A FRAMEWORK FOR QUALITATIVE MULTI-CRITERIA PREFERENCES

Wietske Visser, Reyhan Aydoğan, Koen V. Hindriks and Catholijn M. Jonker
Man-Machine Interaction Group, Delft University of Technology, Delft, The Netherlands
{*wietske.visser, r.aydogan, k.v.hindriks, c.m.jonker*}*@tudelft.nl*

Keywords: Qualitative multi-criteria preferences.

Abstract: A key challenge in the representation of qualitative, multi-criteria preferences is to find a compact and expressive representation. Various frameworks have been introduced, each of which with its own distinguishing features. In this paper we introduce a new representation framework called qualitative preference systems (QPS), which combines priority, cardinality and conditional preferences. Moreover, the framework incorporates knowledge that serves two purposes: to impose (hard) constraints, but also to define new (abstract) concepts. In short, QPS offers a rich and practical representation for qualitative, multi-criteria preferences.

1 INTRODUCTION

A key challenge in the representation of qualitative, multi-criteria preferences is to find a *compact* and at the same time *expressive* representation. A framework for preference representation provides an adequate tool if it is sufficiently expressive to compactly represent a broad range of preference orderings. To this end, various frameworks have been introduced, each of which with its own distinguishing features. For example, in lexicographic approaches (e.g. (Andréka et al., 2002)) preference over outcomes is determined by combining multiple criteria according to *priority*. Goal-based approaches (e.g. (Brewka, 2004)) use *cardinality* and compare alternatives by the number of goals they satisfy. CP-nets (Boutilier et al., 2004) are well-known for their ability to represent *conditional preferences*. A CP-net is a qualitative graphical representation of preferences that reflects conditional preference statements under a *ceteris paribus* (all else being equal) interpretation. (Brafman and Domshlak, 2002) extend CP-nets to so-called TCP-nets which also allow for expressing relative importance of preference variables. In general, however, (T)CP-nets are not able to represent lexicographic orderings (Wilson, 2004).

In this paper we introduce a rich and practical new representation framework for qualitative multi-criteria preferences called *qualitative preference systems* (QPS). This framework enables preference representation by using *priority*, *cardinality* and *conditional preferences*. Moreover, the framework incorporates *knowledge* that serves two purposes: as usual, knowledge can be used to impose (hard) *constraints*, but also to define new (abstract) *concepts*. To illustrate, it can represent facts about the world (e.g. Barcelona is in Spain), the feasibility of options (e.g. hotel X is fully booked in July), and definitions (e.g. the cost of a holiday is the sum of the costs of the flight, hotel and food).

QPSs are based on the lexicographic rule studied in (Andréka et al., 2002). This rule is a fundamental part of the framework presented as it offers a principled tool for *combining basic preferences*. We believe this ability to combine preferences is essential for any practical approach to representing qualitative preferences. It is needed in particular for constructing *multi-criteria* preferences. It is not sufficient, however, since more expressivity is needed and useful in practice. Therefore, QPSs in addition provide a tool for representing knowledge, for abstraction, for counting, and provide a layered structure for representing preference orderings. We show that QPSs are able to represent various strategies for defining preference orderings, and are able to handle conditional preferences. To be precise, we show in Section 3 that Logical Preference Description language (LPD; (Brewka, 2004)) can be embedded into the QPS framework and that there is an order preserving embedding of CP-nets in the QPS framework. In addition we consider the key issue of compact preference representation and show that these embeddings provide a representation that is just as succinct as the LPD expressions and CP-nets.

2 QPS

The main aim of a QPS is to determine preferences between *outcomes* in a purely qualitative way. An outcome is an assignment of values to a set of relevant variables. Every variable has its own domain of possible values. Constraints on the assignments of values to variables are expressed in a knowledge base. Outcomes are defined as variable assignments that respect the constraints in the knowledge base.

The preferences between outcomes are based on multiple *criteria*. Every criterion can be seen as a *reason* for preference, or as a preference from one particular *perspective*. We distinguish between simple criteria that are based on a single variable and compound criteria that combine multiple criteria in order to determine an overall preference.

