

On the Fixed-Parameter Tractability of Composition-Consistent Tournament Solutions*

Felix Brandt Markus Brill Hans Georg Seedig

Technische Universität München
85748 Garching bei München, Germany
{brandtf,brill,seedigh}@in.tum.de

Abstract

Tournament solutions, *i.e.*, functions that associate with each complete and asymmetric relation on a set of alternatives a non-empty subset of the alternatives, play an important role within social choice theory and the mathematical social sciences at large. Laffond *et al.* have shown that various tournament solutions satisfy *composition-consistency*, a structural invariance property based on the similarity of alternatives. We define the *decomposition degree* of a tournament as a parameter that reflects its decomposability and show that computing any composition-consistent tournament solution is fixed-parameter tractable with respect to the decomposition degree. Furthermore, we experimentally investigate the decomposition degree of two natural distributions of tournaments and its impact on the running time of computing the tournament equilibrium set.

1 Introduction

Many problems in multiagent decision making can be addressed using tournament solutions, *i.e.*, functions that associate with each complete and asymmetric relation on a set of alternatives a non-empty subset of the alternatives. Tournament solutions are most prevalent in social choice theory, where the binary relation is typically assumed to be given by the simple majority rule (*e.g.*, Moulin, 1986; Laslier, 1997). Other application areas include multi-criteria decision analysis (*e.g.*, Arrow and Raynaud, 1986), zero-sum games (*e.g.*, Fisher and Ryan, 1995), coalition formation (*e.g.*, Brandt and Harrenstein, 2010), and argumentation theory (*e.g.*, Dunne, 2007).

Recent years have witnessed an increasing interest in the computational complexity of tournament solutions by the multiagent systems and theoretical computer science commu-

nities. A number of concepts such as the Banks set (Woeginger, 2003), the Slater set (Alon, 2006; Conitzer, 2006), and the tournament equilibrium set (Brandt *et al.*, 2010). have been shown to be computationally intractable. For others, including the minimal covering set and the bipartisan set, algorithms that run in polynomial time have been provided (Brandt and Fischer, 2008). The class of all tournaments is excessively rich and it is well-known that only a fraction of these tournaments occur in realistic settings (see, *e.g.*, Feld and Grofman, 1992). Therefore, an important question is whether there are natural distributions of tournaments that admit more efficient algorithms for computing specific tournament solutions.

In this paper, we study tournaments that are *decomposable* in a natural well-defined way. A set of alternatives forms a *component* if all alternatives in this set bear the same relationship to all outside alternatives. Elements of a component can thus be seen as variants of the same type of an alternative. Laslier (1997) has shown that every tournament admits a unique natural decomposition into components, which may themselves be decomposable into subcomponents. A tournament solution is *composition-consistent* if it chooses the best alternatives of the best components (Laffond *et al.*, 1996).¹ In other words, a composition-consistent tournament solution can be computed by recursively determining the winning components. All of the tournament solutions mentioned earlier except the Slater set are composition-consistent.

In this paper, we provide a formalization of the recursive decomposition of tournaments and a detailed analysis of the speed-up that can be achieved when computing composition-consistent tournament solutions. In particular, we define the *decomposition degree* of a tournament as a parameter that reflects its decomposability. Intuitively, a low decomposition degree indicates that the tournament admits a particularly well-behaved decomposition and therefore allows the efficient computation of composition-consistent tournament solutions. Within our analysis, we leverage a recently proposed linear-time algorithm for the modular decomposition of directed graphs (McConnell and de Montgolfier, 2005; Capelle *et al.*, 2002).

In related work, Betzler *et al.* (2010) proposed data reduc-

*This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/6-1 (within the European Science Foundation's EUROCORES program LogICCC) and BR 2312/7-1. The authors gratefully acknowledge the support of the TUM's Faculty Graduate Center CeDoSIA at Technische Universität München.

