CP-nets and Nash equilibria

Krzysztof R. Apt[1,2,3], Francesca Rossi[4], and Kristen Brent Venable[4]

[1] School of Computing, National University of Singapore
[2] CWI, Amsterdam
[3] University of Amsterdam, the Netherlands
[4] Department of Pure and Applied Mathematics, University of Padova, Italy

Abstract. CP-nets are a natural way to express qualitative and conditional preferences. Here we relate them to a natural extension of the classical notion of a strategic game in which parametrized strict linear orderings are used instead of payoff functions. We show then that the optimal outcomes of a CP-net are exactly the Nash-equilibria of an appropriately defined strategic game in the above sense. This allows us to use the techniques of game theory to search for optimal outcomes of CP-nets and vice-versa, to use techniques developed for CP-nets to search for Nash equilibria of the considered games. We believe this is a first promising step towards a fruitful cross-fertilization between these two research areas, from artificial intelligence and microeconomic theory.

1 Introduction

CP-nets (Conditional Preference nets) are an elegant formalism for representing conditional and qualitative preferences, see [4, 7, 3]. They model such preferences under a *ceteris paribus* (that is, 'all else being equal') assumption. Preference elicitation in such a framework appears to be natural and intuitive.

Research on CP-nets focused on its modeling capabilities and algorithms for solving various natural problems related to their use. Also, computational complexity of these problems was extensively studied. An outcome of a CP-net is an assignment of values to its variables. One of the fundamental problems is that of finding an optimal outcome, i.e., the one that cannot be improved in presence of the adopted preference statements. This is in general a complex problem since it was found that finding optimal outcomes and testing for their existence is NP-hard in general. In contrast, for acyclic CP-nets this is an easy problem which can be solved by a linear time algorithm.

The aim of this paper to show the relationship between CP-nets and game theory, and how game-theoretic techniques developed for the analysis of strategic games can be fruitfully used to study CP-nets. To this end, we introduce a generalization of the customary strategic games (see, e.g., [10],) in which each player has to his disposal a strict preference relation on his set of strategies, parametrized by a joint strategy of his opponents. We call such games *strategic games with parametrized preferences*.

The cornerstone of our approach are two results closely relating CP-nets to such games. They show that the optimal outcomes of a CP-net are exactly the Nash-equilibria of an appropriately defined strategic game with parametrized preferences. This allows us to transfer techniques of game theory to CP-nets. These results are based on the observation that the ceteris-paribus principle, typical of CP-nets, implies that any optimal outcome is worsened if any worsening change to any variable is made. This is exactly the idea of a Nash equilibrium, and thus the results follow easily once this is clear.

Notice that these results do not hold if we consider other ways to express preferences, for example soft constraints [2] rather than CP-nets. Consider for instance fuzzy constraints [8], where the preference of a solution is the minimal preference over all the constraints, and solutions with higher values are better. It is possible to show that, in a fuzzy constraint problem, there can be optimal solutions which are not Nash equilibria of corresponding games, and viceversa.

To find Nash equilibria in strategic games, many reduction techniques have been studied which reduce the game by eliminating some players' strategies, thus obtaining a smaller game. We introduce a counterpart of one of such game-theoretic technique that allows us to reduce a CP-net while maintaining its optimal outcomes. We also introduce a method of simplifying a CP-net by eliminating so-called redundant variables from the variables parent sets. Both techniques simplify the search for optimal outcomes of a CP-net.

In the other direction, we can use the techniques developed to reason about optimal outcomes of a CP-net in search for Nash equilibria of strategic games with parametrized preferences. These techniques, as recently shown in [5, 12], involve the use of the customary constraint solving techniques. In fact, it has been shown that the optimal outcomes of any CP-nets, even a cyclic one, can be found by just solving a set of hard constraints. Thus hard constraint solving is enough to find also Nash equilibria in strategic games. In particular, when the CP-net corresponding to a given game is acyclic, we know it has a unique optimal outcome that can be found in linear time. This allows us to find easily the unique Nash equlibrium of the given game.

The paper is organized as follows. Section 2 provides the basic definitions of CP-nets. Then, Section 3 introduces our generalized notion of games, Section 4 shows how to pass from a CP-net to a game, and Section 5 handles the opposite direction. Then, Section 6 introduces the concept of reduced CP-nets, and Sections 7 and 8 show how to exploit techniques developed for CP-nets in games and vice-versa. Finally, Section 9 gives an informal idea of the relationship between soft constraints and game theory, and 10 summarizes the main contributions of the paper and gives some hints for future work.

2 CP-nets

CP-nets [4, 3] (for Conditional Preference nets) are a graphical model for compactly representing conditional and qualitative preference relations. They exploit conditional preferential independence by decomposing an agent's preferences via the **ceteris paribus** (cp) assumption. Informally, CP-nets are sets of ***ceteris paribus (cp)*** preference statements. For instance, the statement *"I prefer red wine to white wine if meat is served."* asserts that, given two meals that differ *only* in the kind of wine served *and* both containing meat, the meal with a red wine is preferable to the meal with a white wine. On the other hand, this statement does not order two meals with a different main course. Many users' preferences appear to be of this type.

CP-nets bear some similarity to Bayesian networks. Both utilize directed graphs where each node stands for a domain variable, and assume a set of *features* (variables) $F = \{X_1, \ldots, X_n\}$ with the corresponding finite domains $\mathcal{D}(X_1), \ldots, \mathcal{D}(X_n)$. For each feature X_i, a user specifies a (possibly empty) set of ***parent*** features $Pa(X_i)$ that can affect her preferences over the values of X_i. This defines a ***dependency graph*** in which each node X_i has $Pa(X_i)$ as its immediate predecessors.

