Hypercubewise Preference Aggregation in Multi-issue Domains

Vincent Conitzer
Department of Computer Science
Duke University
Durham, NC 27708, USA
conitzer@cs.duke.edu

Jérôme Lang
LAMSADE
Université Paris-Dauphine
75775 Paris Cedex, France
lang@lamsade.dauphine.fr

Lirong Xia
Department of Computer Science
Duke University
Durham, NC 27708, USA
lxia@cs.duke.edu

Abstract

We consider a framework for preference aggregation on multiple binary issues, where agents' preferences are represented by (possibly cyclic) CP-nets. We focus on the majority aggregation of the individual CP-nets, which is the CP-net where the direction of each edge of the hypercube is decided according to the majority rule. First we focus on hypercube Condorcet winners (HCWs); in particular, we show that, assuming a uniform distribution for the CP-nets, the probability that there exists at least one HCW is at least $1 - 1/e$, and the expected number of HCWs is 1. Our experimental results confirm these results. We also show experimental results under the *Impartial Culture* assumption. We then generalize a few tournament solutions to select winners from (weighted) majoritarian CP-nets, namely Copeland, maximin, and Kemeny. For each of these, we address some social choice theoretic and computational issues.

1 Introduction

In many multi-agent scenarios, the space of alternatives has a combinatorial structure: there are p issues to decide on, each issue i takes a value from a set D_i, and n agents (voters) generally have preferential dependencies among these issues. In classical voting theory, voters submit their preferences as *linear orders* over the set of alternatives, and then a *voting rule* is applied to select a set of winning alternatives. However, when the set of alternatives has a multi-issue structure, the number of alternatives is exponential in the number of issues, and therefore, as soon as the number of issues is not very small, it is not realistic to ask voters to specify their preferences as explicit linear orders.

Several other ways of proceeding have been considered: (1) *voting separately on each issue simultaneously*, which is known to lead to severe "multiple election paradoxes" [2]; (2) *limiting the set of alternatives that voters may vote for*, which is quite arbitrary (who decides which alternatives are allowed?), and leads the voters to express their preferences on only a tiny fraction of the alternatives; (3) *asking voters to report only a (small) part of their preference relation and applying a voting rule that needs this information only, such as plurality*, which is not always a bad idea, but the smaller the part of their preference relation voters express, the more likely it is that results will not be significant as soon as the number of issues is large ($2^p \gg n$); (4) *imposing a domain restriction* such as separability (which allows for using method (1) above), or a weaker restriction such as $\mathcal{O}$-legality, which allows for deciding on the issues one after the other [7]; (5) *using a compact preference representation language* in which the voters' preferences are represented in a concise way; this method does not require any domain restriction, but leads to nontrivial computational issues.

In this paper, we take an approach that is intermediate between (3) and (5): we elicit only a part of the voters' preferences — however, not a *small* part of it, but still, a part significantly smaller than the explicit specification of the full preferences; and we use it to draw a (partial) preference relation, using the semantics of a preference representation language, namely CP-nets [1].

Group decision making in multi-issue domains via CP-net aggregation has been considered in a number of papers, which we briefly review in a structured (nonchronological) order. Rossi *et al.* [11] were the first to address the aggregation of CP-nets; given a collection of acyclic CP-nets, they define several aggregation functions mapping the preference relations induced by the individual CP-nets to a collective preference relation. This approach was pushed further by Li *et al.* [8], who give algorithms for computing Pareto-optimal alternatives with respect to the preference relations induced by the CP-nets, and fair alternatives with respect to a cardinalization of these preference relations. Another path is followed by Lang and Xia [7]; they also assume that the individual CP-nets are acyclic, and furthermore share the same acyclic dependence graph, and study a family of sequential voting rules that consider the issues one after another, following the dependence graph. Xia *et al.* [12] take still another direction: they do not make any domain restriction on the individual CP-nets, rather they consider separately every set of "neighboring" alternatives differing only in the value of one issue, use a local voting rule for deciding the common preferences over this set, and finally, optimal outcomes are defined based on the aggregated CP-net. They also introduce the notion of *local Condorcet winners*, which are alternatives that beat each of their neighbors in a pairwise majority duel. The latter notion was studied further by Li *et al.* [9], who study some of its properties and propose (and implement) a SAT-based algorithm for computing them.

In this paper, we follow the direction of [12; 9]. The input consists of arbitrary consistent CP-nets on a multi-issue domain whose issues are all binary. Because every preference relation on a multi-issue domain extends the preference relation induced by some consistent CP-net, this method does not require any domain restriction. Then we go further in studying the properties of *local Condorcet winners*, which we rename *hypercube Condorcet winners* (HCWs) for reasons that will be made clear. After giving some background on CP-nets in Section 2 and introducing the majoritarian aggregation of CP-nets in Section 3, we recall the notion of HCWs in Section 4 and give a simple complexity result about the existence of HCWs. Then we focus on an important problem, namely the worst-case and expected numbers of HCWs under various assumptions. We give a theoretical analysis in Section 5 and an experimental analysis in Section 6. Lastly, we show how a few standard tournament solutions can be generalized in a natural way to inputs consisting of CP-nets, and focus on three of them (Copeland, maximin and Kemeny).

2 Multi-issue Domains and CP-nets

Let $\mathcal{X}$ be a finite set of *alternatives* (or *candidates*). A *vote* V is a linear order on $\mathcal{X}$, *i.e.*, a transitive, antisymmetric, and total relation on $\mathcal{X}$. The set of all linear orders on $\mathcal{X}$ is denoted by $L(\mathcal{X})$. An n-voter profile P is a collection of n votes, that is, $P = (V_1, \ldots, V_n)$, where $V_j \in L(\mathcal{X})$ for every $j \leq n$. The set of all profiles on $\mathcal{X}$ is denoted by $P(\mathcal{X})$. A *(voting) rule* $r : P(\mathcal{X}) \to 2^{\mathcal{X}}$ maps any profile to a subset of alternatives.