¹Composition-consistency is related to *cloning-consistency*, which was introduced by Tideman (1987) in the context of voting.

tion rules that facilitate the computation of *Kemeny rankings*. One of these rules, the “Condorcet-set rule”, corresponds to a (rather limited) special case of composition-consistency where tournaments are decomposed into exactly *two* components. Furthermore, a preprocessing technique that resembles the one proposed in this paper has been used by Conitzer (2006) to speed up the computation of *Slater rankings*. Interestingly, even though Slater’s solution is *not* composition-consistent, decompositions of the tournament can be exploited to identify a *subset* of the optimal rankings.

Our results, on the other hand, allow us to compute *complete choice sets* and are applicable to *all composition-consistent tournament solutions*, including the uncovered set, the minimal covering set, the bipartisan set, the Banks set, the tournament equilibrium set, and the minimal extending set (Laslier, 1997; Brandt, 2011). The former three admit polynomial-time algorithms whereas the latter three are computationally intractable. None of the concepts is known to admit a linear-time algorithm.

We show that computing any composition-consistent tournament solution is *fixed-parameter tractable* with respect to the decomposition degree of the tournament, *i.e.*, there are algorithms that are only superpolynomial in the decomposition degree. We conclude the paper with an experimental investigation of the decomposition degree and the actual running time of computing the tournament equilibrium set for two natural distributions of tournaments. The first one is a well-studied model that assumes the existence of a true linear ordering of the alternatives that has been perturbed by binary random inversions. The other one is a spatial voting model based on the proximity of voters and alternatives in a multi-dimensional space.

2 Preliminaries

In this section, we provide the terminology and notation required for our results (see Laslier (1997) for an excellent overview of tournament solutions and their properties).

2.1 Tournaments

A *tournament* T is a pair $(A, >)$, where A is a finite set of *alternatives* and $>$ is an asymmetric and complete (and thus irreflexive) binary relation on A , usually referred to as the *dominance relation*. Intuitively, $a > b$ signifies that alternative a is preferable to b . The dominance relation can be extended to sets of alternatives by writing $X > Y$ when $x > y$ for all $x \in X$ and $y \in Y$. The *order* $|T|$ of a tournament $T = (A, >)$ refers to its number of alternatives $|A|$. For a subset of alternatives $B \subseteq A$, let $T|_B$ denote the induced subtournament on B . Finally, a *tournament isomorphism* of two tournaments $T = (A, >)$ and $T' = (A', >')$ is a bijective mapping $\pi : A \rightarrow A'$ such that $a > b$ if and only if $\pi(a) >' \pi(b)$.

2.2 Components and Decompositions

An important structural concept in the context of tournaments is that of a *component*. A component is a subset of alternatives that bear the same relationship to all alternatives not in the set.

Definition 1 Let $T = (A, >)$ be a tournament and B . A non-empty subset B of A is a component of T if for all $a \in A \setminus B$ either $B > a$ or $a > B$. A decomposition of T is a set of pairwise disjoint components $\{B_1, \dots, B_k\}$ of T such that $A = \bigcup_{i=1}^k B_i$.

The *null decomposition* of a tournament $T = (A, >)$ is $\{A\}$; the *trivial decomposition* consists of all singletons of A . Any other decomposition is called *proper*. A tournament is said to be *decomposable* if it admits a proper decomposition. Given a particular decomposition, the *summary* of a tournament is defined as the tournament on the individual components rather than the alternatives.

Definition 2 Let $T = (A, >)$ be a tournament and $\tilde{B} = \{B_1, \dots, B_k\}$ a decomposition of T . The summary of T with respect to \tilde{B} is defined as $\tilde{T} = (\{1, \dots, k\}, \tilde{>})$, where

$$i \tilde{>} j \quad \text{if and only if} \quad B_i > B_j.$$

A tournament is called *reducible* if it admits a decomposition into *two* components. Otherwise, it is *irreducible*. Laslier (1997) has shown that there exist a natural unique way to decompose any tournament. Call a decomposition \tilde{B} *finer* than another decomposition \tilde{B}' if $\tilde{B} \neq \tilde{B}'$ and for each $B \in \tilde{B}$ there exists $B' \in \tilde{B}'$ such that $B \subseteq B'$. \tilde{B}' is said to be *coarser* than \tilde{B} . A decomposition is *minimal* if its only coarser decomposition is the null decomposition.