Given this structural information, the user explicitly specifies her preference over the values of X_i for *each complete assignment* on $Pa(X_i)$. This preference is assumed to take the form of a linear ordering over $\mathcal{D}(X_i)$ [4, 3]. Each such specification is called below a ***preference statement*** for the variable X_i. These conditional preferences over the values of X_i are captured by a *conditional preference table* which is annotated with the node X_i in the CP-net. An ***outcome*** is an assignment of values to the variables with each value taken from the corresponding domain.

As an example, consider a CP-net whose features are A, B, C and D, with binary domains containing f and $\overline{f}$ if F is the name of the feature, and with the following preference statements:

$d : a \succ \overline{a},\ \overline{d} : a \succ \overline{a},$
$a : b \succ \overline{b},\ \overline{a} : \overline{b} \succ b,$
$b : c \succ \overline{c},\ \overline{b} : \overline{c} \succ c,$
$c : d \succ \overline{d},\ \overline{c} : \overline{d} \succ d.$

Here the preference statement $d : a \succ \overline{a}$ states that $A = a$ is preferred to $A = \overline{a}$, given that $D = d$. ¿From the structure of these preference statements we see that $Pa(A) = \{D\}, Pa(B) = \{A\}, Pa(C) = \{B\}, Pa(D) = \{C\}$ so the dependency graph is cyclic.

An ***acyclic*** CP-net is one in which the dependency graph is acyclic. As an example, consider a CP-net whose features and domains are as above and with the following preference statements:

$a \succ \overline{a},$
$b \succ \overline{b},$
$(a \wedge b) \vee (\overline{a} \wedge \overline{b}) : c \succ \overline{c},\ (a \wedge \overline{b}) \vee (\overline{a} \wedge b) : \overline{c} \succ c,$
$c : d \succ \overline{d},\ \overline{c} : \overline{d} \succ d.$

Here, the preference statement $a \succ \overline{a}$ represents the unconditional preference for $A = a$ over $A = \overline{a}$. Also each preference statement for the variable C is a actually an abbreviated version of two preference statements. In this example we have $Pa(A) = \emptyset, Pa(B) = \emptyset, Pa(C) = \{A, B\}, Pa(D) = \{C\}$.

The semantics of CP-nets depends on the notion of a ***worsening flip***. A worsening flip is a transition between two outcomes that consists of a change in the value of a single variable to one which is less preferred in the unique preference statement for that variable. By analogy we define an ***improving flip***. For example, in the acyclic CP-net above, passing from $abcd$ to $ab\overline{c}d$ is a worsening flip since c is better than $\overline{c}$ given a and b. We say that an outcome α is ***better*** than the outcome β (or, equivalently, β is ***worse*** than α), written as $\alpha \succ \beta$, iff there is a chain of worsening flips from α to β. This definition induces a strict preorder over the outcomes. In the above acyclic CP-net the outcome $\overline{a}b\overline{c}\overline{d}$ is worse than $abcd$.

An ***optimal*** outcome is one for which no better outcome exists. In general, a CP-net does not need to have an optimal outcome. As an example consider two features A and B with the respective domains $\{a, \overline{a}\}$ and $\{b, \overline{b}\}$ and the following preference statements:

$a : b \succ \overline{b},\ \overline{a} : \overline{b} \succ b,$

$b : \overline{a} \succ a,\ \overline{b} : a \succ \overline{a}.$

It is easy to see that then

$ab \succ a\overline{b} \succ \overline{a}\overline{b} \succ \overline{a}b \succ ab.$

Finding optimal outcomes and testing for optimality is NP-hard. However, in acyclic CP-nets there is a unique optimal outcome and it can be found in linear time [4, 3]. We simply sweep through the CP-net, following the arrows in the dependency graph, assigning at each step the most preferred value in the preference relation. For instance, in the CP-net above, we would choose $A = a$ and $B = b$, then $C = c$ and then $D = d$. The optimal outcome is therefore $abcd$.

Hard constraints are enough to find optimal outcomes of a CP-net and to test whether a CP-net has an optimal outcome. In fact, given a CP-net one can define a set of hard constraints (called ***optimality constraints***) such that their solutions are the optimal outcomes of the CP-net [5, 12].

Indeed, take a CP-net N and consider a linear ordering $\succ$ over the elements of the domain of a variable X used in a preference statement for X. Let φ be the disjunction of the corresponding assignments used in the preference statements that use $\succ$. Then for each of such linear ordering $\succ$ the corresponding optimality constraint is $\varphi \rightarrow X = a_j$, where a_j is the undominated element of $\succ$. The optimality constraints $opt(N)$ corresponding to N consist of the entire set of such optimality constraints, each for one such linear ordering $\succ$.

For example, the preference statements $a \succ \overline{a}$ and $(a \wedge \overline{b}) \vee (\overline{a} \wedge b) : \overline{c} \succ c$ from the above CP-net map to the hard constraints $A = a$ and $(A = a \wedge B = \overline{b}) \vee (A = \overline{a} \wedge B = b) \rightarrow C = \overline{c}$, respectively.

It has been shown that an outcome is optimal in the strict preorder over the outcomes induced by a CP-net N iff it is a satisfying assignment for $opt(N)$.

A CP-net is ***eligible*** iff it has an optimal outcome. Even if the strict preorder induced by a CP-net has cycles, the CP-net may still be useful if it is eligible. All acyclic CP-nets are trivially eligible as they have a unique optimal outcome. We can thus test eligibility of any (even cyclic) CP-net by testing the consistency of the optimality constraints $opt(N)$. That is, a CP-net N is eligible iff $opt(N)$ is consistent.

3 Strategic games with parametrized preferences

In this section we introduce a generalization of the notion of a strategic game used in game theory, see, e.g., [10].

First we need the concept of a ***preference*** on a set A which in this paper denotes a strict linear ordering on A. If $\succ$ is a preference, we denote by $\succeq$ the corresponding ***weak preference*** defined by: $a \succeq b$ iff $a \succ b$ or $a = b$.

