

Voting on Multiattribute Domains with Cyclic Preferential Dependencies

Lirong Xia

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

Vincent Conitzer

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

Jérôme Lang

IRIT
Université Paul Sabatier
31062 Toulouse Cedex, France
lang@irit.fr

Abstract

In group decision making, often the agents need to decide on multiple attributes at the same time, so that there are exponentially many alternatives. In this case, it is unrealistic to ask agents to communicate a full ranking of all the alternatives. To address this, earlier work has proposed decomposing such voting processes by using local voting rules on the individual attributes. Unfortunately, the existing methods work only with rather severe domain restrictions, as they require the voters' preferences to extend acyclic CP-nets compatible with a common order on the attributes. We first show that this requirement is very restrictive, by proving that the number of linear orders extending an acyclic CP-net is exponentially smaller than the number of all linear orders. Then, we introduce a very general methodology that allows us to aggregate preferences when voters express CP-nets that can be cyclic. There does not need to be any common structure among the submitted CP-nets. Our methodology generalizes the earlier, more restrictive methodology. We study whether properties of the local rules transfer to the global rule, and vice versa. We also address how to compute the winning alternatives.

Introduction

In many real-life group decision making problems, the space of alternatives has a multiattribute (or combinatorial) structure. For instance, in multiple referenda (Brams, Kilgour, & Zwicker 1998), the inhabitants of some local community have to make a common decision over several related issues of local interest. For example, the inhabitants may need to decide on whether a swimming pool is built and on whether a tennis court is built. The decisions are not independent, because, perhaps, if a tennis court is built there is less time to go swimming; hence we cannot decide on the issues separately. As another example, the members of an association may have to elect a steering committee, composed of a president, a vice-president and a treasurer (Benoit & Kornhauser 1991). Again, the decisions are not independent: the voters may not like the president and the treasurer to be close friends (nor enemies). In both cases, the space of alternatives has a combinatorial structure.

When voting is used in artificial intelligence, these issues become even more pronounced. For example, agents may

have to vote over a joint plan or an allocation of tasks or resources. These alternative spaces are also combinatorial, and they are generally much larger than those considered in human domains. This is one of the problems that is driving the burgeoning field of *computational social choice* (for an introduction, see (Chevalerey *et al.* 2007)).

In classical social choice theory, voters are supposed to submit their preferences as *linear orders* over the set of alternatives, and then a *voting rule* is applied to select one alternative as the winner. If the set of alternatives has a multiattribute structure, then the number of alternatives is exponentially large, so it is unrealistic to ask voters to specify their preferences as (explicit) linear orders; hence, we cannot apply traditional voting rules in a straightforward way. A simple idea to cope with this problem consists of decomposing an election into a set of independent elections, each of which bears on a single attribute. This works to some extent when the preferences of voters are separable (that is, if voters' preferences over each attribute are independent from the values of other attributes), but it is impractical when they are not, because in this case a voter cannot specify preferences over a single attribute without knowing the values of the other attributes.

Instead of decomposing the election in parallel, it was proposed in (Lang 2007) to compose local voting rules *sequentially*: given a fixed directed acyclic graph G whose vertices are the attributes, these rules work by holding an election for each attribute based on a local voting rule, after the decision on its parents has been taken (which is always possible, because G is acyclic). For this sequential procedure to be applied, it is sufficient that each voter expresses an acyclic CP-net (Boutilier *et al.* 1999) whose DAG is G . This procedure is elicitation-friendly (it can be executed with an elicitation protocol which asks each voter only polynomially many queries) and easy to compute (provided that the local voting rules are). Its properties are studied in (Xia, Lang, & Ying 2007a). Unfortunately, this sequential procedure is applicable only to *G-legal profiles*, that is, profiles for which each vote extends some acyclic CP-net whose DAG is G . This shortcoming was partially addressed in (Xia, Lang, & Ying 2007b), which defines *order-independent* sequential composition of voting rules, allowing the profile of votes to be compatible with *any* linear order (or equivalently, any DAG) on the set of attributes. More precisely, the profile

must be compatible with *some* order over the attributes, but this order is not specified in the definition of the rule.

However, the domain restriction imposed by this order-independent sequential composition of voting rules is still severe: there must exist some (unspecified) directed acyclic graph G such that the profile is G -legal. (Other work in which CP-nets are used to combine the preferences of multiple agents also assumes acyclicity (Rossi, Venable, & Walsh 2004). That work allows voters to submit acyclic partial CP-nets—an extension of CP-nets. Then, a graph representing the group preference is built based on the submitted partial CP-nets and pairwise aggregation rules.) This can be seen as the conjunction of the following two restrictions: (1) each vote i is compatible with some acyclic graph G_i ; (2) for all voters i, j we have $G_i = G_j$. For (1), it is known that some linear orders on the set of alternatives do not extend any CP-net whose associated graph is acyclic, but so far there has been no real quantitative study of this lack of expressivity of acyclic CP-nets. It is only known that the *expressivity ratio* of separable preferences (a separable preference extends a CP-net whose associated graph has no edges) is exponentially small (Hodge 2006). That is, the fraction of linear orders that correspond to separable preferences is exponentially small. To our knowledge, there is no previous result about how small the expressivity ratio is for general acyclic CP-nets. In this paper, we prove that this ratio is exponentially small to show that restriction (1) is indeed very severe.

