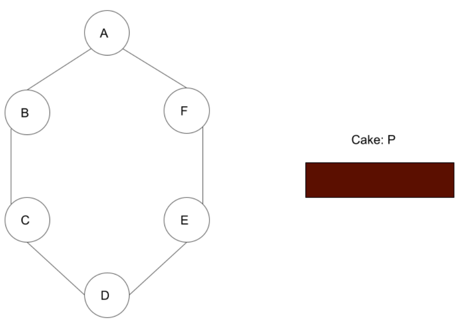

Figure 11: Structure C_6

6.1.1 *Length six cycle: C_6.*

Phase 1:

- Step 1: Perform an Austin cut w.r.t any two agents who are maximally separated (separated via maximum number of edges) on the initial cake P to produce 6 pieces $P_1, P_2, P_3, P_4, P_5, P_6$. WLOG assume the agents performing the Austin cut are A and D. As a result of the Austin cut, both A and D value each of the above pieces equally according to them.
- Step 2: Given the six pieces, let either B or C trim its most preferred piece to equate its valuation to its second-best piece. This essentially creates a two-way tie for the best piece for the agent making the trim just as in the Selfridge-Conway procedure [7]. Assume that WLOG B performs the trim on its most preferred piece P_1 which gives T_1, T_2, and consequently, according to B the valuation of T_2 is equal to the valuation of P_2(its second best preference). Additionally, note that the trimmed piece T_1 is kept aside for now and will be allocated later in the next phase.
- Step 3: Let C choose its most preferred piece among $T_2, P_2, P_3, P_4, P_5, P_6$. If C chooses T_2, then B gets P_2, else, B gets T_2. Assume WLOG that B gets T_2 and C gets P_2.
- Step 4: Given the remaining four pieces, let either E or F trim its most preferred piece to equate its valuation to its second-best piece. This essentially creates a two-way tie for the best piece for the agent making the trim just as in the Selfridge-Conway procedure [7]. Assume that WLOG E performs the trim on its most preferred piece P_3 which gives T_3, T_4, and consequently, according to E the valuation of T_4 is equal to the valuation of P_4(its second best preference). Additionally, note that the trimmed piece T_3 is kept aside for now and will be allocated later in the next phase.
- Step 5: Let F choose its most preferred piece among T_4, P_4, P_5, P_6. If F chooses T_4, then E gets P_4, else, E gets T_4. Assume WLOG that F gets T_4 and E gets P_4.
- Step 6: Finally, the two leftover pieces are arbitrarily allocated to A and D.

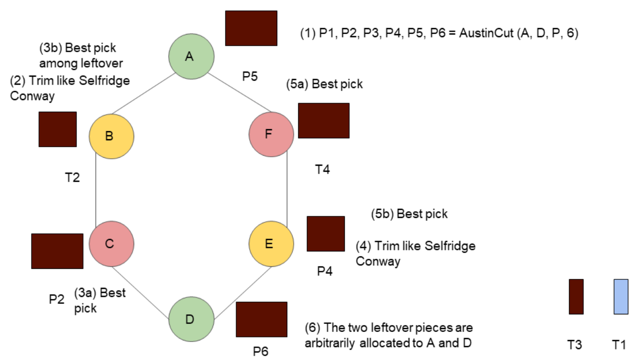

Figure 12: Structure C_6: Phase 1

After Phase 1, trimmed piece T_1 and T_3 remain to be allocated. However, note that under the current interim allocation agents are network envy-free of others' pieces just as in the Brams-Taylor-Zwicker procedure [6]. Additionally, we now have agent B(F) as anchor agent w.r.t A (neighbor and one of the agents who performed the Austin cut in phase 1) for trimmed piece $T_1(T_3)$. As a result of this, agent A will never envy B(F) even if the whole of $T_1(T_3)$ were allocated to it. This change in envy structure w.r.t B(F) after phase 1 helps us stop the further trimming of the trimmed piece $T_1(T_3)$ eventually.