Given a sequence of non-empty sets $S_1, \ldots, S_n$ and $s \in S_1 \times \ldots \times S_n$ we denote the ith element of s by s_i and use the following standard notation of game theory, where $I := i_1, \ldots, i_k$ is a subsequence of $1, \ldots, n$:

- $s_{-i} := (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$,
- $s_I := (s_{i_1}, \ldots, s_{i_k})$,
- $(s'_i, s_{-i}) := (s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots, s_n)$, where we assume that $s'_i \in S_i$,
- $S_{-i} := S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_n$,
- $S_I := S_{i_1} \times \ldots \times S_{i_k}$.

In game theory it is customary to study strategic games in which the outcomes are numerical values provided by means of the payoff functions. A notable exception is [11] in which instead of payoff functions the linear quasi-orderings on the sets of joint strategies are used.

In our setup we adopt a different approach according to which each player has to his disposal a strict preference relation $\succ(s_{-i})$ on his set of strategies *parametrized* by a joint strategy s_{-i} of his opponents. So in our approach

- for each $i \in [1..n]$ player i has a finite, non-empty, set S_i of strategies available to him,
- for each $i \in [1..n]$ and $s_{-i} \in S_{-i}$ player i has a preference relation $\succ(s_{-i})$ on his set of strategies S_i.

In what follows such a ***strategic game with parametrized preferences*** (in short a ***game with parametrized preferences***, or just a ***game***) for n players is represented by a sequence

$$(S_1, \ldots, S_n, \succ(s_{-1}), \ldots, \succ(s_{-n})),$$

where each s_{-i} ranges over S_{-i}.

It is straightforward to transfer to the case of games with parametrized preferences the basic notions concerning strategic games. The following notions will be of importance for us (for the original definitions see, e.g., [11]), where G is a game with parametrized preferences specified as above.

- A strategy s_i is a ***best response*** for player i to a joint strategy s_{-i} of his opponents if $s_i \succeq(s_{-i})\ s'_i$, for all $s'_i \in S_i$.
- A strategy s_i is ***never a best response*** for player i if it is not a best response to any joint strategy s_{-i} of his opponents.
- A joint strategy s is a (pure) ***Nash equilibrium*** of G if each s_i is a best response to s_{-i}. Equivalently, s a Nash equilibrium if for all $i \in [1..n]$ and all $s'_i \in S_i$

$$s_i \succeq(s_{-i})\ s'_i.$$

- A strategy s'_i is ***strictly dominated*** by a strategy s_i if $s_i \succ(s_{-i})\ s'_i$, for all $s_{-i} \in S_{-i}$.

To clarify these definitions consider the classical Prisoner's dilemma strategic game represented by the following bimatrix representing the payoffs to both players:

	C_2	N_2
C_1	$3,3$	$0,4$
N_1	$4,0$	$1,1$

So each player i has two strategies, C_i (cooperate) and N_i (not cooperate), the payoff to player 1 for the joint strategy (C_1, N_2) is 0, etc. To represent this game as a game with parametrized preferences we simply stipulate that

$\succ(C_2) :=\ N_1 \succ C_1,\ \succ(N_2) :=\ N_1 \succ C_1,$

$\succ(C_1) :=\ N_2 \succ C_2,\ \succ(N_1) :=\ N_2 \succ C_2.$

These orderings reflect the fact that for each strategy of the opponent each player considers his 'not cooperate' strategy better than his 'cooperate' strategy. So for each player i his strategy C_i is strictly dominated by N_i, or (here) equivalently, his strategy C_i is never a best response. Further, (N_1, N_2) is a unique Nash equilibrium of this game with parametrized preferences.

Given a game with parametrized preferences

$$G := (S_1, \ldots, S_n, \succ(s_{-1}), \ldots, \succ(s_{-n})),$$

where each s_{-i} ranges over S_{-i}, and sets of strategies $S'_1, \ldots, S'_n$ such that $S'_i \subseteq S_i$ for $i \in [1..n]$, we say that

$$G' := (S'_1, \ldots, S'_n, \succ(s_{-1}), \ldots, \succ(s_{-n})),$$

where each s_{-i} now ranges over S'_{-i}, is a ***subgame*** of G, and identify in the context of G' each preference relation $\succ(s_{-i})$ with its restriction to S'_i.

We now introduce the following two notions of reduction between a game

$$G := (S_1, \ldots, S_n, \succ(s_{-1}), \ldots, \succ(s_{-n})),$$

where each s_{-i} ranges over S_{-i} and its subgame

$$G' := (S'_1, \ldots, S'_n, \succ(s_{-1}), \ldots, \succ(s_{-n})),$$

where each s_{-i} ranges over S'_{-i}:

- $G \rightarrow_{NBR} G'$
  when $G \neq G'$ and for all $i \in [1..n]$ each $s_i \in S_i \setminus S'_i$ is never a best response for player i in G,
- $G \rightarrow_S G'$
  when $G \neq G'$ and for all $i \in [1..n]$ each $s'_i \in S_i \setminus S'_i$ is strictly dominated in G by some $s_i \in S_i$.

In the literature it is customary to consider more specific reduction relations in which, respectively, *all* never best responses or *all* strictly dominated strategies are eliminated. The advantage of using the above versions is that we can prove the relevant property of both reductions by just one simple lemma, since by definition a strictly dominated strategy is never a best response and consequently $G \rightarrow_S G'$ implies $G \rightarrow_{NBR} G'$.

Lemma 1. *Suppose that $G \rightarrow_{NBR} G'$. Then s is a Nash equilibrium of G iff it is a Nash equilibrium of G'.*

Proof. ($\Rightarrow$) By definition each s_i is a best response to s_{-i} to G. So no s_i is eliminated in the reduction of G to G'.