Restriction (2) was justified intuitively in (Lang 2007) by the belief that in many practical contexts, there exists a natural order of the attributes for the voters to decide on the attributes in that order. Unfortunately, this is not always true, and the approach fails as soon as one voter disagrees with the order.

In this paper, we drop these two assumptions. We define a new family of voting rules on multiattribute domains that are applicable to *any* profile of linear orders. Each voter only needs to submit a CP-net; the CP-net can be cyclic, and there does not need to be any relationship among the voters' CP-nets.¹ These voting rules are parameterized by: (1) the local voting rules that are used on individual attributes—we will use these to define a particular graph on the set of alternatives; and (2) a choice set function T that selects the winners based on this induced graph. We show that if T satisfies a very natural assumption, then the voting rules induced by T extend the order-independent sequential composition of local rules from (Xia, Lang, & Ying 2007b). We study whether properties of the local rules transfer to the global rule, and vice versa. Then, we focus on a particular choice set function, namely the *Schwartz set* (Schwartz 1970), which has been argued to be the largest reasonable choice set for tournament graphs (Laslier 1997). For the Schwartz set, we study how to compute the winners under this methodology.

¹Earlier work has also considered social choice for potentially cyclic CP-nets (Purrington & Durfee 2007). However, that approach does not apply to all possible (cyclic) CP-nets.

Preliminaries

Let \mathcal{X} be a set of alternatives (or candidates). A *linear order* is a transitive, antisymmetric, and total relation. The set of all linear orders on \mathcal{X} is denoted by $L(\mathcal{X})$. In contrast, a *partial order* is a transitive and antisymmetric relation. A profile on \mathcal{X} of n votes consists of n linear orders on \mathcal{X} . A *voting rule* maps each profile to an alternative, and a *voting correspondence* maps each profile to a subset of the alternatives. In this paper, the set of all alternatives is a *multiattribute domain*. That is, let $\mathfrak{A} = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$ be a set of attributes, where each attribute \mathbf{x}_i takes values in a *local attribute domain* D_i . An alternative is uniquely identified by the combination of its attribute values, that is, $\mathcal{X} = D_1 \times \dots \times D_p$.

CP-nets (Boutilier *et al.* 1999) are a compact representation for partial orders. A CP-net \mathcal{N} over \mathcal{X} consists of two parts: (a) a directed graph $G = (\mathfrak{A}, E)$ and (b) a set of conditional linear preferences $\succeq_{\vec{u}}^i$ over D_i , for any setting \vec{u} of the parents of \mathbf{x}_i in G . Let $CPT(\mathbf{x}_i)$ be the set of all conditional linear preferences on D_i ; this is called a *conditional preference table* (CPT). When G is acyclic, \mathcal{N} is said to be an *acyclic CP-net*.

CP-net \mathcal{N} induces a partial preorder $\succeq_{\mathcal{N}}$ such that for any $a_i, b_i \in D_i$, any setting \vec{u} of the set of parents of \mathbf{x}_i (denoted by $Par_G(\mathbf{x}_i)$), and any setting \vec{z} of $\mathfrak{A} - Par_G(\mathbf{x}_i) - \{\mathbf{x}_i\}$, $(a_i, \vec{u}, \vec{z}) \succeq_{\mathcal{N}} (b_i, \vec{u}, \vec{z})$ if and only if $a_i \succeq_{\vec{u}}^i b_i$. We note that when \mathcal{N} is acyclic, $\succeq_{\mathcal{N}}$ is transitive and asymmetric, that is, a strict partial order. Given a DAG G on \mathfrak{A} , a CP-net \mathcal{N} is *compatible* with G if its graph $G_{\mathcal{N}}$ is compatible with G , which means that $G_{\mathcal{N}} \subseteq G$.

A linear order V *extends* a CP-net \mathcal{N} , denoted by $V \sim \mathcal{N}$, if it extends the partial order that \mathcal{N} induces. For any setting \vec{u} of $Par_G(\mathbf{x}_i)$, let $V|_{\mathbf{x}_i:\vec{u}}$ be the restriction of V to \mathbf{x}_i , given \vec{u} . That is, $V|_{\mathbf{x}_i:\vec{u}}$ is the linear order $\succeq_{\vec{u}}^i$. Given a DAG G on \mathfrak{A} , V is compatible with G if there exists a CP-net \mathcal{N} such that $V \sim \mathcal{N}$ and \mathcal{N} is compatible with G . If V is compatible with G , we also say that V is *G -legal*; we say V is *legal*, if it is G -legal for some acyclic graph G . The set of all G -legal votes is denoted by $Legal(G)$. A profile is G -legal if all of its votes are G -legal. For any linear order \mathcal{O} on \mathfrak{A} , we define $G_{\mathcal{O}}$ to be the *graph induced by \mathcal{O}* —that is, there is an edge $(\mathbf{x}_i, \mathbf{x}_j)$ in $G_{\mathcal{O}}$ if and only if $\mathbf{x}_i \succ_{\mathcal{O}} \mathbf{x}_j$. We note that for any DAG G , a linear order \mathcal{O} can be found such that $G \subseteq G_{\mathcal{O}}$, which means that any G -legal profile is also $G_{\mathcal{O}}$ -legal.