In this paper, the set of all alternatives $\mathcal{X}$ is a *multi-issue domain*. We have a set of *issues* $\mathcal{I} = \{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$ ($p \geq 2$), where each issue $\mathbf{x}_i$ takes values in a finite domain D_i. *In this paper, we assume that all issues are binary*: for every i we have $D_i = \{0_i, 1_i\}$. The set of alternatives is $\mathcal{X} = D_1 \times \cdots \times D_p$: an alternative is uniquely identified by its values on all issues. For any alternative $\vec{d} = (d_1, \ldots, d_p)$ and any issue $\mathbf{x}_i$, we let $\vec{d}|_{\mathbf{x}_i} = d_i$ and $\vec{d}_{-i} = (d_1, \ldots, d_{i-1}, d_{i+1}, \ldots, d_p)$. For any $I \subseteq \mathcal{I}$, we let $D_I = \prod_{\mathbf{x}_i \in I} D_i$, and $D_{-i} = D_{\mathcal{I} \setminus \{\mathbf{x}_i\}}$.

A CP-net $\mathcal{N}$ over $\mathcal{X}$ consists of two components: (a) a directed graph $G = (\mathcal{I}, E)$ and (b) a set of conditional linear preferences $\succ^i_{\vec{u}}$ over D_i, for any $i \leq p$ and any setting $\vec{u}$ of the parents of $\mathbf{x}_i$ in G (denoted by $Par_G(\mathbf{x}_i)$). These conditional linear preferences $\succ^i_{\vec{u}}$ over D_i form the *conditional preference table* for issue $\mathbf{x}_i$, denoted by $CPT(\mathbf{x}_i)$. When G is acyclic, $\mathcal{N}$ is said to be an *acyclic CP-net*. The size of a CP-net is the cumulative size of all its conditional preference tables.

The preference relation $\succ_{\mathcal{N}}$ induced by $\mathcal{N}$ is the transitive closure of $\{(a_i, \vec{u}, \vec{z}) \succ (b_i, \vec{u}, \vec{z}) \mid i \leq p; \vec{u} \in D_{Par_G(\mathbf{x}_i)}; a_i, b_i \in D_i, a_i \succ^i_{\vec{u}} b_i; \vec{z} \in D_{-(Par_G(\mathbf{x}_i) \cup \{\mathbf{x}_i\})}\}$. If $\succ_{\mathcal{N}}$ is asymmetric then $\mathcal{N}$ is *consistent*. If G is acyclic, then we know that $\mathcal{N}$ is consistent [1].

Because all issues are binary, a CP-net $\mathcal{N}$ can be visualized as a hypercube with directed edges in a p-dimensional space, where each vertex is an alternative, and any two neighboring vertices differ in only one component (issue): for any $i \leq p$ and $\vec{d}_{-i} \in D_{-i}$, there is a directed edge connecting $(0_i, \vec{d}_{-i})$ and $(1_i, \vec{d}_{-i})$, and the direction of the edge is from $(0_i, \vec{d}_{-i})$ to $(1_i, \vec{d}_{-i})$ if and only if $(0_i, \vec{d}_{-i}) \succ_{\mathcal{N}} (1_i, \vec{d}_{-i})$. For any alternative $\vec{d} \in \mathcal{X}$ and any $i \leq p$, we let $\vec{d}[\leftrightarrow i]$ denote the neighbor of $\vec{d}$ that only differs from $\vec{d}$ in the ith issue.

Example 1 *Let $p = 3$ and let $\mathcal{N}$ be a CP-net defined as follows: the directed graph has an edge from $\mathbf{x}_1$ to $\mathbf{x}_2$ and an edge from $\mathbf{x}_2$ to $\mathbf{x}_3$; the CPTs are $CPT(\mathbf{x}_1) = \{0_1 \succ 1_1\}$, $CPT(\mathbf{x}_2) = \{0_1 : 0_2 \succ 1_2, 1_1 : 1_2 \succ 0_2\}$, $CPT(\mathbf{x}_3) = \{0_2 : 0_3 \succ 1_3, 1_2 : 1_3 \succ 0_3\}$. $\mathcal{N}$ is illustrated in Figure 1.*

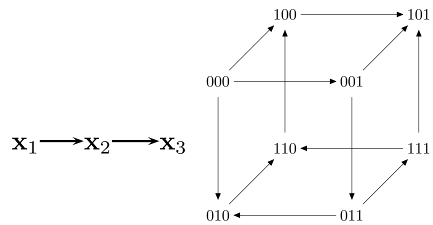

Figure 1: The hypercube representation of a CP-net. (For simplicity, 000 represents the alternative $0_1 0_2 0_3$, etc.)

A linear order V *extends* a CP-net $\mathcal{N}$, denoted by $V \sim \mathcal{N}$, if it extends $\succ_{\mathcal{N}}$.

3 Majoritarian Hypercube Aggregation

We assume that what is known about the voters' preferences is their underlying (consistent) CP-nets, *i.e.*, for every voter, we know the direction of every edge in the hypercube. Such a collection of consistent CP-nets will be called a *hypercube profile*, or for short, an *H-profile*. For the sake of simplicity, we also assume that there is an odd number of agents.