Phase 2 for trimmed piece T_1:

- Step 1: Perform an Austin cut w.r.t agents C and F on the trimmed piece T_1 to produce 6 pieces $Q_1, Q_2, Q_3, Q_4, Q_5, Q_6$. As a result of the Austin cut, both C and F value each of the above pieces equally according to them.
- Step 2: Given the six pieces and the fact that B is an anchor agent, we can let B choose its best pick and then let agent A choose its best piece among the remaining. WLOG assume B chose Q_1 and A chose Q_2.
- Step 3: Given the remaining four pieces, let either D or E trim its most preferred piece to equate its valuation to its second-best piece. This essentially creates a two-way tie for the best piece for the agent making the trim just as in the Selfridge-Conway procedure [7]. Assume that WLOG D performs the trim on its most preferred piece Q_3 which gives T_5, T_6, and consequently, according to D the valuation of T_6 is equal to the valuation of Q_4(its second best preference).

Additionally, note that the trimmed piece T_5 is kept aside for now and will be allocated later in the next phase.

- step 4: Let E choose its most preferred piece among T_6, Q_4, Q_5, Q_6. If E chooses T_6, then D gets Q_4, else, D gets T_6.
- Now there are two cases possible: (1) E gets the trimmed piece T_6 or (2) D gets it. Either way, we can proceed with phase 2 step 5 below, however, the above cases become significant for phase 3 which we cover shortly.
- Step 5: The two leftover pieces are arbitrarily allocated to C and F.

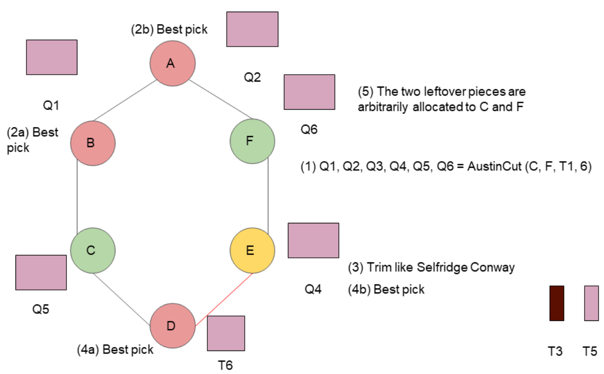

Figure 13: Structure C_6: Phase 2 for T_1

Phase 3 for trimmed piece T_5 ($T_1 \rightarrow T_5$:)

Case 1: D got the trimmed piece T_6 in phase 2 for T_1. After "Phase 2 for trimmed piece T_1", trimmed piece T_5 remains to be allocated. However, note that under the current interim allocation agents are network envy-free of others' pieces just as in the Brams-Taylor-Zwicker procedure [6]. Additionally, we now have agent B as anchor agent w.r.t A (neighbor and one of the agents who performed the Austin cut in phase 1) for trimmed piece T_5 and agent D as anchor agent w.r.t C (neighbor and one of the agents who performed the Austin cut in phase 2) for trimmed piece T_5. As a result of this, agent A(C) will never envy B(D) even if the whole of T_5 were allocated to it. This change in the envy structure of B and D helps us stop the further trimming of the trimmed piece T_5.

- Step 1: Perform an Austin cut w.r.t agents E and F on the trimmed piece T_5 to produce 6 pieces $R_1, R_2, R_3, R_4, R_5, R_6$. As a result of the Austin cut, both E and F value each of the above pieces equally according to them.
- Step 2: The fact that D is an anchor agent w.r.t C and its other neighbour is agent E which did the Austin cut, we can let D choose its best pick. WLOG assume D picked R_1.
- Since B is anchor agent w.r.t A, we can let it pick before A. However, note B is not an anchor w.r.t C, thus, we resolve the dispute between B and C on a case basis as described below.
- Step 3: We consider the following three cases which are exhaustive:

(1) Case 1: Top preferences of agents B and C among the remaining pieces are different. In this case, we directly allocate agents B and C their best choices, let A pick its best choice among the remaining, and the two leftover pieces are allocated to E and F arbitrarily.