Given an order $\mathcal{O} = \mathbf{x}_1 \succ \dots \succ \mathbf{x}_p$ and a set of local rules $\{r_1, \dots, r_p\}$ —that is, for any $i \leq p$, r_i is a voting rule on D_i —the *(fixed-order) sequential composition* of local rules r_1, \dots, r_p w.r.t. \mathcal{O} (Lang 2007), denoted by $Seq(r_1, \dots, r_p)$, is defined for all $G_{\mathcal{O}}$ -legal profiles as follows: $Seq(r_1, \dots, r_p)(P) = (d_1, \dots, d_p)$, where for each $i \leq p$, $d_i = r_i(P|_{\mathbf{x}_i:d_1 \dots d_{i-1}})$. The *(fixed-order) sequential composition* of local correspondences c_1, \dots, c_p w.r.t. \mathcal{O} is defined similarly. More precisely, for any $G_{\mathcal{O}}$ -legal profile P , $Seq(c_1, \dots, c_p)(P) = (d_1, \dots, d_p)$, where for each $i \leq p$, $d_i \in c_i(P|_{\mathbf{x}_i:d_1 \dots d_{i-1}})$. The *order-independent sequential composition* of local rules (Xia, Lang, & Ying 2007b), denoted by $Seq^{OI}(r_1, \dots, r_p)$, extends the domain of sequential composition of local rules to the set of all legal profiles P , which means that the order \mathcal{O} is not held fixed

in the definition. For any permutation σ on $\{1, \dots, p\}$, let $\mathcal{O} = \mathbf{x}_{\sigma(1)} > \dots > \mathbf{x}_{\sigma(p)}$. Then, for any $G_{\mathcal{O}}$ -legal profile P , $Seq^{OI}(r_1, \dots, r_p)(P) = Seq(r_{\sigma(1)}, \dots, r_{\sigma(p)})(P)$. The order-independent sequential composition of local correspondences is defined similarly. All of the above voting rules (correspondences) are well-defined because it has been shown in (Lang 2007) that for any G -legal profile, the set of winners is the same for all \mathcal{O} such that $G \subseteq G_{\mathcal{O}}$.

To study the properties of a voting rule, some common voting criteria are examined. We say a voting rule r satisfies

- **anonymity**, if the output of the rule is insensitive to the names of the voters;
- **homogeneity**, if for any vote V and any $n \in \mathbb{N}$, $n > 0$, $r(V) = r(nV)$, where nV is the profile composed of n copies of V ;
- **neutrality**, if the output of the rule is insensitive to the names of the alternatives;
- **monotonicity**, if for any profile $P = (V_1, \dots, V_N)$ and another profile $P' = (V'_1, \dots, V'_N)$ such that each V'_i is obtained from V_i by raising only $r(P)$, we have $r(P') = r(P)$;
- **consistency**, if for two disjoint profiles P_1, P_2 , $r(P_1) = r(P_2)$, then $r(P_1 \cup P_2) = r(P_1) = r(P_2)$;
- **participation**, if for any profile P and any vote V , $r(P \cup \{V\}) \succeq_V r(P)$;
- **Pareto efficiency**, if for any profile P , there is no alternative c that is preferred to $r(P)$ by all the voters.

Acyclic CP-nets are restrictive

In this section, we show that even when each local domain is binary, the number of legal linear orders—the set of all linear orders \succ for which there is some acyclic CP-net that \succ extends—is exponentially smaller than the number of all linear orders. Let $CP(\mathcal{X}) = \{V \in L(\mathcal{X}) : \text{There exists a CP-net } \mathcal{N} \text{ such that } V \sim \mathcal{N}\}$. That is, $CP(\mathcal{X}) = \cup_{\mathcal{O}} Legal(G_{\mathcal{O}})$.

Theorem 1 *If $\mathcal{X} = \{0, 1\}^p$, then $\frac{|CP(\mathcal{X})|}{|L(\mathcal{X})|} \leq \frac{p!}{2^{2^{p-2}}}$.*

Sketch of proof. We construct a set of exponentially many permutations on the set of alternatives, and we prove that for any two different linear orders compatible with the same order over attributes, for any two (not necessarily different) permutations in the set, if we apply the first permutation to the first linear order and the second permutation to the second linear order, the results are different. That is, for any linear order compatible with a given order \mathcal{O} , we can find a large set of corresponding linear orders by applying the set of permutations to it; and the sets of linear orders corresponding to different $G_{\mathcal{O}}$ -legal linear orders are disjoint. This completes the proof. Details are omitted due to space constraints. \square

We note that $|\mathcal{X}| = 2^p$. Theorem 1 implies that the expressivity ratio of legal linear orders ($\frac{|CP(\mathcal{X})|}{|L(\mathcal{X})|}$) is $O((2^{0.2})^{-|\mathcal{X}|})$, which is exponentially small.