Definition 1 *Let $\mathcal{I} = \{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$ be a set of binary issues. An* H-profile *over $\mathcal{I}$ is a collection $P = (\mathcal{N}_1, \ldots, \mathcal{N}_n)$ of consistent CP-nets over $\mathcal{I}$. A profile $P_L = (V_1, \ldots, V_n)$ of linear orders over $\mathcal{X}$ extends an H-profile $P = (\mathcal{N}_1, \ldots, \mathcal{N}_n)$ (denoted by $P_L \sim P$) if for every $j \leq n$, V_j extends $\mathcal{N}_j$.*

Definition 2 *Given an H-profile $P = (\mathcal{N}_1, \ldots, \mathcal{N}_n)$, the majoritarian aggregation $M(P)$ of P is the CP-net $\mathcal{N}^*$ where for any pair of neighboring alternatives $\vec{x} = (\vec{x}_{-i}, x_i)$ and $\vec{y} = (\vec{x}_{-i}, x'_i)$, we have $\vec{x} \succ_{\mathcal{N}^*} \vec{y}$ if and only there is a majority of agents j such that $\vec{x} \succ_{\mathcal{N}_j} \vec{y}$.*

Similarly, we can define the majoritarian aggregation for a profile composed of (possibly inconsistent) CP-nets. The majoritarian aggregation has been defined and studied previously under different names [12; 8]. In the *weighted majoritarian aggregation* $W(P)$, each edge in the hypercube is associated with a weight, defined as follows. Suppose $\vec{x} \succ_{\mathcal{N}^*} \vec{y}$, then, the weight on the edge $\vec{x} \to \vec{y}$ is $w(\vec{x} \to \vec{y}) = |\{j \leq n : \vec{x} \succ_{\mathcal{N}_j} \vec{y}\}| - |\{j \leq n : \vec{y} \succ_{\mathcal{N}_j} \vec{x}\}|$.

Example 2 *We have two issues,* **S** *(build a swimming pool) and* **T** *(build a tennis court), and the following H-profile.*

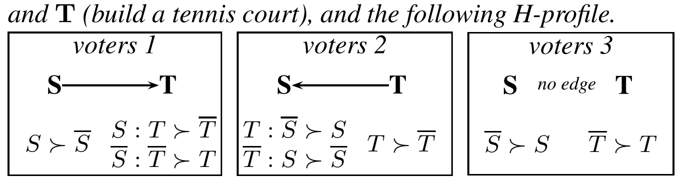

Note that the profile P_L consisting of the three linear preference relations ($ST \succ S\overline{T} \succ \overline{ST} \succ \overline{S}T$, $\overline{S}T \succ ST \succ S\overline{T} \succ \overline{ST}$, $\overline{ST} \succ S\overline{T} \succ \overline{S}T \succ ST$) extends P. The (weighted) majority aggregation of $\mathcal{N}_1$, $\mathcal{N}_2$ and $\mathcal{N}_3$ is the following CP-net, depicted with its induced preference relation.

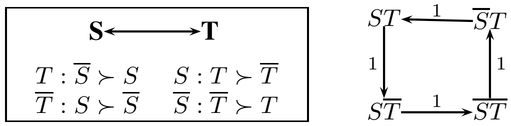

Note that although $\mathcal{N}_1$, $\mathcal{N}_2$, and $\mathcal{N}_3$ are consistent, their majoritarian aggregation is not, as it contains a cycle.

We note that usually the majoritarian aggregation is represented compactly as a CP-net (called the *majoritarian CP-net*), rather than directly as a hypercube. CP-nets are a good representation for majoritarian aggregations because the majoritarian aggregation preserves preferential independencies of the individual CP-nets. Therefore, the more structure the individual CP-nets share, the more compact the majoritarian CP-net is. Given a preference relation $\succ$, $\mathbf{x} \in \mathcal{I}$, $Y \subseteq \mathcal{I} \setminus \{\mathbf{x}\}$ and $Z = \mathcal{I} \setminus (\{\mathbf{x}\} \cup Y)$, we say that $\mathbf{x}$ is preferentially independent of Y given Z with respect to $\succ$, which we denote by $Ind(\mathbf{x}, Y, Z, \succ)$, if for any $x_i, x_i' \in D_{\mathbf{x}}$, $\vec{y}, \vec{y}' \in D_Y$, and $\vec{z} \in D_Z$, we have $(x_i, \vec{y}, \vec{z}) \succ (x_i', \vec{y}, \vec{z})$ if and only if $(x_i, \vec{y}', \vec{z}) \succ (x_i', \vec{y}', \vec{z})$. Similarly, $\mathbf{x}$ is preferentially independent of Y given Z with respect to a CP-net $\mathcal{N}$, which we denote by $Ind(\mathbf{x}, Y, Z, \mathcal{N})$, if $Ind(\mathbf{x}, Y, Z, \succ)$ holds for any $\succ$ extending $\succ_{\mathcal{N}}$. We immediate obtain the following proposition. Most proofs are omitted due to the space constraint.

Proposition 1 *Let $\mathbf{x} \in \mathcal{I}$, $Y \subseteq \mathcal{I} \setminus \{\mathbf{x}\}$ and $Z = \mathcal{I} \setminus (\{\mathbf{x}\} \cup Y)$. If $Ind(\mathbf{x}, Y, Z, \mathcal{N}_j)$ holds for every $j \leq n$, then we have $Ind(\mathbf{x}, Y, Z, M(\mathcal{N}_1, \ldots, \mathcal{N}_n))$.*

On the other hand, $Ind(\mathbf{x}, Y, Z, M(\mathcal{N}_1, \ldots, \mathcal{N}_n))$ may hold even if $Ind(\mathbf{x}, Y, Z, \mathcal{N}_j)$ fails to hold for some j (and even for every j). For instance, take $\mathcal{N}_1 = \{a \succ \bar{a}, b \succ \bar{b}, ab : \bar{c} \succ c, a\bar{b} : c \succ \bar{c}, \bar{a}b : c \succ \bar{c}, \bar{a}\bar{b} : c \succ \bar{c}\}$; $\mathcal{N}_2 = \{a \succ \bar{a}, b \succ \bar{b}, ab : c \succ \bar{c}, a\bar{b} : \bar{c} \succ c, \bar{a}b : c \succ \bar{c}, \bar{a}\bar{b} : c \succ \bar{c}\}$; $\mathcal{N}_3 = \{a \succ \bar{a}, b \succ \bar{b}, ab : c \succ \bar{c}, a\bar{b} : c \succ \bar{c}, \bar{a}b : c \succ \bar{c}, \bar{a}\bar{b} : \bar{c} \succ c\}$. Then $M(\mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3)$ is the CP-net with no edges, where the local preferences are $\{a \succ \bar{a}, b \succ \bar{b}, c \succ \bar{c}\}$. The next proposition shows that the size of the majoritarian aggregation (as a CP-net) can be exponentially larger than the sum of the sizes of the individual CP-nets.