Definition 1 (Qualitative Preference System)**.** *A* qualitative preference system (QPS) *is a tuple* $\langle Var, Dom, K, \mathcal{C}\rangle$*. Var is a finite set of* variables*. Every variable* $X \in Var$ *has a domain* $Dom(X)$ *of possible values. K (a* knowledge base*) is a set of constraints on the assignments of values to the variables in Var. A* constraint *is an equation of the form* $X = Expr$ *where* $X \in Var$ *is a variable and Expr is an algebraic expression that maps to* $Dom(X)$*. An* outcome α *is an assignment of a value* $x \in Dom(X)$ *to every variable* $X \in Var$*, such that no constraints in K are violated.* α_X *denotes the value of variable X in outcome* α*.* $\mathcal{C}$ *is a finite rooted tree of criteria, where leaf nodes are simple criteria and other nodes are compound criteria. Child nodes of a compound criterion are called its subcriteria. Weak preference between outcomes by a criterion c is denoted by the relation* $\succeq_c$*.* $\succ_c$ *denotes the strict subrelation,* $\approx_c$ *the indifference subrelation.*

Example 1. When comparing holidays, some variables could be *Destination*, *NeverBeenThere* and *Cost*, with $Dom(Destination) = \{Barcelona, Rome, NewYork\}$, $Dom(NeverBeenThere) = \{\top, \bot\}$, $Dom(Cost) = \mathbb{Z}^+$. The definition of concepts (e.g. the cost of a holiday is the sum of the costs of the flight, hotel and food) can be straightforwardly represented with the following constraint: $Cost = FlightC + HotelC + FoodC$. Equational constraints are also sufficiently expressive to model different kinds of knowledge. For example, suppose I want to express that I have never been to Barcelona, i.e. in all outcomes where $Destination = Barcelona$, we should have $NeverBeenThere = \top$. To do this, we first introduce an auxiliary variable B with $Dom(B) = \{\top, \bot\}$. Then we add $B = (Destination = Barcelona)$ and $B = B \wedge NeverBeenThere$ to the constraint base K. This ensures that there are no outcomes where $Destination = Barcelona$ and $NeverBeenThere = \bot$.

2.1 Simple Criteria

A simple criterion specifies a preference ordering on the values of a single variable. Its preference between outcomes is based solely on the value of this variable in the considered outcomes.

Definition 2 (Simple Criterion)**.** *A* simple criterion *c is a tuple* $\langle X_c, \trianglerighteq_c\rangle$*, where* $X_c \in Var$ *is a variable, and* $\trianglerighteq_c$*, a preference relation on the possible values of* X_c*, is a preorder on* $Dom(X_c)$*.* $\triangleright_c$ *is the strict subrelation,* $\doteq_c$ *is the indifference subrelation. We call c a* Boolean simple criterion *if* X_c *is Boolean and* $\top \triangleright_c \bot$*. A simple criterion* $c = \langle X_c, \trianglerighteq_c\rangle$ weakly prefers *an outcome* α *over an outcome* β*, denoted* $\alpha \succeq_c \beta$*, iff* $\alpha_{X_c} \trianglerighteq_c \beta_{X_c}$*.*

Example 2. A criterion ‘economy’ can be defined as $\langle Cost, \trianglerighteq\rangle$ where for all $x, x' \in Dom(Cost)$, $x \trianglerighteq x'$ iff $x \leq x'$. An example of a Boolean criterion is ‘exploration’, defined as $\langle NeverBeenThere, \{(\top, \bot)\}\rangle$. The criterion ‘economy’ prefers any holiday with a lower cost over any holiday with a higher cost, irrespective of the values of other variables.

Observation 1. Let $c = \langle X_c, \trianglerighteq_c\rangle$ be a simple criterion. Then $\succeq_c$ is a preorder. If $\trianglerighteq_c$ is total, then so is $\succeq_c$.

2.2 Compound Criteria

Qualitative preference systems offer two ways in which to combine multiple criteria: lexicographic criteria and cardinality criteria. In a lexicographic criterion, preference is determined by the subcriteria with the highest priority; lower priority subcriteria only influence the preference if the higher priority subcriteria are indifferent. In a cardinality criterion, all subcriteria have the same priority, and preference is determined by a kind of voting mechanism that counts the number of subcriteria that support a certain preference and those that do not.

2.2.1 Lexicographic Criteria

A lexicographic criterion consists of a set of subcriteria and an associated priority order (a strict partial order, which means that no two subcriteria can have the same priority). It weakly prefers outcome α over outcome β if for every subcriterion, either this subcriterion weakly prefers α over β, or there is another subcriterion with a higher priority that strictly prefers α over β. This definition of preference by a lexicographic criterion is equivalent to the priority operator as defined by (Andréka et al., 2002). It generalizes the familiar rule used for alphabetic ordering of words, such that the priority can be any partial order and the combined preference relations can be any preorder.