H-composition of local voting rules

In this section, we introduce a new framework for composing local voting rules. We call this *hypercube-wise composition (H-composition) of local voting rules*. The reason is

that the outcome only depends on preferences between alternatives that differ on only one attribute. We can visualize the set of all alternatives as a hypercube, and alternatives that differ on only one attribute are neighbors on this hypercube, as discussed in (Domshlak & Brafman 2002). An H-composition of local rules is defined for all profiles in which for each vote, there exists a (possibly cyclic) CP-net that it extends. In fact, for any linear order V on \mathcal{X} , there exists a CP-net \mathcal{N} such that V extends \mathcal{N} , so we can apply this to any linear orders (but also some partial orders). An H-composition of local rules is defined in two steps. In the first step, an *induced graph* is generated by applying local rules to the input profile. Then, in the second step, a *choice set* is selected based on the induced graph as the set of winners. We first define the induced graph of P w.r.t. local rules (or correspondences) r_1, \dots, r_p .

Definition 1 *Given a profile $P = (V_1, \dots, V_n)$ and local rules (or correspondences) r_1, \dots, r_p , the induced graph of P w.r.t. r_1, \dots, r_p , denoted by $IG(r_1, \dots, r_p)(P) = (\mathcal{X}, E)$, is defined by the following edges between alternatives. For any $i \leq p$, any setting \bar{x}_{-i} , let $C_i = r_i(P|_{\mathbf{x}_{-i}=\bar{x}_{-i}})$; for any $c_i \in C_i$, any $d_i \in D_i$, let there be an edge $(c_i, \bar{x}_{-i}) \rightarrow (d_i, \bar{x}_{-i})$.*

Example 1 *Suppose the multiattribute domain consists of two binary attributes: \mathbf{S} ranging over $\{S, \bar{S}\}$ and \mathbf{T} ranging over $\{T, \bar{T}\}$. The local rules are both the majority rule. Two votes V_1, V_2 and their induced graph $IG(Maj, Maj)(V_1, V_2)$ are illustrated in Figure 1. We note that V_1 is compatible with $\mathbf{S} > \mathbf{T}$, V_2 is compatible with $\mathbf{T} > \mathbf{S}$.*




Figure 1: Two votes and their induced graph.

Next, we define the dominance relation in a directed graph.

Definition 2 *Given a directed graph $G = (V, E)$, for any $v_1, v_2 \in V$, v_1 is said to dominate v_2 , denoted by $v_1 \succ_G v_2$, if and only if:*

1. *There is a directed path from v_1 to v_2 , and*
2. *There is no directed path from v_2 to v_1 .*

Let \succeq_G be the *transitive closure* of E , that is, \succeq_G is the minimum preorder such that if $(v_1, v_2) \in E$, then $v_1 \succeq_G v_2$. Then, another equivalent way to define the dominance relation is: \succ_G is the strict order induced by \succeq_G , that is, $v_1 \succ_G v_2$ if and only if $v_1 \succeq_G v_2$ and $v_2 \not\succeq_G v_1$.

We further define two kinds of special vertices in a directed graph G as follows. The first is a vertex that dominates all the other vertices, and the second is a vertex that dominates all its neighbors. We call the former the *global Condorcet winner* (which must be unique), and the latter a *local Condorcet winner*.

Now, we are ready to define the *choice set function*, which specifies a *choice set* for each graph.

Definition 3 *A choice set function T is a mapping from any graph to a subset of its vertices.*

We recall the definitions of some often-studied choice sets in a graph $G = (V, E)$.

The Schwartz set is the union of all maximal mutually connected subsets. A maximal mutually connected subset is a subset of vertices such that there is a path between any two vertices in the set, but there is no path from a vertex outside the set to a vertex inside the set.

The Smith set is the smallest set of vertices such that every vertex in the set dominates all the vertices outside the set.

The Copeland set: A vertex c 's Copeland score is the number of vertices that are dominated by c minus the number of vertices that dominate c . The vertices with the highest Copeland score are the winners.

Choice sets were originally introduced to make group decisions for tournament graphs. However, the definitions are easily extended to general graphs, as we did above. See (Laffond, Laslier, & Le Breton 1995) and (Brandt, Fischer, & Harrenstein 2007) for more discussion.

We say a choice set function T always chooses the global Condorcet winner, if for any graph $G = (V, E)$ in which c is the global Condorcet winner, we have $T(G) = \{c\}$. We say that T always chooses local Condorcet winners, if every local Condorcet winner is always in $T(G)$. We emphasize that here, the meaning of a Condorcet winner is different from traditional meaning of a Condorcet winner, which refers to an alternative that wins every pairwise election. We say that T is *monotonic*, if for any graph (V, E) , any $c \in T(V, E)$, and any (V, E') that is obtained from (V, E) by only flipping some of the incoming edges of c , we have $c \in T(V, E')$.