($\Leftarrow$) Suppose s is not a Nash equilibrium of G. So some s_i is not a best response to s_{-i} in G. Let s'_i be a best response to s_{-i} in G. (s'_i exists since $\succ(s_{-i})$ is a linear ordering.)

So s'_i is not eliminated in the reduction of G to G' and s'_i is a best response to s_{-i} in G'. But this contradicts the fact that s is a Nash equilibrium of G'. $\square$

Theorem 1. *Suppose that $G \rightarrow^*_{NBR} G'$, i.e., G' is obtained by an iterated elimination of never best responses from the game G.*

(i) Then s is a Nash equilibrium of G iff it is a Nash equilibrium of G'.
(ii) If each player in G' has just one strategy, then the resulting joint strategy is a unique Nash equilibrium of G.

Proof.
(i) By the repeated application of Lemma 1.

(ii) It suffices to note that $(s_1, \ldots s_n)$ is a unique Nash equilibrium of the game in which each player i has just one strategy, s_i. $\square$

The above theorem allows us to reduce a game without affecting its (possibly empty) set of Nash equilibria or even, occasionally, to find its unique Nash equilibrium. In the latter case one says that the original game was ***solved*** by an iterated elimination of never best responses (or of strictly dominated strategies).

As an example let us return to the Prisoner's dilemma game with parametrized preferences defined above. In this game each strategy C_i is strictly dominated by N_i, so the game can be solved by either reducing it in two steps (by removing in each step one C_i strategy) or in one step (by removing both C_i strategies) to a game in which each player i has exactly one strategy, N_i.

Finally, let us mention that [9] and [13] proved that all iterated eliminations of strictly dominated strategies yield the same final outcome. An analogous result for the iterated elimination of never best responses was established in [1].

4 From CP-nets to strategic games

Consider now a CP-net with the set of variables $\{X_1, \ldots, X_n\}$ with the corresponding finite domains $\mathcal{D}(X_1), \ldots, \mathcal{D}(X_n)$. We write each preference statement for the variable X_i as $X_I = a_I : \succ_i$, where for the subsequence $I = i_1, \ldots, i_k$ of $1, \ldots, n$:

- $Pa(X_i) = \{X_{i_1}, \ldots, X_{i_k}\}$,
- $X_I = a_I$ is an abbreviation for $X_{i_1} = a_{i_1} \wedge \ldots \wedge X_{i_k} = a_{i_k}$,
- $\succ_i$ is a preference over $\mathcal{D}(X_i)$.

We also abbreviate $\mathcal{D}(X_{i_1}) \times \ldots \times \mathcal{D}(X_{i_k})$ to $\mathcal{D}(X_I)$.

By definition, the preference statements for a variable X_i are exactly all statements of the form $X_I = a_I : \succ(a_I)$, where a_I ranges over $\mathcal{D}(X_I)$ and $\succ(a_I)$ is a preference on $\mathcal{D}(X_i)$ that depends on a_I.

We now associate with each CP-net N a game $\mathcal{G}(N)$ with parametrized preferences as follows:

- each variable X_i corresponds to a player i,
- the strategies of player i are the elements of the domain $\mathcal{D}(X_i)$ of X_i.

To define the parametrized preferences, consider a player i. Suppose $Pa(X_i) = \{X_{i_1}, \ldots, X_{i_k}\}$ and let $I := i_1, \ldots, i_k$. So I is a subsequence of $1, \ldots, i-1, i+1, \ldots, n$. Given a joint strategy a_{-i} of the opponents of player i, we associate with it the preference relation $\succ(a_I)$ on $\mathcal{D}(X_i)$ where $X_I = a_I : \succ(a_I)$ is the unique preference statement for X_i determined by a_I.

In words, the preference of a player i over his strategies, given the strategies chosen by its opponents, say a_{-i}, coincides with the preference given by the CP-net over the domain of X_i given the assignment to his parents a_I which must coincide with the projection of a_{-i} over I.

This completes the definition of $\mathcal{G}(N)$.

As an example consider the first CP-net of Section 2. The corresponding game has four players A, B, C, D, each with two strategies indicated with f, $\bar{f}$ for player F. The preference of each player on his strategies will depend only on the strategies chosen by the players which correspond to his parents in the CP-net. Consider for example player B. His preference over his strategies b and $\bar{b}$, given the joint strategy of his opponents $s_{-B} = dac$, is $b \succ \bar{b}$. Notice that, for example, the same ordering holds for the opponents joint strategy $s_{-B} = \bar{d}a\bar{c}$, since the strategy chosen by the only player corresponding to his parent, A, has not changed.

We have then the following result.

Theorem 2. *An outcome of a CP-net N is optimal iff it is a Nash equilibrium of the game $\mathcal{G}(N)$.*

Proof. ($\Rightarrow$) Take an optimal outcome o of N. Consider a player i in the game $\mathcal{G}(N)$ and the corresponding variable X_i of N. Suppose $Pa(X_i) = \{X_{i_1}, \ldots, X_{i_k}\}$. Let $I := i_1, \ldots, i_k$, and let $X_I = o_I : \succ(o_I)$ be the corresponding preference statement for X_i. By definition there is no improving flip from o to another outcome, so o_i is the maximal element in the ordering $\succ(o_I)$.

By the construction of the game $\mathcal{G}(N)$, each outcome in N is a joint strategy in $\mathcal{G}(N)$. Also, two outcomes are one flip away iff the corresponding joint strategies differ only in a strategy of one player. Given the joint strategy o considered above, we thus have that, if we modify the strategy of player i, while leaving the strategies of the other players unchanged, this change is worsening in $\succ(o_{-i})$, since $\succ(o_{-i})$ coincides with $\succ(o_I)$. So by definition o is a Nash equilibrium of $\mathcal{G}(N)$.