Proposition 2 *The largest ratio between the size of the majoritarian CP-net and the sum of the sizes of the individual CP-nets is $2^{p-2}/(p-1)$.*

Proposition 2 only states that in the worst case, representing the majoritarian aggregation as a CP-net is costly. However, because of Proposition 1, we can expect the dependency graph for the majoritarian CP-net, in practice, to have a limited number of edges, which can be computed easily by computing the union of the edges in individual CP-nets.

4 Hypercube Condorcet Winners

Let $P_L = (V_1, \ldots, V_n)$ be a profile of linear orders. We recall that an alternative c is a *Condorcet winner (CW)* for P_L if for every alternative $b \neq c$, c is preferred to b in a majority of V_j's. Let $CW(P_L)$ denote the set of Condorcet winner(s).

Definition 3 *$\vec{x}$ is a* hypercube Condorcet winner *(HCW)*[1] *for $P = (\mathcal{N}_1, \ldots, \mathcal{N}_n)$ if for every neighbor $\vec{y}$ of $\vec{x}$, we have $\vec{x} \succ_{M(P)} \vec{y}$.*

[1]This was called "local Condorcet winner" [12]. We use a different name to emphasize that it is for binary multi-issue domains.

Let HCW(P) denote the set of all HCWs in the H-profile P. We know from [9] (Theorem 1 and Corollary 1) that for any H-profile P, $\bigcup_{P_L \sim P} CW(P_L) \subseteq HCW(P)$, and the inclusion is strict for some H-profile P.[2]

The following useful lemma states that any (possibly cyclic) CP-net can be represented as the (weighted) majoritarian CP-net of an H-profile consisting of $2p-1$ consistent CP-nets, whose size is no more than $2p-1$ times larger.

Lemma 1 *Let $\mathcal{N}$ be a CP-net (consistent or not) over p binary variables. There exists an H-profile of $2p-1$ consistent CP-nets $P = (\mathcal{N}_1, \ldots, \mathcal{N}_{2p-1})$, such that (a) for every $j \leq n$, the size of $\mathcal{N}_j$ is no larger than the size of $\mathcal{N}$, (b) $\mathcal{N} = M(P)$, and (c) the weight of each edge in $W(P)$ is 1.*

Proposition 3 *Deciding whether there exists at least one hypercube Condorcet winner for an H-profile is* NP*-complete.*

Proof sketch: Membership is easy; the hardness proof uses the following reduction from EXISTENCE OF NON-DOMINATED OUTCOME IN A CP-NET, which is NP-complete (Theorem 1 in [4]): to any CP-net $\mathcal{N}$ we associate an H-profile P composed of the $2p-1$ CP-nets as in Lemma 1. $\vec{x}$ is an HCW for P iff $\vec{x}$ is undominated in $M(P)$. ■

Note that Li *et al.* address the practical computation of HCWs, via a reduction to SAT with cardinality formulas.

5 How Many HCWs?

HCW is an important solution concept for majoritarian CP-nets. Two questions naturally arise: (1) what is the probability that there exists at least one HCW in the majoritarian CP-net, and (2) what is the average number of HCWs?

The importance of these questions lies in the fact that if the set of HCWs is empty most of the time, then this casts doubt on the usefulness of the notion. On the other hand, if it is likely to contain many alternatives, then it has little decisive power, and listing all HCWs may even result in exponentially large output. Fortunately, we argue that, at least under the following natural assumption on the distribution over profiles, neither is the case. We assume that any CP-net (consistent or not) is drawn with the same probability, which is $1/2^{p \cdot 2^{p-1}}$. This distribution naturally induces a distribution for the majoritarian CP-net, where the direction of each edge is drawn i.i.d. uniformly at random (we recall that n is odd).

Proposition 4 *The maximum number of HCWs is 2^{p-1}.*

Theorem 1 *Suppose each CP-net is drawn i.i.d. uniformly from the set of all CP-nets. The probability that there exists at least one HCW in the majoritarian CP-net is at least $1 - \frac{1}{e}$.*

Proof: Under this distribution the direction of each edge in the majoritarian CP-net is generated i.i.d. uniformly at random. Now, for every $\vec{d} \in \mathcal{X}$, $\vec{d}$ is an HCW if for each $i \leq p$, the edge between $\vec{d}$ and $\vec{d}[\leftrightarrow i]$ goes in the direction $\vec{d} \to \vec{d}[\leftrightarrow i]$. Because the directions of these edges are drawn independently, the probability that $\vec{d}$ is an HCW is $\frac{1}{2^p}$. Let EVEN (respectively, ODD) denote the set of alternatives $\vec{d}$ such that $\sum_i d_i$ is even (respectively, odd). Since no two alternatives in EVEN are neighbours, for any pair of

[2]Their results were established for a slightly different notion, due to the handling of ties, but their proofs also hold for HCWs.

alternatives $\vec{d}, \vec{y}$ in EVEN, the events "$\vec{d}$ is an HCW" and "$\vec{y}$ is an HCW" are independent. Therefore, the probability that EVEN does not contain any HCW is $\left(1-\frac{1}{2^p}\right)^{2^{p-1}}$. Now, $\ln\left(1-\frac{1}{2^p}\right)^{2^{p-1}} = 2^{p-1}\cdot\ln(1-\frac{1}{2^p})$, and $\ln(1-\frac{1}{2^p}) < -\frac{1}{2^p}$. Therefore, $\ln\left(1-\frac{1}{2^p}\right)^{2^{p-1}} < -\frac{1}{2}$ and the probability that EVEN (and, symmetrically, ODD) does not contain any HCW is at most $\frac{1}{\sqrt{e}}$.