Theorem 2 (known/easy) *The Schwartz set, Smith set, and Copeland set are monotonic and always choose the global Condorcet winner; the Smith set and Schwartz set always choose local Condorcet winners.*

We are now ready to define the H-composition of local rules (correspondences).

Definition 4 *Let T be a choice set function. The Hypercubewise- T (H- T) composition of local rules r_1, \dots, r_p , denoted by $H_T(r_1, \dots, r_p)$, is defined as follows. For any profile P of linear orders on \mathcal{X} ,*

$$H_T(r_1, \dots, r_p)(P) = T(IG(r_1, \dots, r_p)(P))$$

That is, for any profile P , $H_T(r_1, \dots, r_p)$ computes the winner in the following two steps. First, the induced graph $IG(r_1, \dots, r_p)(P)$ is generated by applying local rules r_1, \dots, r_p to the restrictions of P to all the local domains. Then, in the second step, the set of winners is selected by the choice set function T from the induced graph $IG(r_1, \dots, r_p)(P)$.

From Theorem 1, the fact that all linear orders are consistent with some CP-net, and all CP-nets can be used under H-composition, we know that the domain of H-composition of local rules is exponentially larger than the domain of order-independent sequential composition. We note that to build the induced graph, only the preferences between adjacent alternatives are necessary. We note that the H-composition of local rules is a correspondence.

One interesting question is, how H-compositions are related to (order-independent) composition of local rules. Because the H-compositions are defined by both local rules and

the choice set, the relationship should also depend on local rules and the properties of the choice set. The next theorem states that if a choice set function T always chooses the global Condorcet winner, then H- T composition of local rules is an extension of order-independent sequential composition of the same local rules.

Theorem 3 *Let T be a choice set function that always chooses the global Condorcet winner. Then, for all legal profiles P , $H_T(r_1, \dots, r_p)(P) = Seq^{OI}(r_1, \dots, r_p)(P)$.*

Proof. We omit the details due to space constraint, but the main idea is that from the winner under order-independent sequential composition, a path to every other alternative can be found, and no path in the other direction exists. Hence this alternative is the global Condorcet winner, and will therefore be elected by H_T . \square

Corollary 1 *If T is the Schwartz set, Smith set, or Copeland set, then $H_T(r_1, \dots, r_p)$ is an extension of $Seq^{OI}(r_1, \dots, r_p)$.*

Local vs. global properties

Two interesting questions for order-independent (or fixed-order) sequential composition of local rules are (Xia, Lang, & Ying 2007a; 2007b): First, if all of the local rules satisfy some criterion, then does their order-independent (or fixed-order) sequential composition satisfy it as well? Second, vice versa, if the order-independent (or fixed-order) sequential composition of local rules satisfies some criterion, does each local rule satisfy it as well? For H-composition of local rules, we can ask the same question. From Theorem 3 we know that if T always chooses the global Condorcet winner, then H_T is an extension of Seq^{OI} . We can use this observation to carry over some of the results in (Lang 2007; Xia, Lang, & Ying 2007a; 2007b) to H_T . Specifically, if T always chooses the global Condorcet winner, and if a criterion transfers from the order-independent sequential composition of local rules to each local rule, then it also transfers for H- T composition; if a criterion does not transfer from local rules to their order-independent sequential composition, then it also does not transfer for H- T composition. Given the results in (Xia, Lang, & Ying 2007b), these observations allow us to resolve everything except how anonymity, homogeneity, monotonicity, and consistency transfer from local rules to their H- T composition. It is easy to see that anonymity and homogeneity always transfer. The next example shows that if T always chooses local Condorcet winners, then consistency does not transfer, even when the votes in the profile extend (possibly different) acyclic CP-nets.

Example 2 *Let $\mathcal{X} = \{x, \bar{x}\} \times \{y, \bar{y}\} \times \{z, \bar{z}\}$, and let all the local rules be the majority rule. Consider the following three CP-nets: (The non-specified parts of the CPTs do not matter.)*

\mathcal{N}_1 : compatible with $\mathbf{x} > \mathbf{y} > \mathbf{z}$, and $x > \bar{x}$, $x : y > \bar{y}$, $xy : z > \bar{z}$, $\bar{x} : \bar{y} > y$, $\bar{x}\bar{y} : \bar{z} > z$.

\mathcal{N}_2 : compatible with $\mathbf{y} > \mathbf{z} > \mathbf{x}$, and $y > \bar{y}$, $y : z > \bar{z}$, $yz : x > \bar{x}$, $\bar{y} : \bar{z} > z$, $\bar{y}\bar{z} : \bar{x} > x$.

\mathcal{N}_3 : compatible with $\mathbf{z} > \mathbf{x} > \mathbf{y}$, and $z > \bar{z}$, $z : x > \bar{x}$, $zx : y > \bar{y}$, $\bar{z} : \bar{x} > x$, $\bar{z}\bar{x} : \bar{y} > y$.