Definition 3 (Lexicographic Criterion)**.** *A lexicographic criterion c is a tuple $\langle C_c, \rhd_c \rangle$, where C_c is a nonempty set of criteria (the* subcriteria *of c) and $\rhd_c$, a* priority relation *among subcriteria, is a strict partial order (a transitive and asymmetric relation) on C_c. A lexicographic criterion $c = \langle C_c, \rhd_c \rangle$* weakly prefers *an outcome α over an outcome β, denoted $\alpha \succeq_c \beta$, iff $\forall s \in C_c (\alpha \succeq_s \beta \vee \exists s' \in C_c (\alpha \succ_{s'} \beta \wedge s' \rhd_c s))$.*

Example 3. Consider a lexicographic criterion $c = \langle \{s_1, s_2\}, \{(s_1, s_2)\} \rangle$ where s_1 is the 'exploration' criterion and s_2 the 'economy' criterion from Example 2. Consider three outcomes such that $\alpha_{Destination} = Rome$, $\alpha_{Cost} = 500$, and $\alpha_{NeverBeenThere} = \bot$; $\beta_{Destination} = Barcelona$, $\beta_{Cost} = 350$, and $\beta_{NeverBeenThere} = \top$; and $\gamma_{Destination} = NewYork$, $\gamma_{Cost} = 700$, and $\gamma_{NeverBeenThere} = \top$. Then we have $\beta \succ_c \gamma \succ_c \alpha$. Note that even though α is cheaper than γ and hence preferred by criterion s_2, criterion c prefers γ to α because subcriterion s_1 has higher priority than s_2 and s_1 prefers γ to α.

Proposition 1. Let $c = \langle C_c, \rhd_c \rangle$ be a lexicographic criterion. If for all subcriteria $s \in C_c$, $\succeq_s$ is a preorder, then the relation $\succeq_c$ is also a preorder.

Proof. Preservation of reflexivity follows directly from the definition of $\succeq_c$ (if all subcriteria are reflexive, then for every outcome α: $\forall s \in C_c (\alpha \succeq_s \alpha)$ and hence $\alpha \succeq_c \alpha$). Preservation of transitivity has been proven by (Andréka et al., 2002). ☐

2.2.2 Cardinality Criteria

Like a lexicographic criterion, a cardinality criterion combines multiple criteria into one preference ordering. Unlike a lexicographic criterion, priority between subcriteria is not a strict partial order, but all subcriteria have the same priority. A cardinality criterion weakly prefers an outcome α over an outcome β if it has at least as many subcriteria that strictly prefer α over β as criteria that do not weakly prefer α over β.

Definition 4 (Cardinality Criterion)**.** *A cardinality criterion c is a tuple $\langle C_c \rangle$ where C_c is a nonempty set of criteria (the* subcriteria *of c). A cardinality criterion $c = \langle C_c \rangle$* weakly prefers *an outcome α over an outcome β, denoted $\alpha \succeq_c \beta$, iff $|\{s \in C_c \mid \alpha \succ_s \beta\}| \geq |\{s \in C_c \mid \alpha \not\succeq_s \beta\}|$.*

Example 4. Consider a cardinality criterion $c = \langle \{s_1, s_2\} \rangle$ where s_1 is the 'exploration' criterion and s_2 the 'economy' criterion from Example 2. For the three outcomes specified in Example 3, we have $\beta \succ_c \alpha \approx_c \gamma$.

Unfortunately, transitivity of $\succeq_c$ is not guaranteed for just any set of subcriteria. For example, consider three outcomes α, β, γ and three subcriteria s_1, s_2, s_3 such that $\alpha \succ_{s_1} \beta \succ_{s_1} \gamma$, $\beta \succ_{s_2} \gamma \succ_{s_2} \alpha$, and $\gamma \succ_{s_3} \alpha \succ_{s_3} \beta$. Then α would be strictly preferred over β, β strictly preferred over γ and γ strictly preferred over α, so the preference would not be transitive. However, there are some conditions under which transitivity *can* be guaranteed. E.g. if every subcriterion is a Boolean simple criterion $s = \langle X_s, \succeq_s \rangle$, they all induce a total preorder of preference that stratisfies the outcome space into two levels: the outcomes where $X_s = \top$ are more preferred and the outcomes where $X_s = \bot$ are less preferred. This also means that $\alpha \succ_s \beta$ iff $\alpha_{X_s} = \top$ and $\beta_{X_s} = \bot$; and $\alpha \not\succeq_s \beta$ iff $\alpha_{X_s} = \bot$ and $\beta_{X_s} = \top$. So in this case the definition preference by a priority class compares the number of 'goals' X_s that α satisfies to the number of goals that β satisfies, just as is done by e.g. the # strategy of (Brewka, 2004).