Given that there is no HCW in ODD, intuitively the probability that there is no HCW in EVEN should be (slightly) less than the (unconditioned) probability that there is no HCW in EVEN. The existence of an HCW in ODD implies that all its neighbors, which are all in EVEN, cannot be HCWs; therefore, there seems to be a negative correlation between the events "at least one HCW in ODD" and "at least one HCW in EVEN," hence a positive correlation between "no HCW in ODD" and "at least one HCW in EVEN," and a negative correlation between the events "no HCW in ODD" and "no HCW in EVEN." Let Pr denote the uniform distribution over all (consistent or not) CP-nets. We have the following claim, which is proved by induction. We omit the proof due to the space constraint.

Claim 1 *Pr(No HCW in* EVEN*|No HCW in* ODD*)* $< \frac{1}{\sqrt{e}}$.

Therefore, the probability that there is no HCW overall is no more than $(\frac{1}{\sqrt{e}})^2 = \frac{1}{e}$, which means that the probability that there exists at least one HCW is at least $1-\frac{1}{e}$. ∎

Theorem 1 is quite positive, because $1-\frac{1}{e}$ is around 0.632, which means that the probability of having at least one HCW is significant. We conjecture that as $p\rightarrow\infty$, the probability actually *tends* towards $1-\frac{1}{e}$.

Proposition 5 *Suppose each CP-net is drawn i.i.d. uniformly from the set of all CP-nets. The expected number of HCWs in the majoritarian CP-net is 1.*

Proof: For each alternative, the probability that it is an HCW is $1/2^p$. Therefore, the expected number of HCWs is $\sum_{\vec{d}} E[\vec{d}$ is an HCW$] = 2^p\cdot\frac{1}{2^p} = 1$. ∎

By Markov's inequality, we obtain that the probability that there are at least k HCWs is at most $1/k$. This shows that the probability that we have many HCWs is low.

An arguably more natural distribution is the one where each voter's vote is a linear order and is drawn i.i.d. uniformly from the set of all linear orders over $\mathcal{X}$. For any profile P_L of linear orders over $\mathcal{X}$, we can still compute the majoritarian CP-net of the H-profile P, where P_L is an extension of P, and count the number of HCWs. This setting is known as *Impartial Culture*, which is by far the most common probability distribution in (computational or not) social choice theory for both theoretical analysis and simulations. The probability of the existence of one Condorcet winner (in non-multi-issue domains) under Impartial Culture has been investigated (see e.g. [10; 5]), but the exact probability is still unknown. Surprisingly, such probabilities are higher than one might expect.[3] For HCWs, we have a similar observation in the next section, where votes are drawn i.i.d. uniformly at random from all linear orders over all 2^p alternatives.

[3]For example, for 20 alternatives and 37 voters, the probability of the existence of a Condorcet winner is 0.33 by simulation [10].

6 Simulation Results

We run simulations to show the probability of the existence of at least one HCW, as well as the average number of HCWs, for the following two settings.

Setting 1 (Figure 2): The CP-nets are drawn i.i.d. uniformly; the number of issues ranges from 2 to 15; we generated 20000 samples for each setting. In Figure 2(a) we observe that the probability that there exists at least one HCW is almost always above $0.632 \approx 1-1/e$, and 0.632 seems to be the limit as p increases. This observation is consistent with Theorem 1, which only proves that $1-1/e$ is a lower bound. In Figure 2(b) we observe that the average number of HCWs is approximately 1 for any number of issues we have investigated. This observation is consistent with Proposition 5.

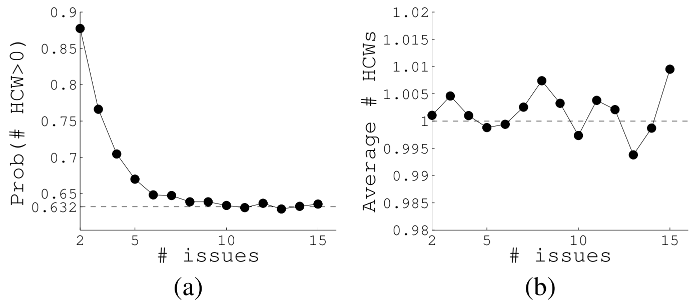

Figure 2: The direction of each edge in the hypercube is drawn i.i.d. uniformly. (a) The probability that there exists at least one HCW. (b) The average number of HCWs.

Li *et al.* [9] ran similar simulations to find the probability that there exists at least one weak HCW (an alternative that does not lose to any of its neighbors in pairwise elections). However, they randomly generated CP-nets where each issue has no more than 6 parents, while we do not have such a constraint on the number of parents of any issue.

