

Tiered Coalition Formation Games

Cory Siler

University of Kentucky
Lexington, Kentucky, USA

Abstract

In competitive videogame communities, a tier list is a hierarchical ranking of playable characters that, despite its simplicity, tries to capture an often nuanced metagame where matchups between characters do not follow a transitive ordering. We model the creation of tier lists as a coalition formation game, based on hedonic games, where the set of agents is partitioned into a hierarchy and an agent has preferences over the set of agents at and below its level of the hierarchy. We prove the computational complexity of determining whether there exists a stable partition under two stability notions borrowed from hedonic games.

1 Introduction

In many scenarios involving pairwise comparisons between objects, *intransitivity* prevents a precise best-to-worst ranking of those objects. A classic example is the game of Rock-Paper-Scissors: Paper beats rock, and rock beats scissors, but this does not entail that paper beats scissors. Similar relationships occur in voting (consider the Condorcet paradox), preferences (Tversky (1969) observes that human choices are often inconsistent with transitive preference orderings), and sports. Recently, there have been efforts to model these relationships quantitatively, e.g., Chen and Joachims (2016) and Saarinen, Tovey, and Goldsmith (2014).

Of particular interest is character selection in videogames, a major strategic decision beyond the gameplay itself. Previous game-theoretic perspectives on character selection include the work of Jaffe (2013), who analyzes it as a zero-sum game based on win probabilities in the ensuing matches, and Spradling et al. (2013; 2015), who consider matchmaking in team-based videogames incorporating gamers’ preferences over character “roles” and group compositions. Our model instead focuses on the notion of a *tier list*. A tier list, in the casual sense, is a ranking of playable characters into levels representing their competitive viability, with higher-tier characters dominating the competitive scene. The creation of a tier list (by the collective decision of a videogame community or by an authority within it) is subjective and often controversial, owing in part to intransitivity; a character that is perceived as weak may perform exceedingly well against certain characters generally considered “better”.

Some competitive videogame communities use tier lists as a game balancing tool: The character selection in a tournament may be limited only to characters at and below a certain position on the community’s tier list. For instance, the *Pokémon* community Smogon¹ uses this model; matches take place on a predefined tier level, with “Ubers” matches allowing all characters, “OverUsed” matches allowing all but a select few dominant characters, “UnderUsed” matches allowing only a subset of that, and so on. (A character may be played in a match on a level higher than its own tier, and may even be an effective counter to a character or strategy popular on that level, but the character is likely to underperform on that level in general.) The tier system allows players who enjoy a particular character to have an environment where they can use that character competitively even if the character is not considered a viable choice with respect to the entire pool of characters.

Tier list formation is an iterative process; Smogon’s tiers are periodically updated according to a formula based on frequency of usage, and other communities adjust based on human deliberation. In our case, we wish to model tier list formation as a computational problem to study it from a complexity-theoretic point of view.

We base our model on hedonic games (Bogomolnaia and Jackson 2002), in which a set of agents is partitioned into a set of coalitions, with agents having preferences over the coalition they join; relevant problems include finding *stable* partitions, where no agent (or colluding group of agents) would prefer to move to a different coalition with respect to some set of allowable moves. In the model that we introduce, *tiered coalition formation games*, agents represent characters being placed into tiers; these tiers are analogous to the coalitions in hedonic games, but rather than being independent of each other, they are ordered from high-tier to low-tier. Since, in a Smogon-like competitive environment, a character played on its own tier level may compete against opponents at *and* below its own tier, an agent has preferences over the combined membership of the tier list from its own tier downward, with more-preferred sets of agents representing more-favorable environments.

In Section 2 we formalize the tiered coalition formation game model. In Section 3 we explore two notions of stability

¹<http://www.smogon.com/>

and the computational complexity of deciding the existence of stable partitions. In Section 4 we introduce a preference representation that trades expressivity for certain guarantees about stability. In Section 5 we summarize and discuss directions for future work.