Proposition 2. Let $c = \langle C_c \rangle$ be a cardinality criterion such that for all $s \in C_c$, s is a Boolean simple criterion. Then $\succeq_c$ is a preorder.

Proof. Since all subcriteria of c are reflexive (Observation 1), for any outcome α both $|\{s \in C_c \mid \alpha \succ_s \alpha\}|$ and $|\{s \in C_c \mid \alpha \not\succeq_s \alpha\}|$ are 0, so $\alpha \succeq_c \alpha$, hence $\succeq_c$ is reflexive. Since all subcriteria are Boolean simple criteria, $\alpha \succeq_c \beta$ iff $|\{s = \langle X_s, \succeq_s \rangle \in C_c \mid \alpha_{X_s} = \top\}| \geq |\{s = \langle X_s, \succeq_s \rangle \in C_c \mid \beta_{X_s} = \top\}|$. This is just a comparison between two integers, and hence is transitive. ☐

(Andréka et al., 2002) showed that the only operator to combine *any arbitrary* preference relations that satisfies the desired properties IBUT (independence of irrelevant alternatives, based on preferences only, unanimity with abstentions, and preservation of transitivity) is the priority operator, which assumes that priority is a partial order. We observe here that if only *Boolean* preference relations (such as those resulting from Boolean simple criteria) are combined, the cardinality-based rule, in which all combined relations have equal priority, can also be applied. Requiring antisymmetry in this case would unneccessarily restrict the expressivity.

2.3 Conditional Preferences

A QPS can be used to express *conditional preferences*, i.e. preferences between values of one variable that depend on the values of other variables.

Example 5. If Anne goes on a holiday to Barcelona (b), she would like to go together with her friend Juan (j), but if she goes to Rome (r), she prefers to go with Mario (m). To express this conditional preference in a QPS, we use an auxiliary variable L, whose domain consists of all combinations of the variables D (destination) and C (company), i.e. $Dom(L) =$

$\{(b,j),(b,m),(r,j),(r,m)\}$. To keep the outcomes consistent, the constraint $L=(D,C)$ is added to the knowledge base. Finally, the following simple criterion expresses the conditional preference: $c=\langle L,\geq_c\rangle$ where $\geq_c=\{((b,j),(b,m)),((r,m),(r,j))\}$.

Instead of representing this kind of preference as conditional preferences on the values of variables, it would be more natural to model the underlying *reason* for the conditional preference, as was argued in (Visser et al., 2011). This is possible in a QPS, but outside the scope of this paper.

3 COMPARISON WITH OTHER FRAMEWORKS

3.1 LPD

(Brewka, 2004) presents a rank-based description language for qualitative preferences called *logical preference description* language (LPD). The basic expressions of LPD are called *basic preference descriptions* which are pairs $\langle s,R\rangle$ with s one of the *strategy identifiers* $\top$, κ, $\subseteq$, # and R a *ranked knowledge base* (RKB). An RKB is a set F of propositional formulas together with a *total* preorder $\geq$ on F, representing the relative importance of the formulas. Alternatively, an RKB can be represented as a set of ranked formulas (f,i) where f is a propositional formula and i, the rank of f, is a non-negative integer such that $f_1 \geq f_2$ iff $rank(f_1) \geq rank(f_2)$.

The four strategy identifiers refer to different strategies to obtain preferences over outcomes from an RKB. Outcomes in this context are propositional models, i.e. the variables used are Boolean. $\subseteq$ prefers α over β if there is a rank where α satisfies a superset of the formulas that β satisfies, and α and β satisfy the same more important formulas. # prefers α over β if there is a rank where α satisfies more formulas than β, and for all more important ranks, α and β satisfy the same number of formulas. Since (Brewka, 2004) shows that basic preference descriptions $\langle\top,R\rangle$ and $\langle\kappa,R\rangle$ can be transformed into equivalent basic preference descriptions of the form $\langle\subseteq,R'\rangle$, we do not discuss these strategies here.