(2) Case 2: Top preferences of agents B and C among the remaining pieces are the same (say R_2), however, their second most preferred piece is different. In this case, we allocate R_2 to agent E, allocate agents B and C their second best choices, then let A pick its best choice among the remaining, and the leftover piece is allocated to F.

(3) Case 3: Top two preferences of agents B and C among the remaining pieces are the same (say R_2, R_3). In this case, we perform an Austin cut w.r.t agent C and D on $R_3 \cup R_2$ to produce two new pieces K_1, K_2. As a result of the Austin cut, both C and D value each of the above two pieces equally according to them. Next, we let agent B choose its best pick with the other going to C. Finally, we let A pick its best choice among the remaining, and the two leftover pieces are allocated to E and F arbitrarily.

(4) Therefore, in the case where T_6 went to D, T_5 is fully allocated without further trimming, while at the same time ensuring network envy-freeness.

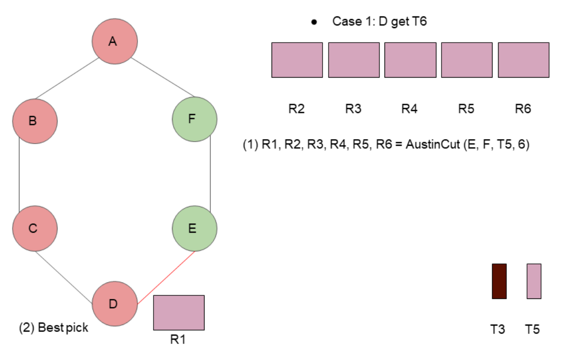

Figure 14: Structure C_6: Phase 3 for trimmed piece T_5: Case 1 (step 1 and 2)

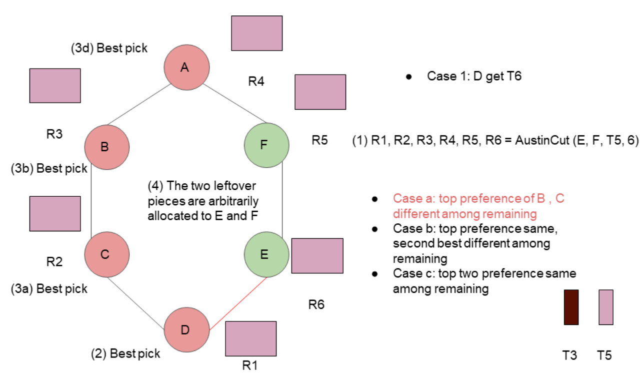

Figure 15: Structure C_6: Phase 3 for trimmed piece T_5: Case 1: Step 3: Sub-case 1

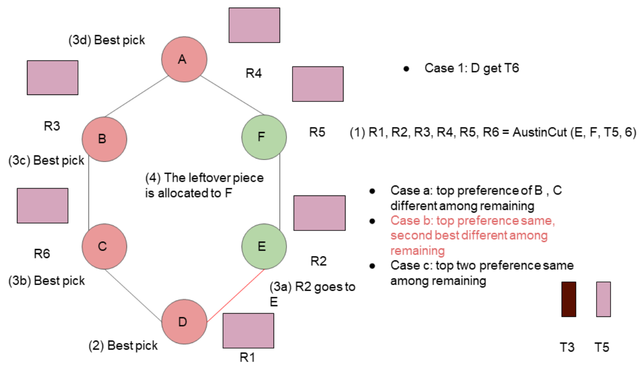

Figure 16: Structure C_6: Phase 3 for trimmed piece T_5: Case 1: Step 3: Sub-case 2