Proposition 1 (Laslier (1997)) Every irreducible tournament with more than one alternative admits a unique minimal decomposition.

This is obviously not true for *reducible* tournaments, as witnessed by the tournament $T = (\{1, 2, 3\}, >)$ with $1 > 2$, $1 > 3$, and $2 > 3$, which admits two minimal decompositions, namely $\{\{1\}, \{2, 3\}\}$ and $\{\{1, 2\}, \{3\}\}$. Nevertheless, there is a unique way to decompose any reducible tournament. A *scaling decomposition* is a decomposition with a transitive summary.

Proposition 2 (Laslier (1997)) Every reducible tournament admits a unique scaling decomposition such that each component is irreducible.

This scaling decomposition into irreducible components is also the finest scaling decomposition. In graph-theoretic terms, this decomposition partitions the tournament into its strongly connected components.

2.3 Tournament Solutions

Since the dominance relation may contain cycles and thus fail to have a maximal element, a variety of so-called tournament solutions have been suggested to take over the role of singling out the “best” alternatives of a tournament. Following Laslier (1997), we require a tournament solution to be invariant under tournament isomorphisms and to select the maximum (or *Condorcet winner*) whenever it exists.

Definition 3 A tournament solution is a function S that associates with each tournament $T = (A, >)$ a non-empty subset $S(T)$ of A such that

- (i) $S((\pi(A), >')) = \pi(S((A, >)))$ for all tournaments $(A, >)$, $(A', >')$, and every tournament isomorphism $\pi : A \rightarrow A'$ of $(A, >)$ and $(A', >')$; and
- (ii) $S(T) = \{a\}$ whenever there is some $a \in A$ such that $a > A \setminus \{a\}$.

A tournament solution is composition-consistent if it chooses the “best” alternatives from the “best” components (Laffond *et al.*, 1996).

Definition 4 A tournament solution S is composition-consistent if for all tournaments T and \tilde{T} such that \tilde{T} is the summary of T with respect to some decomposition $\{B_1, \dots, B_k\}$,

$$S(T) = \bigcup_{i \in S(\tilde{T})} S(T|_{B_i}).$$

2.4 Fixed-Parameter Tractability

We briefly introduce the most basic concepts of parameterized complexity theory (see, *e.g.*, Niedermeier, 2006). In contrast to classical complexity theory, where only the size of problem instances is taken into account, parameterized complexity allows for a more fine-grained analysis by considering arbitrary parameters of the instances. A problem with parameter k is said to be *fixed-parameter tractable* (or to belong to the class FPT) if there exists an algorithm that solves the problem in time $f(k) \cdot \text{poly}(|I|)$, where $|I|$ is the size of the input and f is some computable function independent of $|I|$.

For example, each (computable) problem is trivially fixed-parameter tractable with respect to the parameter $|I|$. The crucial point is to identify a parameter that is reasonably small in realistic instances and to devise an algorithm that is only superpolynomial in this parameter.

3 Decomposition Trees

Propositions 1 and 2 offer a straightforward method to iteratively decompose tournaments. If the tournament is reducible, take the finest scaling decomposition. If it is irreducible, take the minimal decomposition. The repeated application of these decompositions leads to the *decomposition tree* of a tournament.

Definition 5 The decomposition tree $D(T)$ of a tournament $T = (A, >)$ is defined as a rooted tree whose nodes are non-empty subsets of A . The root of $D(T)$ is A and for each node B with $|B| \geq 2$, the children of B are defined as follows:

- If $T|_B$ is reducible, the children of B are the components of the finest scaling decomposition of $T|_B$.
- If $T|_B$ is irreducible, the children of B are the components of the minimal decomposition of $T|_B$.

It follows from Propositions 1 and 2 that every tournament has a *unique* decomposition tree. By definition, each node in $D(T)$ is a component of T and each leaf is a singleton. However, not all components of T need to appear as nodes in $D(T)$. An example of a decomposition tree is provided in Figure 1.




Figure 1: Example tournament with corresponding decomposition tree. Nodes $\{f, c\}$ and $\{d, e\}$ are reducible, all other nodes are irreducible.