2 Definitions

Definition 1. A tiered coalition formation game (TCFG) instance is a pair $(A, \succeq)$ consisting of a finite set of agents $A = \{a_1, a_2, \dots, a_n\}$ and a preference profile $\succeq = \{\succeq_1, \succeq_2, \dots, \succeq_n\}$, where $\succeq_i$ is a weak total order over subsets of A that contain a_i .

An outcome (tier list) consists of a partition of A into a totally ordered set of k disjoint coalitions (tiers) $T = \{T_1, T_2, \dots, T_k\}$.

In a tier list T , we say that a_i sees a_j if a_i is in same tier as a_j or a higher tier: $\text{Seen}(a_i, T) = \bigcup_{l=1}^m T_l, a_i \in T_m$. For tier lists T and T' , $\succeq_i$ specifies a_i 's preference between $\text{Seen}(a_i, T)$ and $\text{Seen}(a_i, T')$.

Example 1. Consider a TCFG with two agents, a_1 and a_2 . The possible tier lists are as follows: $T = \{\{a_2\}, \{a_1\}\}$; $T' = \{\{a_1, a_2\}\}$; and $T'' = \{\{a_1\}, \{a_2\}\}$. Note that with respect to a_1 's preferences $\succeq_1$, T (where a_1 is in a higher tier than a_2) and T' (where both are in the same tier) are indistinguishable: $\text{Seen}(a_1, T) = \text{Seen}(a_1, T') = \{a_1, a_2\}$. On the other hand, a_1 does not see a_2 from its lower tier position in T'' : $\text{Seen}(a_1, T'') = \{a_1\}$.

The initial specification for a TCFG consists of a set of agents and a preference specification, just like a hedonic game. However, unlike the coalitions resulting from a hedonic game, the tiers resulting from a TCFG are arranged in a hierarchy; we will analyze an agent's strategic behavior under the assumption that its preferences are based not only on the membership of its own tier but on the combined membership of all tiers up to its own on the hierarchy.

3 Stability

For the theorems and proofs in this section, we assume TCFG preferences are expressed in terms of a hedonic coalition net (Elkind and Wooldridge 2009), a fully-expressive representation consisting of rules that award agents utility for being in a coalition that satisfies a given propositional formula. This allows the effects of single-agent deviations to be checked in polynomial time; furthermore, the agent "priorities" described in this section's reductions can be reformulated in terms of a polynomially-sized set of hedonic coalition net rules, assigning higher utilities to higher-priority conditions.

In a hedonic game, a partition is Nash stable if there exists no move by any individual agent (joining another existing coalition or forming a new coalition by itself) that would result in a higher-utility coalition for that agent. We introduce a similar notion of Nash stability for TCFGs:

Definition 2. A tier list T is Nash stable if there exists no move (to an existing tier, or by forming a new tier in the hierarchy) by any individual agent a_i that would yield a new tier list T' such that $\text{Seen}(a_i, T') \succ_i \text{Seen}(a_i, T)$.

Theorem 1. The problem of deciding whether there exists a Nash stable tier list for a given TCFG is NP-complete.

Proof. For membership in NP, we observe that given a Nash stable tier list, we can verify its Nash stability in polynomial time by checking each of the no more than $n + 1$ possible movements for each of the n agents.

We will show NP-hardness with a reduction from the problem of deciding the existence of a Nash stable partition in additively separable hedonic games, for which Olsen (2009) proves NP-completeness. An additively separable hedonic game instance is a pair (A, U) consisting of a set of agents $A = \{a_1, a_2, \dots, a_n\}$ and a utility function $U : A \times A \rightarrow \mathbb{R}$. An outcome consists of a partition of A into an unordered set of disjoint coalitions. An agent a_i 's utility for being in coalition C is the sum of its utilities for the individual members, $\sum_{a_j \in C} U(a_i, a_j)$.

Given an additively separable hedonic game instance (A, U) , we can derive a TCFG instance $(A', \succeq)$ in polynomial time such that (A, U) has a Nash stable partition iff $(A', \succeq)$ has a Nash stable tier list. Rather than explicitly generate the entire preference profile $\succeq$, which is exponential in the number of agents, we will represent the preference orderings implicitly in the form of rules described below.