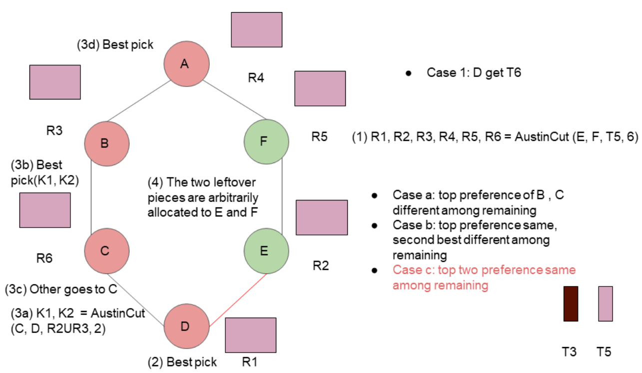

Figure 17: Structure C_6: Phase 3 for trimmed piece T_5: Case 1: Step 3: Sub-case 3

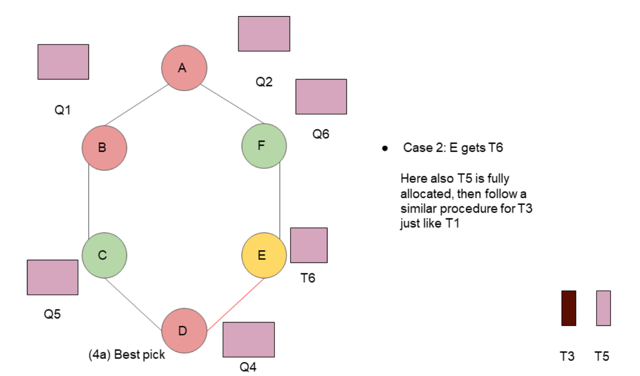

Figure 18: Structure C_6: Phase 3 for trimmed piece T_5: Case 2

Case 2: E got the trimmed piece T_6 in phase 2 for T_1. After "Phase 2 for trimmed piece T_1", trimmed piece T_5 remains to be allocated. However, note that under the current interim allocation agents are network envy-free of others' pieces just as in the Brams-Taylor-Zwicker procedure [6]. Additionally, we now have agent B as anchor agent w.r.t A (neighbor and one of the agents who performed the Austin cut in phase 1) for trimmed piece T_5 and agent E as anchor agent w.r.t F (neighbor and one of the agents who performed the Austin cut in phase 2) for trimmed piece T_5. As a result of this, agent A(F) will never envy B(E) even if the whole of T_5 were allocated to it. This change in the envy structure of B and E help us stop the further trimming eventually.

- Step 1: Perform an Austin cut w.r.t agents B and E on the trimmed piece T_5 to produce 6 pieces $R_1, R_2, R_3, R_4, R_5, R_6$. As a result of the Austin cut, both B and E value each of the above pieces equally according to them.
- Step 2: Given the six pieces, let either A or F trim its most preferred piece to equate its valuation to its second-best piece. This essentially creates a two-way tie for the best piece for the agent making the trim just as in the Selfridge-Conway procedure [7]. Assume that WLOG A performs the trim on its most preferred piece R_1 which gives T_7, T_8, and consequently, according to A the valuation of T_8 is equal to the valuation of R_2(its second best preference). Additionally, note that the trimmed piece T_7 is kept aside for now and will be allocated later in the next phase.
- Step 3: Let F choose its most preferred piece among $T_8, R_2, R_3, R_4, R_5, R_6$. If F chooses T_8, then A gets R_2, else, A gets T_8. WLOG F chooses R_2 and A gets the trimmed piece.
- Step 4: Given the four remaining pieces, let either C or D trim its most preferred piece to equate its valuation to its second-best piece. This essentially creates a two-way tie for the best piece for the agent making the trim just as in the Selfridge-Conway procedure [7]. Assume that WLOG C performs the trim on its most preferred piece R_3 which gives T_9, T_{10}, and consequently, according to C the valuation of T_{10} is equal to the valuation of R_4(its second best preference). Additionally, note that the trimmed piece T_9 is kept aside for now and will be allocated later in the next phase.
- Step 5: Let D choose its most preferred piece among T_{10}, R_4, R_5, R_6. If D chooses T_{10}, then C gets R_4, else, C gets T_{10}. WLOG D chooses R_5 and C gets the trimmed piece T_{10}.
- Step 6: Finally, the two leftover pieces are arbitrarily allocated to B and E.