An internal (*i.e.*, non-leaf) node B of $D(T)$ with children B_1, \dots, B_k corresponds to the tournament $T_B = (\{1, \dots, k\}, \succ)$ where $i \succ j$ if and only if $B_i > B_j$, *i.e.*, T_B is the summary of $T|_B$ with respect to the decomposition $\{B_1, \dots, B_k\}$. The order of T_B is thus equal to the number of children of node B . Moreover, we call an internal node B *reducible* (respectively, *irreducible*) if the tournament T_B is reducible (respectively, irreducible).² If B is reducible, we assume without loss of generality that the children B_1, \dots, B_k are labelled according to their transitive summary, *i.e.*, $B_i > B_j$ if and only if $i < j$. In particular, the maximum of T_B is 1.

Recent results on the modular decomposition of directed graphs (Capelle *et al.*, 2002; McConnell and de Montgolfier, 2005) imply that the decomposition tree of a tournament can be computed in linear time.³

Proposition 3 The decomposition tree of a tournament T can be computed in time $O(|T|^2)$.

The proof consists of two steps. In the first step, a *factorizing permutation* of the tournament is constructed. A factorizing permutation of $T = (A, >)$ is a permutation of the alternatives in A such that each component of T is a contiguous interval in the permutation. McConnell and de Montgolfier (2005) provide a simple algorithm that computes a factorizing permutation of a tournament in linear time. Furthermore, there exists a fairly complicated linear-time algorithm by Capelle *et al.* (2002) that, given a tournament T and a factorizing permutation of T , computes the decomposition tree $D(T)$.

4 Computing Solutions via Decompositions

Let S be a composition-consistent tournament solution and consider an arbitrary tournament $T = (A, >)$ together with its decomposition tree $D(T)$. For an internal node B of

² $T|_B$ is reducible (respectively, irreducible) if and only if its summary T_B is.

³The size of the representation of a tournament is already quadratic in the number of its alternatives.

$D(T)$, let $B_i(D(T), B)$ denote the i th children of B in $D(T)$. Composition-consistency implies that

$$S(T|_B) = \bigcup_{i \in S(T_B)} S(T|_{B_i(D(T), B)}).$$

The choice set $S(T)$ can thus be computed by starting at the root of $D(T)$ and iteratively applying the equation above. If B is reducible, we immediately know that $S(T|_B) = S(T|_{B_i(D(T), B)})$, since 1 is the maximum of the transitive tournament T_B . A straightforward implementation of this approach is given in Algorithm 1.

Algorithm 1 Compute $S(T)$ via decomposition tree

```

1: Compute  $D(T)$ 
2:  $S \leftarrow \emptyset$ 
3:  $Q \leftarrow (A)$ 
4: while  $Q \neq ()$  do
5:    $B \leftarrow \text{Dequeue}(Q)$ 
6:   if  $|B| = 1$  then
7:      $S \leftarrow S \cup B$ 
8:   else
9:     if  $B$  is reducible then
10:       $\text{Enqueue}(Q, B_1(D(T), B))$ 
11:     else //  $B$  is irreducible
12:       for all  $i \in S(T_B)$  do
13:          $\text{Enqueue}(Q, B_i(D(T), B))$ 
14: return  $S$ 

```

Algorithm 1 visits each node of $D(T)$ at most once. The algorithm for computing S is only invoked for tournaments T_B for which B is irreducible and $|B| \geq 2$. The order of such a tournament T_B is equal to the number of children of node B in $D(T)$. The *decomposition degree* of T is defined as an upper bound of this number.

Definition 6 Let $\text{Irr}(D(T))$ be the set of irreducible internal nodes of $D(T)$. The decomposition degree $\delta(T)$ of a tournament T is given by

$$\delta(T) = \begin{cases} \max\{|T_B| : B \in \text{Irr}(D(T))\}, & \text{if } \text{Irr}(D(T)) \neq \emptyset \\ 0, & \text{otherwise.} \end{cases}$$

Proposition 3 implies that $\delta(T)$ can be computed efficiently. The decomposition degree of the example tournament in Figure 1 is 3.