Setting 2 (Figure 3): Each linear order over $\mathcal{X}$ is drawn i.i.d. uniformly (i.e., Impartial Culture); the number of issues ranges from 2 to 15; we tested the cases where the number of voters is 501, 601, 701, 801, 901, and 1001; for each setting we generated 10000 samples. In Figure 3(a) we observe that the probability that there exists at least one HCW is almost 1 when there are 6 issues or more, and these probabilities are insensitive to the number of voters (as long as it is larger than 501). In Figure 3(b) we show the log average number of HCWs when the number of voters is 1001 (results are similar for other numbers of voters we have tested). We observe that this number increases as the number of issues increases, and there seems to be a linear correlation. These results justify (experimentally) our conjecture that the probability of the existence of at least one HCW under Impartial Culture is larger than the probability of the existence of at least one HCW when the direction of each edge in the weighted majority graph is drawn i.i.d. uniformly. They also suggest that under Impartial Culture, we face the problem of having too many HCWs.

7 Hypercube-tournament Solution Concepts

All common tournament solution concepts can be naturally extended to hypercubes. Note that hypercubes correspond to partial tournaments on a multi-issue domain, and that extending tournament solution concepts to partial tournaments has

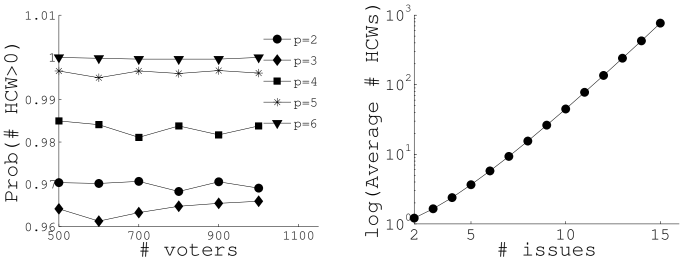

(a) (b)

Figure 3: The linear orders over $\mathcal{X}$ are drawn i.i.d. uniformly (Impartial Culture). (a) The probability that there exists at least one HCW. When the number of issues is larger than 6, the probability is almost one. (b) The log average number of HCWs.

been investigated in [3]. Therefore, the definitions below correspond to those in [3] if we ignore the hypercube structure of our partial tournaments. For each tournament solution T, we can define a voting rule over H-profiles by computing the (weighted) majoritarian CP-net and applying the tournament solution T to it. In previous work [12], we explored this idea only for *Schwartz*. Here, we focus on *Copeland*, *maximin*, and *Kemeny*. Let P be an H-profile, $\succ_P^{maj} = M(P)$, and $\vec{x} \in \mathcal{X}$.

7.1 H-Copeland

We recall that the Copeland score of an alternative, with respect to a profile P, is the number of alternatives it beats in pairwise elections,[4] and a Copeland winner is an alternative with the maximum Copeland score. Let Copeland(P) denote the set of all Copeland winners for P.

Given an H-profile P, we define the *hypercube Copeland score* of an alternative to be the number of its neighbors it beats in pairwise elections. This seems the most intuitive way of defining a variant of the Copeland score when we know only the hypercube. In both cases (Copeland and H-Copeland), we maximize the number of outgoing edges in the dominance graph, the difference is that in H-Copeland we use the majority hypercube instead of the full majority graph.

Definition 4 (H-Copeland) *Let $HC(\vec{x}, P) = \#\{i \le p : \vec{x} \succ_P^{maj} \vec{x}[\leftrightarrow i]\}$. $\vec{x}$ is a* hypercube Copeland winner *for P if it maximizes $HC(\vec{x}, P)$. Let $H_{\text{Copeland}}(P)$ denote the set of hypercube Copeland winners for P.*

Proposition 6 *If $HCW(P) \neq \emptyset$, then $H_{\text{Copeland}}(P) = HCW(P)$.*

The next proposition studies the relationship between H-Copeland winners and Copeland winners of the same profile.

Proposition 7 *There exists an H-profile P and an extension P_L of P such that Copeland$(P_L) \not\subseteq H_{\text{Copeland}}(P)$. There exists an H-profile P' such that $H_{\text{Copeland}}(P') \not\subseteq \bigcup_{P'_L \sim P'}$ Copeland(P'_L).*

Proof: Let P be composed of the following three CP-nets.

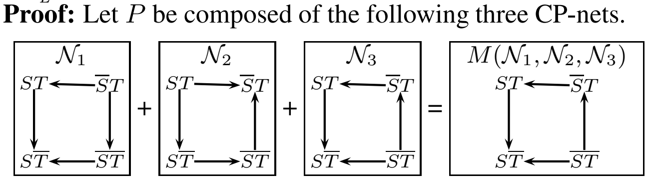

[4]There are various methods for counting pairwise ties, but here we assume an odd number of voters, so ties are excluded.

Let $V_1 = \overline{S}T \succ ST \succ \overline{ST} \succ S\overline{T}$ be the extension of $\mathcal{N}_1$; the only extension of $\mathcal{N}_2$ is $V_2 = ST \succ S\overline{T} \succ \overline{ST} \succ \overline{S}T$; and the only extension of $\mathcal{N}_3$ is $V_3 = \overline{ST} \succ \overline{S}T \succ ST \succ S\overline{T}$. Therefore, ST is a Copeland winner for P_L but not an $\text{H}_{\text{Copeland}}$ winner for P. Let $P' = (\mathcal{N}'_1, \mathcal{N}'_2, \mathcal{N}'_3)$, where $\mathcal{N}'_1$ is $S : T \succ \bar{T}; \bar{S} : \bar{T} \succ T; T : S \succ \bar{S}; \bar{T} : \bar{S} \succ S$, $\mathcal{N}'_2$ is $S \succ \bar{S}; S : T \succ \bar{T}; \bar{S} : \bar{T} \succ T$ and $\mathcal{N}'_3$ is $\bar{S} \succ S; T \succ \bar{T}; \bar{S} : \bar{T} \succ T$. $\overline{ST}$ is an $\text{H}_{\text{Copeland}}$ winner but not a Copeland winner in any extension of P' (P' only has one extension). ■

Proposition 8 *Checking whether $\vec{x}$ is an H-Copeland winner is* coNP*-complete.*

Proof: Membership in coNP is straightforward. Hardness is proved by a reduction from EXISTENCE OF NON-DOMINATED OUTCOME IN A CP-NET, which is NP-complete (Theorem 1 in [4]). Let $\mathcal{N}$ be the (possibly cyclic) CP-net in an instance of EXISTENCE OF NON-DOMINATED OUTCOME IN A CP-NET. We define a CP-net $\mathcal{N}'$ over $p+1$ issues as follows. Let $I = \{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$.