We construct the TCFG's set of agents as $A' = A \cup L \cup X \cup Y$, containing the original set of agents A plus new agents $L = \{l_1, l_2, \dots, l_n\}$, $X = \{x_1, x_2, \dots, x_n\}$, and $Y = \{y_1, y_2, \dots, y_n\}$.

We define each x_i 's preferences such that x_i adheres to the following rules, in descending order of priority (i.e., earlier rules take precedence, with later rules breaking ties):

1. Avoid seeing y_i .
2. See y_{i-1} (for $i > 1$).
3. Avoid seeing more than i A -agents.
4. Maximize the number of A -agents seen (up to i).

Similarly, y_i 's priorities:

1. Avoid seeing x_{i+1} (for $i < n$).
2. See x_i .
3. Avoid seeing more than i A -agents.
4. Maximize the number of A -agents seen (up to i).

Together, the rules 1 and 2 for X and Y create an alternating sequence of X -agents and Y -agents in any Nash stable tier list T , with subscripts in ascending tier order.

Priorities for a_i :

1. See more X -agents than Y -agents.
2. Evaluate the set of seen A -agents with respect to the original utility function U , except that if l_j is seen, do not receive the utility for a_j .

In a Nash stable tier list, rule 1 for A -agents, combined with rules 3 and 4 for X - and Y -agents, ensures that each A -agent shares a tier with an X -agent.

The L -agents act as "utility locks" for their corresponding A agents. Priorities for l_i :

1. See the same number of X -agents and Y -agents.

2. See a_i .
3. Minimize the number of A -agents seen.

In a Nash stable tier list, these rules constrain l_i to be on a tier above a_i , but below the next tier containing A -agents (since sharing a tier with A -agents would result in k_i seeing more X -agents than Y -agents, contrary to its rule 1). Thus, only A -agents that share a tier with a_i see a_i but not k_i and thus receive utility for a_i .

Any Nash stable tier list for the TCFG maps to a Nash stable partition for the original instance — namely, one where the A -agents that shared a tier now share a coalition. Likewise, a Nash stable partition has corresponding Nash stable tier lists in the reduction, where the A -agents that shared a coalition share a tier. $\square$

We likewise define core stability for TCFGs based on the hedonic game notion, wherein a coalition structure is core stable if there is no set of agents that would prefer to form a new coalition together:

Definition 3. A tier list T is core stable if there exists no nonempty set of agents $B \subset A$ that could form a new tier together in the hierarchy such that for the resulting tier list T' , $\text{Seen}(a_i, T') \succ_i \text{Seen}(a_i, T)$ for all $a_i \in B$.

Unlike in hedonic coalition structures (Bogomolnaia and Jackson 2002), there is an entailment relationship between core and Nash stability in tier lists:

Theorem 2. If a tier list is core stable, it is also Nash stable.

Proof. Suppose a tier list T is core stable but not Nash stable; there exists an agent a_i that can improve its utility by moving to another tier. If the tier has other agents, a_i could obtain the same set of seen agents by forming a new tier directly above that tier instead. Letting $B = \{a_i\}$ in Definition 3, this contradicts the assumption that T is core stable. $\square$

Like Nash stable tier lists, and like core stable partitions in hedonic games (Ballester 2004) determining whether there is a core stable tier list is a hard problem:

Theorem 3. The problem of deciding whether there exists a core stable tier list for a given TCFG is NP-hard.

Proof. We will show NP-hardness with a reduction from exact cover by 3-sets, for which Garey and Johnson (1979) prove NP-completeness. An instance is a pair (A, S) , where A and S are sets, $|A|$ is a multiple of 3, and the elements of S are size-3 subsets of A . The instance is satisfiable if some subset of S constitutes a partition of A .

Given (A, S) , we can derive a TCFG instance $(A', \succeq)$ in polynomial time such that A can be partitioned into a subset of S iff $(A', \succeq)$ has a core stable tier list.