Criteria	Global to local	Local to global
Anonymity	Y	Y
Homogeneity	Y	Y
Neutrality	Y	N
Monotonicity	Y	Y for monotonic T
Consistency	Y	N if T always chooses local Condorcet winner
Participation	Y	N
Pareto efficiency	Y	N

Table 1: Local vs. global

For any V_1, V_2, V_3 extending $\mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3$, respectively, let $P = (V_1, V_2, V_3)$. Let $H_T(M) = H_T(Maj, Maj, Maj)$. Then $(\bar{x}, \bar{y}, \bar{z})$ is a local Condorcet winner, so it is in $H_T(M)(P)$. However, $H_T(M)(V_1) = H_T(M)(V_2) = H_T(M)(V_3) = (x, y, z)$, so $H_T(M)$ does not satisfy consistency (because otherwise, we must have $H_T(M) = \{(x, y, z)\}$, which we know is not the case).

The next proposition states that for any monotonic choice set function T , the monotonicity is transferred from local rules to their H- T composition. We omit the proof due to space constraint.

Proposition 1 *Let T be a monotonic choice set function. If all local rules $\{r_1, \dots, r_p\}$ satisfy monotonicity, then $H_T(r_1, \dots, r_p)$ also satisfies monotonicity.*

For choice sets T that always choose the global Condorcet winner, whether properties of local rules transfer to their H- T composition and vice versa is summarized in Table 1.

Computing H-Schwartz winners

As we mentioned in the introduction, among all choice sets, we are most interested in the Schwartz set, because first, it has been argued that the Schwartz set is the “largest” reasonable choice set for tournaments (Laslier 1997), and second, it corresponds to the *nondominated set* previously considered in the context of CP-nets (Boutilier *et al.* 2004). Recent work on the complexity of computing dominance relations in CP-nets shows that the *dominance* problem in a CP-net is hard (Goldsmith *et al.* 2008). More precisely, given a CP-net \mathcal{N} and two alternatives a and b , it is PSPACE-complete to compute whether or not $a \succ_{\mathcal{N}} b$. This can be used to show that checking membership in the Schwartz set is PSPACE-complete (Goldsmith *et al.* 2008).

Although computing the Schwartz set is hard in general, if the preferences are more structured it can be easy. As an extreme example, if the voters’ preferences extend an acyclic CP-net G , then H-Schwartz is equivalent to order-independent sequential composition of local rules, under which computing the winner is easy. In this section, we introduce a technique to exploit more limited independence information in the submitted votes for the purpose of computing the set of H-Schwartz winners.

Definition 5 *Let $\mathfrak{A}_1 \dots \mathfrak{A}_q$ ($q \leq p$) be a partition of the set of attributes \mathfrak{A} . We say a CP-net \mathcal{N} is compatible with $\mathfrak{A}_1 > \dots > \mathfrak{A}_q$ if for any $l \leq q$ and any $\mathbf{x} \in \mathfrak{A}_l$, $Par_G(\mathbf{x}) \subseteq \mathfrak{A}_1 \cup \dots \cup \mathfrak{A}_l$. A linear order V is compatible with $\mathfrak{A}_1 > \dots > \mathfrak{A}_q$ if there exists a CP-net \mathcal{N} such that V extends \mathcal{N} and \mathcal{N} is compatible with $\mathfrak{A}_1 > \dots > \mathfrak{A}_q$.*

One special case is the following: if the input profile is G -legal, and G is compatible with $\mathbf{x}_1 > \dots > \mathbf{x}_p$, then we can use the partition $\mathfrak{A}_1 = \{\mathbf{x}_1\}, \dots, \mathfrak{A}_p = \{\mathbf{x}_p\}$. We can use the following algorithm to find a partition with which the input profile P is compatible. Suppose that we already know the graphs of the CP-nets that the votes in P extend.

Algorithm 1

1. Let G_P be the union of all the graphs of the CP-nets that the votes in P extend.
2. Let $q = 0$; repeat step 3 until $G_P = \emptyset$.
3. Let $q \leftarrow q + 1$. Find a maximal mutually connected subset of G_P , and call it \mathfrak{A}_q . Remove \mathfrak{A}_q and all edges connecting it to G_P .
4. Output the partition $\mathfrak{A}_1 \cup \dots \cup \mathfrak{A}_q$.

This algorithm runs in time $O(p^3)$. Now we are ready to present the technique for computing the set of H-Schwartz winners more efficiently. Suppose the set of attributes can be partitioned into $\mathfrak{A}_1 \cup \mathfrak{A}_2$ so that P is compatible with $\mathfrak{A}_1 > \mathfrak{A}_2$.

Process 1

1. Compute the Schwartz set $H_{Schwartz}(r_{\mathfrak{A}_1})(P|_{\mathfrak{A}_1}) = W_1^1 \cup \dots \cup W_1^k$, where the W_1^i are the maximal mutually connected subsets in $IG(r_{\mathfrak{A}_1})(P|_{\mathfrak{A}_1})$.
2. For each $i \leq k$, let $IG(r_{\mathfrak{A}_2})(P|_{\mathfrak{A}_2:W_1^i}) = \bigcup_{w \in W_1^i} IG(r_{\mathfrak{A}_2})(P|_{\mathfrak{A}_2:w})$; then, compute the Schwartz set W_2^i for $IG(r_{\mathfrak{A}_2})(P|_{\mathfrak{A}_2:W_1^i})$.
3. Output $W_P = \bigcup_{i=1}^k W_1^i \times W_2^i$.