- The restriction of $\mathcal{N}'$ on D_I when $\mathbf{x}_{p+1} = 0$ is $\mathcal{N}$. That is, for any $i \le p$ and any $\vec{u} \in D_{I \setminus \{\mathbf{x}_i\}}$, $(0_i, 0_{p+1}, \vec{u}) \succ_{\mathcal{N}'} (1_i, 0_{p+1}, \vec{u})$ if and only if $(0_i, \vec{u}) \succ_{\mathcal{N}} (1_i, \vec{u})$.
- The restriction of $\mathcal{N}'$ on D_I when $\mathbf{x}_{p+1} = 1$ is a CP-net with no edges, such that for any $i \le p$, $1_i \succ 0_i$.
- $\mathbf{x}_{p+1}$ has no incoming edges, and $0_{p+1} \succ 1_{p+1}$.

We note that the size of $\mathcal{N}'$ is two times the size of $\mathcal{N}$ plus 1.

Now, by Lemma 1, there exists an H-profile P composed of $2p+1$ consistent CP-nets such that $M(P) = \mathcal{N}'$ and the size of P is polynomial in $\mathcal{N}'$. Let $\vec{x} = (1, \ldots, 1)$. Because $\vec{x}$ only loses to $(1, \ldots, 1, 0)$ in their pairwise election, we have $HC(\vec{x}, P) = p - 1$. Note that for any HCW $\vec{d} \in D_I$ in $\mathcal{N}$, $(\vec{d}, 0)$ is an HCW for P, so that $HC((\vec{d}, 0), P) = p$. Hence, $\vec{x}$ is an H-Copeland winner if and only if $\mathcal{N}$ has no HCW. ■

7.2 H-maximin

Let $N_P(c_i, c_j)$ denote the number of votes that rank c_i ahead of c_j in the profile P. We recall that the *maximin* rule selects the alternatives c maximizing $\min\{N_P(c, c') : c' \in \mathcal{X}, c' \neq c\}$. Let MM$(P)$ denote the set of all maximin winners for P. We define H-maximin as follows:

Definition 5 (H-maximin) *For any H-profile P, let $Hmm(\vec{x}, P) = \min_{i \le p} |\{j \le n : \vec{x} \succ_{\mathcal{N}_j} \vec{x}[\leftrightarrow i]\}|$. $\vec{x}$ is a* hypercube maximin winner *for P if it maximizes $Hmm(\vec{x}, P)$. Let $H_{\text{Maximin}}(P)$ denote the set of hypercube maximin winners for P.*

Note that $\text{H}_{\text{Maximin}}$ is a *weighted* H-tournament solution: it is determined from the weighted majoritarian CP-net.

Proposition 9 *If HCW$(P) \neq \emptyset$, then $H_{\text{Maximin}}(P) \subseteq$ HCW(P).*

The inclusion can be strict. Take the following H-profile $P = (\mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3)$, where $\mathcal{N}_1$ and $\mathcal{N}_2$ are $a : b \succ \bar{b}; \bar{a} : \bar{b} \succ b; b : a \succ \bar{a}; \bar{b} : \bar{a} \succ a$, and $\mathcal{N}_3$ is $a \succ \bar{a}; b \succ \bar{b}$. The weighted majoritarian CP-net is $a : b \succ \bar{b}$ (weight 3); $\bar{a} : \bar{b} \succ b$ (weight 2); $b : a \succ \bar{a}$ (weight 3); $\bar{b} : \bar{a} \succ a$ (weight 2). Both ab and $\bar{a}\bar{b}$ are HCWs, while only ab is an H-maximin winner.

Proposition 10 *There exists an H-profile P and an extension P_L of P such that MM$(P_L) \not\subseteq H_{\text{Maximin}}(P)$. There exists an H-profile P' such that $H_{\text{Maximin}}(P') \not\subseteq \bigcup_{P'_L \sim P'}$ MM(P'_L).*

Proof: The proposition can be proved with the same profiles shown in the proof of Proposition 7. ■

Proposition 11 *Checking whether $\vec{x}$ is an H-maximin winner is* coNP*-complete.*

Proof: Membership in coNP is straightforward. The hardness proof uses the same reduction as for H-Copeland. ■

7.3 H-Kemeny

For any pair of linear orders V and V', let $d(V, V')$ denote the number of pairs of alternatives $\{c, c'\}$ on which V and V' disagree. Given a profile $P = (V_1, \ldots, V_n)$, a *Kemeny consensus* for P is a linear order V that minimizes $\sum_{j=1}^{n} d(V, V_j)$; an alternative c is a *Kemeny winner* for P if c is ranked in the top position in some Kemeny consensus for P. Let Kemeny(P) denote the set of all Kemeny winners for P.

Now, we adapt the Kemeny rule to multi-issue domains as follows. The main difference from Kemeny is that the distance function for H-Kemeny only counts the number of edges in the hypercube on which two CP-nets differ.