Let $|A| = n$. We construct the TCFG's set of agents as $A' = A \cup L \cup X \cup Y$, where $L = \{l_1, l_2, \dots, l_n\}$, $X = \{x_1, x_2, \dots, x_n\}$, and $Y = \{y_1, y_2, \dots, y_n\}$.

We define the preferences of the L , X , and Y agents to be the same as in the proof of Theorem 1. We define the preferences of agents $a_i \in A$ such that a_i follows these priorities:

1. See the same number of X -agents and Y -agents.

2. Among A -agents for which the corresponding L -agent is not seen, see exactly two other A -agents a_j and a_k such that $\{a_i, a_j, a_k\} \in S$.
3. If #2 cannot be satisfied, see no other agents.

If the exact-cover-by-3-sets instance has a satisfying partition, the A -agents will be grouped into its constituent subsets in the resulting TCFG's core stable tier lists. Otherwise, any tier list will have A -agents whose priority #2 is not satisfied; for tier lists where no more A -agents could group together to satisfy #2, there will be multiple (specifically, a multiple of 3) remaining A -agents with incompatible preferences to satisfy priority #3, preventing Nash stability and hence core stability. $\square$

4 Simple Preferences

It is common in the hedonic-games community to study preference criteria that are not fully expressive (i.e., cannot encode all possible preference orderings for an agent) but that give guarantees about the existence or computability of stable partitions. (For a recent survey, see Aziz and Savani (2016).) We now initiate a similar course of study for TCFGs by proposing a simple preference representation, loosely inspired by the friend- and enemy-oriented hedonic preferences of Dimitrov et al. (2006) and based explicitly on the idea of pairwise matchups between competitors and the preference for a competitor to avoid unfavorable matchups.

Definition 4. A matchup-oriented preference representation consists of an antisymmetric, antireflexive relation between pairs of agents. For a pair (a_i, a_j) in this relation, we say that a_j is a “good matchup” for a_i and that a_i is a “bad matchup” for a_j . For an agent a_i $\text{Seen}(a_i, T) \succ_i \text{Seen}(a_i, T')$ if $\text{Seen}(a_i, T)$ has a higher difference of good matchups — bad matchups.

Whereas core stability implies Nash stability for tier lists in general, with this representation the two notions are equivalent:

Theorem 4. If a tier list is Nash stable under matchup-oriented preferences, it is also core stable.

Proof. Suppose a tier list T is Nash stable but not core stable; there exists a set of agents B of size 2 or greater (as size 1 would violate Nash stability) that would prefer to form a new tier together.

Suppose the new tier were above the current tiers of every agent in B ; then, the set of agents seen from that tier would be the same as the set of agents seen by any individual agent in B if that agent were to move to the new tier position on its own. This would contradict the fact that T is Nash stable, so the new tier position must be below the current tier of at least one agent in B .

Consider the set $B' \subseteq B$ of agents that would move downward to form the new tier. Since T is Nash stable, it must be the case that each $a_i \in B'$ sees at least as many good matchups as bad matchups in T from its current tier down to the new tier position (or else a_i would be incentivized to move down to the new tier position on its own).

Meanwhile, in order to be incentivized to form the new tier with the other agents in B , each $a_i \in B'$ must lose more bad matchups than good matchups in the group movement. Since a_i sees at least as many good matchups as bad matchups in its current tier, and loses more bad matchups than good matchups when forming the new tier, it must have more good matchups than bad matchups with other agents in B' . But this cannot be the case for every $a_i \in B'$ (as every good matchup within the group for one agent entails a bad matchup for another). Thus, the new tier position cannot be below any agent's current position as required.

So assuming a Nash stable but not core stable tier list leads to a contradiction. $\square$

While a general TCFG may have no Nash or core stable tier list, one is guaranteed to exist under this representation:

Theorem 5. *Under matchup-oriented preferences, there always exists a Nash (and, hence, core) stable tier list.*

Proof. Given a non-Nash-stable tier list T where $|T_k| = 1$ for all $T_k \in T$, we can move some agent a_i to a new tier in the hierarchy resulting in a tier list T' such that $\text{Seen}(a_i, T') \succ_i \text{Seen}(a_i, T)$ and $|T'_k| = 1$ for all $T'_k \in T'$.