Let $f(n)$ be an upper bound on the running time of an algorithm that computes $S(T)$ for tournaments of order $|T| \leq n$. Then, the running time of Algorithm 1 can be upper-bounded by $f(\delta(T))$ times the number of irreducible nodes of $D(T)$. We thus obtain the following theorem.

Theorem 1 Let S be a composition-consistent tournament solution and let $f(k)$ be an upper bound on the running time of an algorithm that computes S for tournaments of order at most k . Then, $S(T)$ can be computed in $O(n^2) + f(\delta) \cdot (n - 1)$ time, where δ is the decomposition degree of T and n is the order of T .

Proof: Let T be a tournament and $n = |T|$. We show that Algorithm 1 computes $S(T)$ in $O(n^2) + f(\delta(T)) \cdot |\text{Irr}(D(T))|$ time. Correctness follows from composition-consistency of S . The running time can be bounded as follows. Computing $D(T)$ requires time $O(n^2)$ (Proposition 3). During the execution of the while-loop, each node B of $D(T)$ is visited at most once. If B is reducible or a singleton, there is no further computation. If $B \in \text{Irr}(D(T))$, $S(T_B)$ is computed. As $|T_B|$ is upper-bounded by $\delta(T)$, this can be done in $f(\delta(T))$ time. Finally, as the number of internal nodes in a tree with n leafs is bounded by $n - 1$, we have that $|\text{Irr}(D(T))| \leq n - 1$. Summing up, this yields a running time of at most $O(n^2) + f(\delta(T)) \cdot (n - 1)$. \square

In particular, Theorem 1 shows that the computation of $S(T)$ is fixed-parameter tractable with respect to the parameter $\delta(T)$.

To get a better understanding of this theorem, consider a composition-consistent tournament solution S such that $f(n)$ is in $\text{E} = \text{DTIME}(2^{O(n)})$. This holds, for example, for the Banks set. For any given tournament T of order n , Theorem 1 then implies that $S(T)$ can be computed efficiently (*i.e.*, in time polynomial in n) whenever $\delta(T)$ is in $O(\log^k n)$. Theorem 1 is of course also applicable to tractable tournament solutions such as the minimal covering set and the bipartisan set. Although computing these solutions is known to be in P , existing algorithms rely on linear programming and may be too time-consuming for very large tournaments. For both concepts, a significant speed-up can be expected for distributions of tournaments that admit a small decomposition degree.

Generally, decomposing a tournament asymptotically never harms the running time, as the time required for computing the decomposition tree is only linear in the input size.⁴

5 Experimental Results

It has been shown in the previous section that computing composition-consistent tournament solutions is fixed-parameter tractable with respect to the decomposition degree of a tournament. While the clustering of alternatives within components has some natural appeal by itself, an important question concerns the value of the decomposition degree for reasonable and practically motivated distributions of tournaments. In this section, we will explore this question experimentally using two probabilistic models from social choice theory. Both models are based on a set of an odd number of voters who entertain preferences over candidates. Given a finite set of candidates A and an odd number of voters with linear preferences over A , the *majority tournament* is defined as the tournament $(A, >)$, where $a > b$ if and only if the number of voters preferring a to b is greater than the number of voters preferring b to a .

One of the most studied and mathematically simplest preference model in social choice theory is called *impartial culture* and assumes that all preference profiles are equally likely and hence uniformly distributed. However, impartial culture is generally considered as fairly unrealistic (see, *e.g.*, Tsetlin

⁴Checking whether there exists a maximum already requires $O(n^2)$ time.

et al., 2003) as it does not impose *any* structure on the preferences and it is precisely this structure that we seek to exploit. It can be shown that under impartial culture, the probability that just a single non-trivial component exists goes to zero. Hence, there would be little to no speedup from exploiting composition-consistency. We consider the following alternative models.

Noise model The first model we consider is a standard model in social choice theory where it is usually attributed to Condorcet (see, *e.g.*, Young, 1988). Condorcet assumed that there exists a “true” ranking of the candidates and that the voters possess noisy estimates of this ranking. In particular, he assumed that there is a probability $p > \frac{1}{2}$, such that for each pair a, b of candidates, each voter ranks a and b according to the true ranking with probability p and ranks them incorrectly with probability $1 - p$.