Definition 6 *Given two CP-nets $\mathcal{N}$ and $\mathcal{N}'$ over p binary variables, the distance $d_H(\mathcal{N}, \mathcal{N}')$ between $\mathcal{N}$ and $\mathcal{N}'$ is the number of edges in the hypercube on which $\mathcal{N}$ and $\mathcal{N}'$ differ. Given an H-profile $P = (\mathcal{N}_1, \ldots, \mathcal{N}_n)$ and a CP-net $\mathcal{N}$, the distance between $\mathcal{N}$ and P is defined by $d_H(\mathcal{N}, P) = \sum_{1 \leq j \leq n} d_H(\mathcal{N}, \mathcal{N}_j)$. A* hypercube Kemeny consensus *for P is a consistent CP-net $\mathcal{N}$ over $\{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$ that minimizes the distance $d_H(\mathcal{N}, P)$; an alternative $\vec{x}$ is a* hypercube Kemeny winner *(HKW) for P if $\vec{x}$ is undominated in some hypercube Kemeny consensus for P. We denote by $H_{\text{Kemeny}}(P)$ the set of all HKWs for P.*

We note that H-Kemeny is a weighted H-tournament solution.

Proposition 12 $H_{\text{Kemeny}}(P) \supseteq$ *HCW(P).*

The inclusion can be strict, even when HCW$(P) \neq \emptyset$. (See the profile P in the proof of Proposition 7.)

Proposition 13 *There exists an H-profile P and an extension P_L of P such that Kemeny$(P_L) \not\subseteq H_{\text{Kemeny}}(P)$. There exists an H-profile P' such that $H_{\text{Kemeny}}(P') \not\subseteq \bigcup_{P'_L \sim P'}$ Kemeny(P'_L).*

Proof: The proposition can be proved with the same profiles shown in the proof of Proposition 7. ■

Next, we study the computational complexity of deciding whether an alternative $\vec{x}$ is an H-Kemeny winner. We adopt a compact representation for CP-net entries (already used in [4]) that uses "else" to represent all the valuations of the parents that are not mentioned in the CPT.

Proposition 14 *Checking whether $\vec{x}$ is an H-Kemeny winner is $\mathsf{P}_{||}^{\mathsf{NP}}$-hard.*

Proof sketch: Hardness is proved by reduction from KEMENY WINNER, which is $\mathsf{P}_{||}^{\mathsf{NP}}$-complete [6]. In an instance of KEMENY WINNER, we are given a profile $P_L = (V_1, \ldots, V_n)$ over $\mathcal{C}$ and an alternative $c \in \mathcal{C}$. We are asked whether c is a Kemeny winner. The idea behind the construction is that for $j \leq |\mathcal{C}|$, we first identify an alternative $\vec{e}_j$ in $\mathcal{X}$. Then, we "embed" the weighted majority graph of P_L into the weighted majoritarian CP-net, such that each edge $c_i \rightarrow c_j$ in the weighted majority graph of P_L corresponds to a path from $\vec{e}_i$ to $\vec{e}_j$. The length of each of these paths is 2 and they are disjoint (except the start point and the end point). With this construction, c_i is a Kemeny winner if and only if $\vec{e}_i$ is an H-Kemeny winner. The details are omitted. ■

8 Future Work

We can extend our majoritarian approach to multi-issue domains with non-binary issues as follows: for any issue i and any $\vec{d}_{-i}$ in D_{-i}, we compute the local (weighted) majority graph based on agents' local preferences over $\mathbf{x}_i$.

There are some open questions left for future research. Figure 2(a) suggests that when each CP-net is generated i.i.d. uniformly at random, the probability that there exists at least one HCW in the majoritarian CP-net goes to $1 - 1/e$ as $p \to \infty$. Figure 3(b) suggests that under the impartial culture assumption, the log average number of HCWs is linear in the number of issues. It would be desirable to find theoretical proofs for these observations.

Another important direction is to investigate other tournament solutions, such as *Slater*, *Banks*, and the *uncovered set*.

Acknowledgements

Vincent Conitzer and Lirong Xia acknowledge NSF CAREER 0953756 and IIS-0812113, and an Alfred P. Sloan fellowship for support. Jérôme Lang thanks the ANR project ComSoc (ANR-09-BLAN-0305). Lirong Xia is supported by a James B. Duke Fellowship. We thank all IJCAI-11 reviewers for their helpful comments and suggestions.

References

[1] C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. CP-nets: a tool for representing and reasoning with conditional ceteris paribus statements. *JAIR*, 21:135–191, 2004.

[2] S. Brams, D. Kilgour, and W. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.

[3] F. Brandt, F. Fischer, and P. Harrenstein. The computational complexity of choice sets. *Math. Log. Q.*, 55(4):444–459, 2009.

[4] C. Domshlak, S. Prestwich, F. Rossi, B. Venable, and T. Walsh. Hard and soft constraints for reasoning about qualitative conditional preferences. *J. Heuristics*, 12(4-5):263–285, 2006.

[5] W. Gehrlein and P. Fishburn. The probability of the paradox of voting: A computable solution. *Journal of Economic Theory*, 13(1):14–25, 1976.

[6] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *TCS*, 349(3):382–391, 2005.

[7] J. Lang and L. Xia. Sequential composition of voting rules in multi-issue domains. *MSS*, pages 304–324, 2009.

[8] M. Li, Q. B. Vo, and R. Kowalczyk. An efficient procedure for collective decision-making with cp-nets. In *Proc. ECAI-10*, pages 375–380, 2010.

[9] M. Li, Q. B. Vo, and R. Kowalczyk. Majority-rule-based preference aggregation on multi-attribute domains with structured preferences. In *Proc. AAMAS-11*, 2011. To appear.

[10] J. Pomeranz and R. Weil. The cyclical majority problem. *Commun. ACM*, 13:251–254, 1970.

[11] F. Rossi, K. Venable, and T. Walsh. mCP nets: representing and reasoning with preferences of multiple agents. In *Proc. AAAI-04*, pages 729–734, 2004.

[12] L. Xia, V. Conitzer, and J. Lang. Voting on multiattribute domains with cyclic preferential dependencies. In *Proc. AAAI-08*, pages 202–207, 2008.