($\Leftarrow$) Take a Nash equilibrium s of the game $\mathcal{G}(N)$. Consider a variable X_i of N. Suppose $Pa(X_i) = \{X_{i_1}, \ldots, X_{i_k}\}$. Let $I := i_1, \ldots, i_k$, and let $X_I = s_I : \succ(s_I)$ be the corresponding preference statement for X_i.

By definition for every strategy $s'_i \neq s_i$ of player i, we have $s_i \succ(s_{-i})\ s'_i$, so $s_i \succ(s_I)\ s'_i$ since $\succ(s_{-i})$ coincides with $\succ(s_I)$. So by definition s is an optimal outcome for N. $\square$

5 From strategic games to CP-nets

We now associate with each game G with parametrized preferences a CP-net $\mathcal{N}(G)$ as follows:

- each player i corresponds to a variable X_i,
- the domain $\mathcal{D}(X_i)$ of the variable X_i consists of the set of strategies of player i,
- we stipulate that $Pa(X_i) = \{X_1, X_{i-1}, \ldots, X_{i+1}, \ldots, X_n\}$, where n is the number of players in G.

Next, for each joint strategy s_{-i} of the opponents of player i we take the preference statement $X_{-i} = s_{-i} : \succ(s_{-i})$, where $\succ(s_{-i})$ is the preference relation on the set of strategies of player i associated with s_{-i}.

This completes the definition of $\mathcal{N}(G)$. As an example of this construction let us return to the Prisoner's dilemma game with parametrized preferences from Section 3. In the corresponding CP-net we have then two variables X_1 and X_2 corresponding to players 1 and 2, with the respective domains $\{C_1, N_1\}$ and $\{C_2, N_2\}$. To explain how each parametrized preference translates to a preference statement take for example $\succ(C_2) := N_1 \succ C_1$. It translates to $X_2 = C_2 : N_1 \succ C_1$.

We have now the following counterpart of Theorem 2.

Theorem 3. *A joint strategy is a Nash equilibrium of the game G iff it is an optimal outcome of the CP-net $\mathcal{N}(G)$.*

Proof. $(\Rightarrow)$ Assume G has a Nash equilibrium s. Thus, for every player i, joint strategy s_{-i}, and strategy $s'_i \neq s_i$ for player i, we have $s_i \succ(s_{-i})\ s'_i$. This means that, if we only change the strategy of any player i, this change is worsening for that player. In the CP-net $\mathcal{N}(G)$, s is an outcome, and the ordering in the conditional preference table of variable X_i coincides with $' \succ(s_{-i})$. Thus all the flips from s are worsening. Thus s is optimal for $\mathcal{N}(G)$.

$(\Leftarrow)$ Assume $\mathcal{N}(G)$ is eligible, and take an optimal outcome o of $\mathcal{N}(G)$. By definition of optimal outcome, for every other outcome o', $o' \not\succ o$, which means that there is no sequence of improving flips from o to any other outcome. Thus there is no improving flip from o to any other outcome. Therefore every flip modifying variable X_i in o is worsening in the preference statement for the variable X_i.

Given the construction from game G to CP-net $\mathcal{N}(G)$, an outcome in $\mathcal{N}(G)$ is a joint strategy in G. Also, two outcomes one flip away in $\mathcal{N}(G)$ are two joint strategies of G which differ only for the strategy of one player. Given the joint strategy o, we thus have that, if we modify the strategy of player i, while leaving the strategies of the other players unchanged, this change is worsening in $\succ_i(o_{-i})$, since the preference of variable X_i given o_{-i} coincides with $\succ_i(o_{-i})$.

$\square$

6 Reduced CP-nets

The disadvantage of the above construction of the CP-net $\mathcal{N}(G)$ from a game G is that it always produces a CP-net in which all sets of parent features are of size $n - 1$ where

n is the number of features of the CP-net. This can be rectified by reducing each set of parent features to a minimal one as follows.

Given a CP-net N, consider a variable X_i with the parents $Pa(X_i)$, and take a variable $Y \in Pa(X_i)$. Suppose that for all assignments a to $Pa(X) - \{Y\}$ and any two values $y_1, y_2 \in \mathcal{D}(Y)$, the orderings $\succ(a, y_1)$ and $\succ(a, y_2)$ on $\mathcal{D}(X_i)$ coincide.

We say then that Y is *redundant* in the set of parents of X_i. It is easy to see that by removing all redundant variables from the set of parents of X_i and by modifying the corresponding preference statements for X_i accordingly, the strict preorder $\succ$ over the outcomes of the CP-nets is not changed.

Given a CP-net, if for all its variable X_i the set $Pa(X_i)$ does not contain any redundant variable, we say that the CP-net is ***reduced***.

By iterating the above construction every CP-net can be transformed to a reduced CP-net. As an example consider a CP-net with three features, X, Y and Z, with the respective domains $\{a_1, a_2\}$, $\{b_1, b_2\}$ and $\{c_1, c_2\}$. Suppose now that $Pa(X) = Pa(Y) = \emptyset, Pa(Z) = \{X, Y\}$ and that

$\succ(a_1, b_1) = \succ(a_2, b_1),\ \succ(a_1, b_2) = \succ(a_2, b_2),$
$\succ(a_1, b_1) = \succ(a_1, b_2),\ \succ(a_2, b_1) = \succ(a_2, b_2).$

Then both X and Y are redundant, so we can reduce the CP-net by reducing $Pa(Z)$ to $\emptyset$. Z becomes an independent variable in the reduced CP-net with an ordering over its domain which coincides with the unique one given in the original CP-net in terms of the assignments to its parents.

In what follows for a CP-net N we denote by $r(N)$ the corresponding reduced CP-net. The following result summarizes the relevant properties of $r(N)$ and relates it to the constructions of $\mathcal{G}(N)$ and $\mathcal{N}(G)$.