Spatial Model Spatial models of voting are well-studied objects in social choice theory (see, *e.g.*, Austen-Smith and Banks, 2000; Conitzer, 2006). For a fixed natural number d of issues, we assume that candidates as well as voters are located in the space $[0, 1]^d$. The position of candidates and voters can be thought of as their stance on the d issues. Voters’ preferences over candidates are given by the proximity to their own position according to the Euclidian distance. The one-dimensional case coincides with the well-studied model of single-peaked preferences. We generate tournaments by drawing the positions of candidates and voters uniformly at random from $[0, 1]^d$.




Figure 2: Decomposition degree (noise model, $p = 0.55$)

The results of our experiments with regard to the decomposition degree are presented in Figures 2 and 3. The x -axis shows the number of voters (starting at 5 with increments of 30). In order to facilitate the comparison of results for a varying number of candidates, the y -axis shows the *normalized* decomposition degree, *i.e.*, the decomposition degree divided by the number of candidates. Each graph shows the results for a fixed number of candidates, and each data point corresponds to the average value of 30 instances. Whenever the normalized decomposition degree is less than one,

composition-consistency can be exploited, even for tournament solutions that already admit optimal (*i.e.*, linear-time) algorithms. The slower the original algorithm, the more dramatic is the speedup obtained by capitalizing on the decomposition tree.

Figure 2 shows the results for the noise model with parameter $p = 0.55$. For any number of candidates, the decomposition degree goes to zero when the number of voters grows. This is not surprising because the probability that the tournament is transitive tends to 1 for any $p > \frac{1}{2}$ (and a transitive tournament T has $\delta(T) = 0$). Interestingly, the decomposition degree drops abruptly when a certain number of voters is reached.




Figure 3: Decomposition degree (spatial model, $d = 2$)

Figure 3 shows the results for the two-dimensional spatial model ($d = 2$). We also experimented with dimensions between 2 and 20. Surprisingly, the decomposition degree does not significantly increase when moving to a higher-dimensional space. Similar to the noise model discussed above, δ tends to 0 for growing n because a population of voters that is evenly distributed in $[0, 1]^d$ tends to produce transitive tournaments.

The results of these experiments show that, even for moderately-sized electorates, tournaments in both distributions are highly decomposable and therefore allow significantly faster algorithms for computing composition-consistent tournament solutions.

In order to examine the actual impact on the running time, we compared the naive implementation of the tournament equilibrium set (Brandt *et al.*, 2010) with Algorithm 1. The results for tournaments of size 30 generated with the noise model are shown in Figure 4 where each data point represents an average over 20 runs. For a small number of voters, when the tournaments are typically non-decomposable, the naive algorithm is already quite quick and using Algorithm 1 does not lead to a significant speedup. However, as the number of voters increases, tournaments become more structured and exploiting composition-consistency proves to be extremely beneficial. In case of the spatial model, tournaments tend to be decomposable even for a small number of voters (see Figure 3) and consequently, Algorithm 1 achieves a *minimal* average speedup of around 10^3 for 30 candidates,

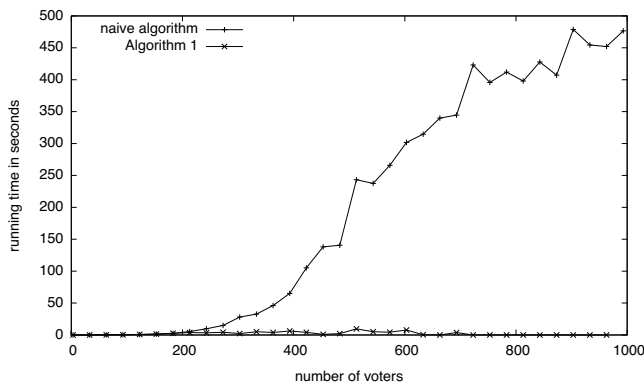


Figure 4: Comparison of running times for computing the tournament equilibrium set on a 2.66GHz Core i5 machine (noise model, $p = 0.55$, 30 candidates)

independently of the number of voters and dimensions.