For each agent a_j that is seen by a_i in T but not in T' , a_i is seen by a_j in T' but not in T ; likewise, for a_j that is seen by a_i in T' but not in T , a_i is seen by a_j in T but not in T' . As such, each gain in a good/bad matchup by a_i results in the loss of a bad/good matchup (respectively) by some other agent, and vice versa; so since the movement is a net improvement for a_i , it is also causes an average net improvement for the other agents.

If we continue making such changes as long as some a_i is incentivized to move, either we will eventually arrive at a Nash stable tier list, or the sequence of changes will contain a cycle of tier lists (as there are a finite number of possible tier lists); but since each change yields an average net improvement for the set of agents, a cycle is not possible and a Nash stable tier list will result. $\square$

Corollary 1. *Under matchup-oriented preferences, a Nash (and, hence, core) stable tier list can be found in polynomial time.*

Proof. The difference of the highest and lowest possible total numbers of *good matchups* – *bad matchups* for all n agents in a tier list is bounded by $O(n^2)$. The current total increases in each iteration of the process described in the proof of Theorem 5, and each iteration takes $O(n^2)$ time (checking n agents against $O(n)$ possible tier positions) to look for an agent with an improving movement, so the process gives an $O(n^4)$ algorithm for finding a stable tier list. $\square$

5 Conclusions

Our hardness results reflect the observation, well-known among videogamers, that capturing the nuances of a competitive metagame in tier list form is no simple affair, and may not always have a satisfactory solution. However, simplified TCFG representations that allow efficient algorithms

are worth further exploration. The matchup-oriented preferences presented here have room for improvement; they allow for counterintuitive stable tier lists, like Rock-Paper-Scissors having three separate tiers (whereas it would be sensible for only the tier list where all three options share a tier to be stable). This can be counteracted by adding a second preference criterion that, all else being the same, an agent prefers seeing more agents to fewer agents, but this also destroys the guarantee that stable partitions exist.

A thorough assessment of TCFG representations will also require an empirical component — i.e., investigating the relationship between computed stable tier lists and human-made tier lists for real competitive environments.

References

- Aziz, H., and Savani, R. 2016. Hedonic games. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 15, 356–376.
- Ballester, C. 2004. NP-completeness in hedonic games. *Games and Economic Behavior* 49(1):1–30.
- Bogomolnaia, A., and Jackson, M. O. 2002. The stability of hedonic coalition structures. *Games and Economic Behavior* 38(2):201–230.
- Chen, S., and Joachims, T. 2016. Modeling intransitivity in matchup and comparison data. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, 227–236. ACM.
- Dimitrov, D.; Borm, P.; Hendrickx, R.; and Sung, S. C. 2006. Simple priorities and core stability in hedonic games. *Social Choice and Welfare* 26(2):421–433.
- Elkind, E., and Wooldridge, M. 2009. Hedonic coalition nets. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 417–424. International Foundation for Autonomous Agents and Multiagent Systems.
- Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, New York.
- Jaffe, A. B. 2013. Understanding game balance with quantitative methods. Ph.D. thesis, University of Washington.
- Olsen, M. 2009. Nash stability in additively separable hedonic games and community structures. *Theory of Computing Systems* 45(4):917–925.
- Saarinen, S.; Tovey, C.; and Goldsmith, J. 2014. A model for intransitive preferences. In *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- Spradling, M., and Goldsmith, J. 2015. Stability in role based hedonic games. In *The Twenty-Eighth International FLAIRS Conference*, 85–90.
- Spradling, M.; Goldsmith, J.; Liu, X.; Dadi, C.; and Li, Z. 2013. Roles and teams hedonic game. In *International Conference on Algorithmic Decision Theory*, 351–362. Springer.
- Tversky, A. 1969. Intransitivity of preferences. *Psychological Review* 76(1):31.