Theorem 4.
(i) *Each CP-net N and its reduced form $N' = r(N)$ have the same ordering $\succ$ over the outcomes.*
(ii) *For each CP-net N and its reduced form $N' = r(N)$ we have $\mathcal{G}(N) = \mathcal{G}(N')$.*
(iii) *Each reduced CP-net N is a reduced CP-net corresponding to the game $\mathcal{G}(N)$. Formally: $N = r(\mathcal{N}(\mathcal{G}(N)))$.*

Proof.
(i) Consider N and $r(N)$, an outcome o and an improving flip in N from o to o' which modifies the value of variable X_i. Then this change is improving in the conditional preference table of X_i in N. Let us now consider the conditional preference table of X_i in $r(N)$. In this table there could be a subsequence of parents but with the same corresponding preference orderings. Thus the change is improving in this conditional preference table as well. Thus any chain of improving flips in N remain a chain of improving flips also in $r(N)$, and therefore the orderings over the outcomes of N and $r(N)$ are the same.

(ii) It follows directly from the construction of the game corresponding to a CP net, since the preference of player i over its strategies depends on the strategies of all the other opponents, even if variable X_i has just a few parents in N.

(iii) Given a reduced CP-net N, consider the CP-net $\mathcal{N}(\mathcal{G}(N))$. For each variable X_i, $Pa(X_i)$ in N is a subset of $Pa(X_i)$ in $\mathcal{N}(\mathcal{G}(N))$, which is the set of all variables except

X_i. However, by the construction of the game corresponding to a CP-net and of the CP-net corresponding to a game, in each conditional preference table, if the assignments to the common parents are the same, the preference orderings over X_i are the same.

Let us now reduce $\mathcal{N}(\mathcal{G}(N))$ to obtain $N' = r(\mathcal{N}(\mathcal{G}(N)))$. Then $Pa(X_i)$ in N' coincides with $Pa(X_i)$ in N. In fact, assume there is a parent of X_i in N which is not in N'. Since N is reduced, such a parent cannot be redundant in N. Thus the reduction, when applied to $\mathcal{N}(\mathcal{G}(N))$, cannot remove it since the orderings in the conditional preference tables of N and $\mathcal{N}(\mathcal{G}(N))$ are the same. On the other hand, assume there is a parent of X_i in N' which is not in N. Since N' is reduced, such a parent cannot be redundant in N'. Thus it is not redundant in $\mathcal{N}(\mathcal{G}(N))$ as well. By construction of $\mathcal{N}(\mathcal{G}(N))$, it cannot be redundant in N neither. $\square$

Part (i) states that the reduction procedure preserves the ordering over the outcomes. Part (ii) states that the construction of a game corresponding to a CP-net does not depend on the redundancy of the given CP-net. Finally, part (iii) states that the reduced CP-net N can be obtained 'back' from the game $\mathcal{G}(N)$.

7 Game theoretical techniques in CP-nets

Given the correspondence between CP-nets and games and its properties presented in the previous sections, we can now use them to transfer standard techniques of game theory, used to find Nash equilibria, to CP-nets to find their optimal outcomes, and vice-versa.

To be more specific, given a CP-net N, consider the game $\mathcal{G}(N)$. Let us now modify $\mathcal{G}(N)$ by an iterated elimination of never best responses, obtaining a game G'. By Theorems 1 and 2, an outcome of N is optimal iff it is a Nash equilibrium of G'. Now modify N by

- reducing the variable domains to the corresponding sets of strategies in G',
- removing all preference statements that refer to a removed element,

and call the resulting CP-net N'.

By Theorem 1, the Nash equilibria of G and G' coincide. Also, by Theorem 3, a joint strategy is a Nash equilibrium of G iff it is an optimal outcome of N, and a joint strategy is a Nash equilibrium of G' iff it is an optimal outcome of N'. Thus N and N' have the same set of optimal outcomes.

It is useful to note that the elimination of never best responses, and consequently also the elimination of strictly dominated elements, can be carried out directly on a CP-net by introducing the following notions. Consider a CP-net N.

- We say that an element d_i from the domain $\mathcal{D}(X_i)$ of the variable X_i is a ***best response*** to a preference statement

  $$X_I = a_I : \succ_i$$

  for X_i if $d_i \succeq_i d_i'$ for all $d_i' \in \mathcal{D}(X_i)$.

- We say that an element d_i from the domain of the variable X_i is a ***never a best response*** if it is not a best response to any preference statement for X_i.
- Given two elements d_i, d'_i from the domain $\mathcal{D}(X_i)$ of the variable X_i we say that d'_i is ***strictly dominated*** by d_i if for all preference statements $X_I = a_I : \succ_i$ for X_i we have

$$d_i \succ_i d'_i.$$

By a ***subnet*** of a CP-net N we mean a CP-net obtained from N by removing some elements from some variable domains followed by the removal of all preference statements that refer to a removed element.

Then we introduce the following relation between a CP-net N and its subnet N':

$$N \rightarrow_{NBR} N'$$

when $N \neq N'$ and for each variable X_i each removed element from the domain of X_i is never a best response in N, and introduce an analogous relation $N \rightarrow_S N'$ for the case of strictly dominated elements.

By the same argument as in the case of Theorem 1 we get the following result. Part (iii) can be established by repeating the argument of [1].

Theorem 5. *Suppose that $N \rightarrow^*_{NBR} N'$, i.e., the CP-net N' is obtained by an iterated elimination of never best responses from the CP-net N.*

(i) Then s is an optimal outcome of N iff it is an optimal outcome of N'.
(ii) If each variable in N' has a singleton domain, then the resulting outcome is a unique optimal outcome of N.
(iii) All iterated eliminations of never best responses from the CP-net N yield the same final outcome.