Theorem 1. There is a QPS $\langle Var,Dom,K,\mathcal{C}\rangle$ with a criterion $c\in\mathcal{C}$ that corresponds to a basic preference description $\langle s,R\rangle$ for $s=\#$ or $s=\subseteq$ such that $\alpha\geq_s^R\beta$ iff $\alpha\geq_c\beta$ for arbitrary outcomes α,β.

Proof. A basic preference description $\langle s,R\rangle$ can be translated into a QPS $\langle Var,Dom,K,\mathcal{C}\rangle$. Let $R=\langle F,\geq\rangle$ be an RKB. The propositional variables used in F are collected in *Var*; moreover, for each formula $f\in F$ a new variable X_f is added to *Var* and $X_f=f$ is added to the knowledge base K. Clearly, $Dom(X)=\{\top,\bot\}$ for all $X\in Var$. For every formula $f\in F$, a Boolean simple criterion on the associated variable is defined: $c_f=\langle X_f,\{(\top,\bot)\}\rangle$. If $s=\subseteq$, preference of $\langle s,R\rangle$ is captured by a lexicographic criterion $c=\langle C_c,\rhd_c\rangle$ such that $C_c=\{c_f \mid f\in F\}$ and $c_f\rhd c_{f'}$ iff $f>f'$. Note that Boolean criteria that correspond to formulas with the same rank are incomparable according to the criterion c. This ensures that an outcome α can only be preferred to an outcome β on some rank, if there is no criterion that strictly prefers β over α, i.e. there is no formula that β satisfies but α does not. This means that α satisfies a superset of the formulas that β satisfies, which is the definition of preference by the $\subseteq$ strategy. If $s=\#$, for every rank i in R, a cardinality criterion is defined with as subcriteria all simple criteria associated to a formula of that rank: $c_i=\langle\{c_f \mid (f,i)\in R\}\rangle$. The preference of $\langle s,R\rangle$ is captured by a lexicographic criterion $c=\langle C_c,\rhd_c\rangle$ such that $C_c=\{c_i \mid (f',i)\in R\}$ and $c_i\rhd c_{i'}$ iff $i>i'$. This way, a subcriterion of c corresponds with a rank in the RKB R. Now note that how preferences are induced by c and its subcriteria corresponds with how the strategy # induces preferences over outcomes. □

Note that it follows from Theorem 1 that the QPS corresponding to a basic preference description is *just as succinct* as this description. That is, the size of the QPS is comparable to that of the LPD description (the size differs at most by a constant factor).

In LPD, complex preference descriptions can be built from basic ones with the connectives $\wedge$, $\vee$, $>$ and $-$. The meaning of a complex description is defined in terms of the orderings $\geq_1$ and $\geq_2$ induced by basic preference descriptions d_1 and d_2. The order denoted by $d_1\wedge d_2$ is the intersection $\geq_1\cap\geq_2$ (Pareto ordering), $d_1\vee d_2$ denotes the *transitive closure* of $\geq_1\cup\geq_2$, $-d_1$ denotes the reversed ordering $\geq_1$, and $d_1>d_2$ denotes the lexicographic ordering of $\geq_1$ and $\geq_2$ where α is strictly preferred to β if $\alpha>_1\beta$ or $\alpha\geq_1\beta$ and $\alpha>_2\beta$.

We show that complex descriptions can also be translated into a QPS that is just as succinct. We first introduce the notion of a reversed criterion that induces the reverse of the ordering induced by the original criterion. This can be achieved by reversing the value preferences of all the simple criteria in a QPS.