6 Conclusion

In this paper, we studied the algorithmic benefits of composition-consistent tournament solutions. We defined the decomposition degree of a tournament as a parameter that reflects its decomposability. Intuitively, a low decomposition degree indicates that the tournament admits a particularly well-behaved decomposition. Our main result states that computing any composition-consistent tournament solution is fixed-parameter tractable with respect to the decomposition degree. This is of particular relevance for tournament solutions that are known to be computationally intractable such as the Banks set and the tournament equilibrium set. For example, one corollary of our main result is that the Banks set of a tournament can be computed efficiently whenever the decomposition degree is polylogarithmic in the number of alternatives. We experimentally determined the decomposition degree of two natural distributions of tournaments stemming from social choice theory and found that the decomposition degree in many realistic instances is surprisingly low. As a consequence, the speedup obtained by exploiting composition-consistency when computing tournament solutions for these instances is quite substantial as we showed experimentally for the tournament equilibrium set. Since computing a decomposition tree requires only linear time, decomposing a tournament never hurts, and often helps. Composition-consistency can be further exploited by parallelization and storing the solutions of small tournaments in a lookup table.

References

N. Alon. Ranking tournaments. *SIAM Journal on Discrete Mathematics*, 20(1):137–142, 2006.

K. J. Arrow and H. Raynaud. *Social Choice and Multicriterion Decision-Making*. MIT Press, 1986.

D. Austen-Smith and J. S. Banks. *Positive Political Theory I: Collective Preference*. University of Michigan Press, 2000.

N. Betzler, R. Bredereck, and R. Niedermeier. Partial kernelization for rank aggregation: Theory and experiments. In *Proc. of 5th International Symposium on Parameterized and Exact Computation (IPEC)*, volume 6478 of *Lecture Notes in Computer Science (LNCS)*, pages 26–37. Springer-Verlag, 2010.

F. Brandt and F. Fischer. Computing the minimal covering set. *Mathematical Social Sciences*, 56(2):254–268, 2008.

F. Brandt and P. Harrenstein. Characterization of dominance relations in finite coalitional games. *Theory and Decision*, 69(2):233–256, 2010.

F. Brandt, F. Fischer, P. Harrenstein, and M. Mair. A computational analysis of the tournament equilibrium set. *Social Choice and Welfare*, 34(4):597–609, 2010.

F. Brandt. Minimal stable sets in tournaments. *Journal of Economic Theory*, 2011. Forthcoming.

C. Capelle, M. Habib, and F. de Montgolfier. Graph decompositions and factorizing permutations. *Discrete Mathematics and Theoretical Computer Science*, 5:55–70, 2002.

V. Conitzer. Computing Slater rankings using similarities among candidates. In *Proc. of 21st AAAI Conference*, pages 613–619. AAAI Press, 2006.

P. E. Dunne. Computational properties of argumentation systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 171(10-15):701–729, 2007.

S. L. Feld and B. Grofman. Who’s afraid of the big bad cycle? Evidence from 36 elections. *Journal of Theoretical Politics*, 4:231–237, 1992.

D. C. Fisher and J. Ryan. Tournament games and positive tournaments. *Journal of Graph Theory*, 19(2):217–236, 1995.

G. Laffond, J. Lainé, and J.-F. Laslier. Composition-consistent tournament solutions and social choice functions. *Social Choice and Welfare*, 13:75–93, 1996.

J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer, 1997.

R. M. McConnell and F. de Montgolfier. Linear-time modular decomposition of directed graphs. *Discrete Applied Mathematics*, 145(2):198–209, 2005.

H. Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3:271–291, 1986.

R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.

T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.

I. Tsetlin, M. Regenwetter, and B. Grofman. The impartial culture maximizes the probability of majority cycles. *Social Choice and Welfare*, 21(3):387–398, 2003.

G. J. Woeginger. Banks winners in tournaments are difficult to recognize. *Social Choice and Welfare*, 20:523–528, 2003.

H. P. Young. Condorcet’s theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.