To illustrate the use of this theorem reconsider the first CP-net from Section 2, i.e., the one with the preference statements

$d : a \succ \overline{a},\ \overline{d} : a \succ \overline{a},$
$a : b \succ \overline{b},\ \overline{a} : \overline{b} \succ b,$
$b : c \succ \overline{c},\ \overline{b} : \overline{c} \succ c,$
$c : d \succ \overline{d},\ \overline{c} : \overline{d} \succ d.$

Denote it by N.

We can reason about it using the iterated elimination of strictly dominated strategies (which coincides here with the iterated elimination of never best responses, since each domain has exactly two elements).

We have the following chain of reductions:

$$N \rightarrow_S N_1 \rightarrow_S N_2 \rightarrow_S N_3 \rightarrow_S N_4,$$

where

- N_1 results from N by removing $\overline{a}$ (from the domain of A) and the preference statements $d : a \succ \overline{a},\ \overline{d} : a \succ \overline{a},\ \overline{a} : \overline{b} \succ b,$

– N_2 results from N_1 by removing $\overline{b}$ and the preference statements $a : b \succ \overline{b},\ \overline{b} : \overline{c} \succ c$,
– N_3 results from N_2 by removing $\overline{c}$ and the preference statements $b : c \succ \overline{c}\ \ \overline{c} : \overline{d} \succ d$,
– N_4 results from N_3 by removing $\overline{d}$ from the domain of D and the preference statement $c : d \succ \overline{d}$.

Indeed, in each step the removed element is strictly dominated in the considered CP-net. So using the iterated elimination of strictly dominated elements we reduced the original CP-net to one in which each variable has a singleton domain and consequently found a unique optimal outcome of the original CP-net N.

Finally, the following result shows that the introduced reduction relation on CP-nets is complete for acyclic CP-nets.

Theorem 6. *For each acyclic CP-net N a unique subnet N' with the singleton domains exists such that $N \rightarrow^*_{NBR} N'$.*

Proof. First note that if N is an acyclic CP-net with some non-singleton domain, then $N \rightarrow_{NBR} N'$ for some subnet N' of N. Indeed, suppose N is such a CP-net. Then a variable X exists with a non-singleton domain with no parent variable that has a non-singleton domain. So there exists in N exactly one preference statement for X, say $X_I = a_I : \succ_i$, where X_I is the sequence of parent variables of X. Reduce the domain of X to the maximal element in $\succ_i$. Then for the resulting subnet N' we have $N \rightarrow_{NBR} N'$.

Uniqueness of the outcome is a consequence of Theorem $5(iii)$. □

8 CP-net techniques in strategic games

The established relationship between CP-nets and strategic games with prioritized preferences allows us also to exploit the techniques developed for the CP-nets when reasoning about such games. In particular, to find the Nash equilibria of such a game we can proceed as follows. Take a game G. Apply to it an iterated elimination of never best responses. This yields a subgame G'. Now consider the corresponding CP-net $\mathcal{N}(G')$. Reduce it by eliminating redundant parents as described in Section 6, obtaining a CP-net $N' = r(\mathcal{N}(G))$. Next, apply to it an iterated elimination of never best responses. This yields a CP-net N''. By the theorems established in this paper the Nash equilibria of G coincide with the optimal outcomes of N''.

If each variable in N'' has a singleton domain, then we found a unique Nash equilibrium of G. Otherwise we can construct the optimality constraints *opt*(N'') discussed in Section 2 and use the fact that *opt*(N'') is consistent iff an optimal outcome of N'' exists. So we can now use hard constraint solving techniques to search for Nash equilibria of the original game G, by focusing on solving the constraints in *opt*(N'').

As an example consider the following strategic game G, in which there are two players, Player 1 with the strategies *Top* and *Bottom* and Player 2 with the strategies *Left* and *Right*.

	L	R
T	$2,2$	$1,1$
B	$1,1$	$2,2$

Note that in this game each strategy is a best response, so it cannot be reduced using the elimination of never best responses.

The corresponding CP-net, $\mathcal{N}(G)$, has two variables, each corresponding to a player: X_1 and X_2, with the domains $\mathcal{D}(X_1) = \{T, B\}$ and $\mathcal{D}(X_2) = \{L, R\}$. Moreover $Pa(X_1) = \{X_2\}$ and $Pa(X_2) = \{X_1\}$. The preference statements are the following:

$X_1 = T : L \succ R$,

$X_1 = B : R \succ L$,

$X_2 = L : T \succ B$,

$X_2 = R : B \succ T$.

Notice that $\mathcal{N}(G)$ is already reduced. In fact, X_1 is not redundant as a parent of X_2 and vice-versa.

The optimality constraints corresponding to this CP-net are:

$X_1 = T \rightarrow X_2 = L,\ X_1 = B \rightarrow X_2 = R,$

$X_2 = L \rightarrow X_1 = T,\ X_2 = R \rightarrow X_1 = B.$

These constraints have two solutions: $(X_1 = T, X_2 = L)$ and $(X_1 = B, X_2 = R)$, which are also the only two Nash equilibria of the initial game.

This approach can be used also to discover quickly that a CP-net has a unique Nash equilibrium, and to find it. Given a game G, we construct the reduced CP-net $N' = r(\mathcal{N}(G))$. If this CP-net is acyclic, we know that it has a unique optimal outcome which can be found in linear time. By the Theorems of the previous sections we also know that the Nash equilibria of game G coincide with the optimal outcome of this CP-net. This means that G has a unique Nash equilibrium that can be found in linear time by the usual CP-net techniques applied to N'.