Definition 5 (Reverse of a Criterion)**.** *The reverse of a simple criterion $c=\langle X_c,\geq_c\rangle$ is $c^-=\langle X_c,\geq_{c^-}\rangle$ with $\beta\geq_{c^-}\alpha$ iff $\alpha\geq_c\beta$. The reverse of a cardinality criterion $c=\langle C_c\rangle$ is $c^-=\langle C_{c^-}\rangle$ where $C_{c^-}=\{s_i^- \mid s_i\in C_c\}$. The reverse of a lexicographic criterion $c=\langle C_c,\rhd_c\rangle$ is $c^-=\langle C_{c^-},\rhd_{c^-}\rangle$ where $C_{c^-}=\{s_i^- \mid s_i\in C_c\}$ and $s_1^-\rhd_{c^-}s_2^-$ whenever $s_1\rhd_c s_2$.*

Theorem 2. Let c_1 and c_2 be any two criteria. The lexicographic criterion $c_{1\wedge 2} = \langle\{c_1, c_2\}, \varnothing\rangle$ induces the order $\succeq_{c_1} \cap \succeq_{c_2}$. The lexicographic criterion $c_{1>2} = \langle\{c_1, c_2\}, \{(c_1, c_2)\}\rangle$ induces the order $\alpha \succeq_{c_{1>2}} \beta$ iff $\alpha \succ_{c_1} \beta$ or $\alpha \succeq_{c_1} \beta$ and $\alpha \succ_{c_2} \beta$. The criterion c_1^- induces the order $\beta \succeq_{c_1^-} \alpha$ iff $\alpha \succeq_{c_1} \beta$.

Theorem 2 clearly shows the expressive power of QPSs. It is very easy to represent specific operations for combining preference orderings by means of QPSs such as creating a Pareto order ($\wedge$ operator), refining a preference ordering by means of a second one ($>$ operator), and reversing an ordering ($-$ operator). Moreover, Theorem 2 shows this can be done just as succinctly with QPSs as with RKBs; i.e. the size needed differs at most with a constant factor.

The only operator that cannot be represented in a QPS is disjunction ($\vee$). However, it has been argued convincingly by (Andréka et al., 2002) that this is not a natural operator, since it does not satisfy the desired properties 'indifference to irrelevant alternatives' and 'unanimity with abstentions'. Indifference to irrelevant alternatives means that two outcomes can be compared solely on their own merits; the presence or absence of other possible outcomes does not influence the preference. The disjunction operator is not indifferent to irrelevant alternatives since it considers the transitive closure of the union of preference relations. Unanimity with abstentions means that if all combined preference relations prefer outcome α over outcome β, except possibly some that are indifferent, then the overall preference relation also prefers α over β. The disjunction operator would be indifferent as soon as one of the combined relations is indifferent, even if all others strictly prefer α over β.

We have shown that LPD descriptions (except disjunction) can be represented by QPSs just as succinctly. QPSs are more general, however, than LPD which is based on ranked knowledge bases. Whereas RKBs require a total preorder on formulas, QPSs allow incomparable priority between subcriteria. QPSs are not restricted to Boolean variables as LPD is. Apart from propositional formulas, QPSs support the use of equational constraints over arbitrary domains. In particular, QPSs provide a definitional mechanism in order to introduce new concepts (abstract variables) and it is possible to define preferences over such abstract variables. The knowledge that can be captured in a QPS therefore is more general.

3.2 CP-nets

(Boutilier et al., 2004) introduce CP-nets: qualitative graphical representations of preferences that reflect (conditional) preference statements under a *ceteris paribus* (all else being equal) interpretation.

Definition 6 (CP-net)**.** *(Boutilier et al., 2004) A* CP-net *N over variables $\boldsymbol{V} = \{X_1, \ldots, X_n\}$ is a directed graph G over $X_1, \ldots, X_n$ whose nodes are annotated with conditional preference tables $CPT(X_i)$ for each $X_i \in \boldsymbol{V}$. Each conditional preference table $CPT(X_i)$ associates a total order $\succeq_{\boldsymbol{u}}^i$ with each instantiation $\boldsymbol{u}$ of X_i's parents $Pa(X_i) = \boldsymbol{U}$. A preference ranking $\succeq$ (a total preorder over the set of outcomes)* satisfies *a CP-net N iff for each variable X_i and for each assignment $\boldsymbol{u}$ to the variables in $\boldsymbol{U}$, $yx\boldsymbol{u} \succeq \boldsymbol{y}x'\boldsymbol{u}$ whenever $x \succeq_{\boldsymbol{u}}^i x'$ – for all assignments $\boldsymbol{y}$ to the set of variables $\boldsymbol{Y} = \boldsymbol{V} - (\boldsymbol{U} \cup \{X_i\})$ and all $x, x' \in Dom(X_i)$. N* entails *$\alpha \succ \beta$, written $N \models \alpha \succ \beta$, iff $\alpha \succ \beta$ holds in every preference ordering that satisfies N. (Boutilier et al., 2004) show that $N \models \alpha \succ \beta$ iff there is a sequence of improving flips from β to α. An improving flip of outcome $\boldsymbol{u}x\boldsymbol{y}$ with respect to variable X_i is any outcome $\boldsymbol{u}x'\boldsymbol{y}$ such that $x' \succ_{\boldsymbol{u}}^i x$.*