The next theorem states that we can compute the winners of $H_{Schwartz}(r_1, \dots, r_p)(P)$ by Process 1.

Theorem 4 $W_P = H_{Schwartz}(r_1, \dots, r_p)(P)$.

Sketch of proof. Let w_2 be a setting of \mathfrak{A}_2 and w_1, w'_1 be settings of \mathfrak{A}_1 such that w_1 and w'_1 differ only on one attribute. Since P is compatible with $\mathfrak{A}_1 > \mathfrak{A}_2$, we have that there is an edge from (w_1, w_2) to (w'_1, w_2) in $IG(r_{\mathfrak{A}})(P)$ if and only if there is an edge from w_1 to w'_1 in $IG(r_{\mathfrak{A}_1})(P|_{\mathfrak{A}_1})$. This implies the following claim.

Claim 1 If there is a path from (w_1, w_2) to (w'_1, w'_2) in $IG(r_{\mathfrak{A}})(P)$, then its projection on \mathfrak{A}_1 is a path from w_1 to w'_1 in $IG(r_{\mathfrak{A}_1})(P|_{\mathfrak{A}_1})$.

We note that for any $i \leq k$, any $w_1, w'_1 \in W_1^i$ such that there is a path from w_1 to w'_1 , and any $w_2 \in D_{\mathfrak{A}_2}$, there is a path from (w_1, w_2) to (w'_1, w_2) . Therefore, we have the following claim.

Claim 2 For any $i \leq k$, any $(w_1, w_2), (w'_1, w'_2) \in W_1^i \times D_{\mathfrak{A}_2}$, there is a path from (w_1, w_2) to (w'_1, w'_2) if and only if there is a path from w_2 to w'_2 in $IG(r_{\mathfrak{A}_2})(P|_{\mathfrak{A}_2:W_1^i})$.

Then, based on Claim 1 and Claim 2 we can prove that $W_P \subseteq H_{Schwartz}(r_1, \dots, r_p)(P)$ and $H_{Schwartz}(r_{\mathfrak{A}})(P) \subseteq W_P$, which means that $W_P = H_{Schwartz}(r_{\mathfrak{A}})(P)$. Details are omitted due to space constraints. \square

If the decomposition is $\mathfrak{A}_1 > \dots > \mathfrak{A}_q$ with $q > 2$, then Process 1 can be applied recursively to find the Schwartz set, as follows. First, compute the Schwartz set over $\mathfrak{A}_1 \times \mathfrak{A}_2$ by Process 1, then use this result to compute the Schwartz set over $(\mathfrak{A}_1 \times \mathfrak{A}_2) \times \mathfrak{A}_3$, etc. up to $(\mathfrak{A}_1 \times \dots \times \mathfrak{A}_{q-1}) \times \mathfrak{A}_q$.

The next example shows how Process 1 works.

Example 3 Let $\mathcal{X} = \{0, 1\}^3$, and let three votes V_1, V_2, V_3 extend three CP-nets. $V_1 : \mathbf{x}_1 > \mathbf{x}_2 > \mathbf{x}_3$, $V_2 : \mathbf{x}_2 > \mathbf{x}_1 > \mathbf{x}_3$, and $V_3 : \text{the variables are independent}$. Let the partition be $\mathfrak{A}_1 = \{\mathbf{x}_1, \mathbf{x}_2\}$, $\mathfrak{A}_2 = \{\mathbf{x}_3\}$. Then, $V_i, i = 1, 2, 3$ is compatible with $\mathfrak{A}_1 > \mathfrak{A}_2$. Suppose that $\{w_1, w'_1\} = H_{Schwartz}(r_1, r_2)(P|_{\{\mathbf{x}_1, \mathbf{x}_2\}})$, so that there is no path from w_1 to w'_1 , and vice versa. Also suppose that $\{w_2\} = H_{Schwartz}(r_3)(P|_{\mathbf{x}_3: \mathbf{x}_{-3}=w_1})$ and $\{w'_2\} = H_{Schwartz}(r_3)(P|_{\mathbf{x}_3: \mathbf{x}_{-3}=w'_1})$. Then, the winners are (w_1, w_2) and (w'_1, w'_2) .

The next theorem states that if P is compatible with $\mathfrak{A}_1 > \dots > \mathfrak{A}_q$, then the time required to compute the set of Schwartz winners by applying Process 1 is a polynomial function of the number of winners, the longest time it takes to apply local rules, p, n , and $\max |D_{\mathfrak{A}_i}|$.

Theorem 5 Suppose an n -vote profile P is compatible with $\mathfrak{A}_1 > \dots > \mathfrak{A}_q$. Let $d_{max} = \max_{i \leq q} |D_{\mathfrak{A}_i}|$. Let $t_{max}(n)$ be the longest time it takes to apply local rules on n inputs. Then, the running time of Process 1 is $O(apd_{max}(np + t_{max}(n)p + d_{max}))$, where a is the number of H-Schwartz winners.