Games G such that the CP-net $N' = r(\mathcal{N}(G))$ is acyclic are not uncommon. In fact, they naturally represent multi-agent scenarios where agents (that is, players of the game) can be partitioned into levels $1, 2, \ldots, n$, such that agents at level i can express their preferences (that is, payoff function) without looking at what players at higher levels do. Informally, agents at level i are more important than agents at level j is $j > i$. In particular, agents at level 1 can decide their preferences without looking at the behavior of any other agent.

9 Soft constraints and Nash equilibria

The direct correpondence between the optimal solutions of a CP-net and the Nash equilibria of the corresponding game cannot be easily found in other preference modelling formalisms. In this section we briefly give an informal account of the correspondence between the soft constraint framework and game theory.

Soft constraints, see [2], model problems with preferences via:

- a set of variables with finite domains,
- a set of constraints; each constraint involves a set of variables and assigns to each instantiation of its variables a preference level,

– the ***preference levels*** that are taken from a set A, over which a, possibly partial, ordering $\leq$ and a combination operator $\times$ are defined,
– the ***preference*** of a solution (that is, an instantiation of all the variables) is the combination of all the preferences given by the constraints to the subtuples of the solution.

A widely used instance of this formalism is the class of fuzzy constraints, see [8], where the set of preference levels A is $[0, 1]$, the ordering is the usual $\leq$ over the reals, and the combination operator is the min operation. In other words, in fuzzy constraints the optimal solutions are those that maximize the minimal preference.

A simple example of a fuzzy constraint problem is the following one:

– three variables: x, y, and z;
– two constraints: C_{xy} (over x and y) and C_{yz} (over y and z);
– domains for all variables: $\{a, b\}$;
– definition of the constraints: $C_{xy} := \{(aa, 0.4), (ba, 0.3), (ab, 0.1), (bb, 0.8)\}$ and $C_{yz} := \{(aa, 0.4), (ab, 0.3), (ba, 0.1), (bb, 0.8)\}$.

The only optimal solution of this problem is $x = y = z = b$, which has preference 0.8.

Given a fuzzy constraint problem, let us now consider a corresponding strategic game as follows:

– the players: one for each variable;
– the strategies of a player x: all values in the domain of x;
– the payoff function for player x: given s_{-x}, which is an instantiation of all the variables of the problem minus x, and given a value for x, this complete instantiation selects a tuple in each constraint. Take the minimum of the preferences over the tuples selected in the constraints where x appears.

Then one can prove that a Nash equlibrium of this game is not necessarily an optimal solution of the original fuzzy problem. In the example above, $x = y = z = a$ is a Nash equilibrium of the game but it is not a fuzzy optimum. On the other hand, $x = y = z = b$ is both an optimal solution and a Nash equilibrium.

In general, if the $\times$ operator is strictly monotonic, then an optimal solution is a Nash equilibrium. Since soft constraints always have optimal solutions, in this case the corresponding games always have Nash equilibria.

In fuzzy constraints, we cannot use this result because the min operator is not stricly monotonic. On the other hand, in weighted constraints (where $\times$ is + over the reals) we can use this result. The same holds for the probabilistic constraints ($\times$ is max, $+$ is multiplication and preferences drawn from $[0, 1]$).

10 Conclusions and future work

We showed that optimal outcomes in CP-nets are Nash equilibria in strategic games, and we exploited this to inherit useful techniques from strategic games to CP-nets and vice-versa.

In this paper we assume that payoff functions give a linear order over the strategies of a player. It could be useful to see whether our results can be generalized to games in which players' strategies can be incomparable or indifferent to each other, thus using partial orderings with ties. We are currently studying this scenario.

This paper is just a first step towards what we think is a fruitful cross-fertilization between preferences, constraint solving, and game theory. CP-nets appear to be very amenable for being studied using the existing game theoretical techniques, and also to provide strategic games with new and hopefully more efficient approaches to find Nash equilibria or to solve other game-theoretical tasks. We have seen that unfortunately this is not true in general for other preference modelling formalisms like soft constraints. We are therefore studying the conditions under which soft constraints or other formalisms can be related to game theory.

References

1. K. R. Apt. Rationalizability and order independence. In *Proc. 10th Conference on Theoretical Aspects of Reasoning about Knowledge (TARK)*. National University of Singapore, 2005. To appear.
2. S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based constraint solving and optimization. *Journal of the ACM*, 44(2):201–236, Mar 1997.
3. C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)*, 21:135–191, 2004.
4. C. Boutilier, R. I. Brafman, H. H. Hoos, and D. Poole. Reasoning with conditional ceteris paribus preference statements. In *UAI '99*, pages 71–80. Morgan Kaufmann, 1999.
5. R. Brafman and Y. Dimopoulos. Extended semantics and optimization algorithms for cp-networks. *Computational Intelligence*, 20, 2:218–245, 2004.
6. H. Fargier D. Dubois and H. Prade. The calculus of fuzzy restrictions as a basis for flexible constraint satisfaction. In *IEEE International Conference on Fuzzy Systems*, 1993.
7. C. Domshlak and R. I. Brafman. CP-nets: Reasoning and consistency testing. In *KR-02*, pages 121–132. Morgan Kaufmann, 2002.
8. H. Fargier D. Dubois and H. Prade. The calculus of fuzzy restrictions as a basis for flexible constraint satisfaction. In *IEEE International Conference on Fuzzy Systems*, 1993.
9. I. Gilboa, E. Kalai, and E. Zemel. On the order of eliminating dominated strategies. *Operation Research Letters*, 9:85–89, 1990.
10. R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard Univ Press, Cambridge, Massachusetts, 1991.
11. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, 1994.
12. S. Prestwich, F. Rossi, K. B. Venable, and T. Walsh. Constraint-based preferential optimization. In *Proc. of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*. Morgan Kaufmann, 2005.
13. M. Stegeman. Deleting strictly eliminating dominated strategies. Working Paper 1990/6, Department of Economics, University of North Carolina, 1990.