Note that, for case 2 trimmed pieces T_7 and T_9 remain to be allocated.

Phase 4 for trimmed piece T_7($T_1 \rightarrow T_5 \rightarrow T_7$:) In addition to anchors B(w.r.t A) and E(w.r.t F), we now have a newly created anchor agent A(w.r.t B) for trimmed T_7.This stops the further trimming of the trimmed piece T_7.

- Step 1: Perform an Austin cut w.r.t C and E on the trimmed piece T_7 to produce 6 pieces $S_1, S_2, S_3, S_4, S_5, S_6$. As a result of the Austin cut, both E and C value each of the above pieces equally according to them.
- Step 2: Let D choose its most preferred piece. WLOG assume D picked S_1.
- Step 3: The fact that B is an anchor agent w.r.t A and its other neighbour is agent C which did the Austin cut, we can let B choose its best pick. WLOG assume B picked S_2.
- Now, note that A and F are not anchors w.r.t each other, thus, we resolve the dispute between A and F on a case basis as described below.
- Step 3: We consider the following three cases which are exhaustive:

(1) Case 1: Top preferences of agents A and F among the remaining pieces are different. In this case, we directly allocate agents A and F their best choices, and the two leftover pieces are allocated to E and C arbitrarily.
(2) Case 2: Top preferences of agents A and F among the remaining pieces are the same (say S_3), however, their second most preferred piece is different. In this case, we allocate S_3 to agent C, allocate agents A and F their second best choices, and the leftover piece is allocated to E.
(3) Case 3: Top two preferences of agents A and F among the remaining pieces are the same (say S_3, S_4). In this case, we perform an Austin cut w.r.t agent E and F on $S_3 \cup S_4$ to produce two new pieces L_1, L_2. As a result of the Austin cut, both F and E value each of the above two pieces equally according to them. Next, we let agent A choose its best pick with the other going to F. Note that, we can let A choose its best piece without worrying about agent B's envy towards A due to the fact that A is an anchor agent w.r.t B. Finally, the two leftover pieces are arbitrarily allocated to E and C.
(4) Therefore, in the case where T_6 went to E, T_5 is fully allocated without further trimming ($T_1 \longrightarrow T_5 \longrightarrow T_7$), while at the same time ensuring network envy-freeness.

Phase 4 for trimmed piece T_9($T_1 \longrightarrow T_5 \longrightarrow T_9$:) In addition to anchors B(w.r.t A) and E(w.r.t F), we now have a newly created anchor agent C(w.r.t B) for trimmed T_9.This stops the further trimmed of the trimmed piece T_9.

- Step 1: Perform an Austin cut w.r.t D and F on the trimmed piece T_9 to produce 6 pieces $U_1, U_2, U_3, U_4, U_5, U_6$. As a result of the Austin cut, both D and F value each of the above pieces equally according to them.
- Step 2: Let E choose its most preferred piece. WLOG assume E picked U_1.
- Step 3: Because of the fact that C is an anchor w.r.t B and B is an anchor w.r.t A, we can let C pick first, followed by B and then A.
- Step 4: Finally, the two leftover pieces are arbitrarily allocated to D and F.
- Therefore, in the case where T_6 went to E, T_5 is fully allocated without further trimming($T_1 \longrightarrow T_5 \longrightarrow T_9$), while at the same time ensuring network envy-freeness.

With this, the Trimmed piece T_1 from phase 1 is fully allocated without trimming in addition to ensuring network envy-freeness. We do the exact same procedure for T_3, except with a different set of anchor agents. We skip the procedure for this piece due to space limitations and close the procedure here.