Usually $t_{max}(n)$ is polynomial. Therefore, the computational complexity of Process 1 mainly comes from the number of H-Schwartz winners, and the size of the largest partition d_{max} . Effectively, Process 1 gives a smooth tradeoff between computational efficiency and generality, and is certainly better than applying standard voting rules directly.

Conclusion and future work

In artificial intelligence, agents often need to jointly choose an alternative from a combinatorial domain (for example, the set of all plans or the set of all allocations of tasks or resources). In such settings, it is unrealistic to ask agents to communicate a full ranking of all the alternatives, because there are exponentially many—but this is what standard voting rules require. To address this, earlier work has proposed decomposing such voting processes by using local voting rules on the individual attributes. Unfortunately, the existing methods work only with rather severe domain restrictions: they require the voters' preferences to extend acyclic CP-nets compatible with a common order on the attributes. We showed that this requirement is very restrictive, by proving that the number of linear orders extending an acyclic CP-net is exponentially smaller than the number of all linear orders. Then, we introduced a very general methodology that allows us to aggregate preferences when voters express CP-nets that can be cyclic. There does not need to be any common structure among the submitted CP-nets. Our methodology generalizes the earlier, more restrictive methodology. We studied whether properties of the local rules transfer to the global rule, and vice versa. We also address how to compute the winning alternatives. We believe this methodology constitutes a significant step forward in using voting in realistic AI settings.

Further research can extend the idea of H-composition to other local aggregation functions, for example, local *social welfare functions*, which rank all the alternatives rather than

just produce a winner. Another extension would be to allow voters submit preorders. In that case, we can apply H-composition of *local* choice sets to determine winners. A final possibility is to consider choice sets other than the Schwartz set, such as the Smith and Copeland sets.

Acknowledgements

We thank Paul Harrenstein, Francesca Rossi, Kristen Brent Venable, and anonymous reviewers for helpful discussions and comments. Lirong Xia is supported by a James B. Duke Fellowship, Vincent Conitzer is supported by an Alfred P. Sloan Research Fellowship, and Jérôme Lang is supported by the French ANR project ANR-05-BLAN-0304 "Preference Handling and Aggregation in Combinatorial Domains".

References

- Benoit, J.-P., and Kornhauser, L. 1991. Voting simply in the election of assemblies. Technical report 91-32, C.V. Starr Center for Applied Economics.
- Boutilier, C.; Brafman, R.; Hoos, H.; and Poole, D. 1999. Reasoning with conditional *ceteris paribus* statements. In *UAI'99*, 71–80. Morgan Kaufmann.
- Boutilier, C.; Brafman, R.; Domshlak, C.; Hoos, H.; and Poole, D. 2004. CP-nets: a tool for representing and reasoning with conditional *ceteris paribus* statements. *JAIR* 21:135–191.
- Brams, S.; Kilgour, D.; and Zwicker, W. 1998. The paradox of multiple elections. *Social Choice and Welfare* 15(2):211–236.
- Brandt, F.; Fischer, F.; and Harrenstein, P. 2007. The computational complexity of choice sets. In *TARK'07*.
- Chevalere, Y.; Endriss, U.; Lang, J.; and Maudet, N. 2007. A short introduction to computational social choice. In *SOFSEM'07*.
- Domshlak, C.; Brafman, R.I. 2002. CP-nets: Reasoning and Consistency Testing. In *KR'02*, 121–131.
- Goldsmith, J.; Lang, J.; Truszczynski, M.; and Wilson, N. 2008. The computational complexity of dominance and consistency in CP-nets. Extended version; early version appeared in *IJCAI'05*, 935–940.
- Hodge, J. 2006. Permutations of separable preference orders. *Discrete applied mathematics*. 154:1478–1499.
- Laffond, G.; Laslier, J. F.; and Le Breton, M. 1995. Condorcet choice correspondences: A set-theoretical comparison. *Mathematical Social Sciences* 30:23–35.
- Lang, J. 2007. Vote and aggregation in combinatorial domains with structured preferences. In *IJCAI'07*, 1366–1371.
- Laslier, J.-F. 1997. *Tournament Solutions and Majority Voting*.
- Purrington, Keith, and Durfee, Edmund H. 2007. Making Social Choices from Individuals' CP-nets. In *AAMAS'07*, 129.
- Rossi, F.; Venable, K.; and Walsh, T. 2004. mCP nets: representing and reasoning with preferences of multiple agents. In *AAAI'04*, 729–734.
- Schwartz, T. 1970. On the possibility of rational policy evaluation. *Theory and Decision* 1:89–106.
- Xia, L.; Lang, J.; and Ying, M. 2007a. Sequential voting rules and multiple elections paradoxes. In *TARK'07*.
- Xia, L.; Lang, J.; and Ying, M. 2007b. Strongly decomposable voting rules on multiattribute domains. In *AAAI'07*.