Theorem 3. There is a QPS $\langle Var, Dom, K, \mathcal{C}\rangle$ with a criterion $c \in \mathcal{C}$ that corresponds to an acyclic CP-net N over variables $\mathbf{V} = \{X_1, \ldots, X_n\}$ such that if $N \models \alpha \succ \beta$ then $\alpha \succ_c \beta$ for arbitrary outcomes α, β.

Proof. The CP-net N can be translated to the QPS S as follows. All variables in the CP-net are also variables in the QPS: $\mathbf{V} \subseteq Var$. For every variable $X_i \in \mathbf{V}$, a simple criterion c_i is specified. If X_i is conditionally independent, $c_i = \langle X_i, \succeq_{c_i}\rangle$ such that $x \succeq_{c_i} x'$ iff $x \succ^i x'$. If X_i is conditionally dependent, an auxiliary variable X_i' is added to Var such that $Dom(X_i') = \prod\{Dom(X) \mid X \in X_i \cup Pa(X_i)\}$. The constraint $X_i' = \prod(X_i \cup Pa(X_i))$ is added to K. $c_i = \langle X_i', \succeq_{c_i}\rangle$ such that $x\mathbf{u} \succeq_{c_i} x'\mathbf{u}$ iff $x \succ_{\mathbf{u}}^i x'$. Finally, a lexicographic criterion $c = \langle C_c, \rhd_c\rangle$ is defined such that for every simple criterion c_i thus generated from the CP-net, $c_i \in C_c$, and $\rhd_c$ is the transitive closure of $\rhd_c'$, where $c_i \rhd_c' c_j$ iff $X_i \in Pa(X_j)$ (note that since N is acyclic, $\rhd_c$ is asymmetric).

Suppose that $N \models \alpha \succ \beta$. This means that there is a sequence of improving flips from β to α. First consider the case where this sequence has length 1, i.e. there is a single improving flip w.r.t. some variable X_i from β to α. Since the preference by a simple criterion is taken from the corresponding CPT, $\alpha \succ_{c_i} \beta$. If X_i is not a parent of any variable, all other simple criteria are indifferent between α and β (since they do not involve X_i), so $\alpha \succ_c \beta$. If X_i is a parent of another variable X_j, flipping its value influences the value of the auxiliary variable X_j'. However, c_i has higher priority than c_j, so again we have $\alpha \succ_c \beta$. Since $\succeq_c$ is transitive, we also have $\alpha \succ_c \beta$ if the sequence of improving flips from β to α is longer than 1. □

Note that it follows from Theorem 3 that the QPS corresponding to an acyclic CP-net is *just as succinct* as this description. That is, the size of the QPS is comparable to that of the CP-net (the size differs at most by a constant factor).

There are some things that CP-nets cannot express, but a QPS can. Most importantly, we are able to express abstract preferences based on auxiliary variables whose values are constrained by the knowledge base. Consider the well-known example from game theory called the 'battle of the sexes': a husband and wife have to decide whether to go to the theater or to a football match. The wife prefers the theater and the husband prefers football, but both would rather go together than go to different places. If we let A (resp. B) stand for 'the wife (resp. the husband) goes to the theater' and $\neg A$ (resp. $\neg B$) for 'the wife (resp. the husband) goes to the football match', then the ordering $AB \succ \neg A \neg B \succ A \neg B \succ \neg AB$ represents the wife's preferences. A CP-net cannot express this ordering, since there is no improving flip between $\neg A \neg B$ and AB. In a QPS, this preference can be easily expressed by introducing an auxiliary variable T ('together'), whose values are constrained by $T = A \leftrightarrow B$. A lexicographic criterion with two Boolean simple subcriteria, based on T and A respectively, where the one based on T has higher priority, induces the desired preference ordering $TAB \succ T \neg A \neg B \succ \neg TA \neg B \succ \neg T \neg AB$.

Second, we add priority between criteria, which allows us to express that a good value for one variable is more important than a good value for another variable. TCP-nets (Brafman and Domshlak, 2002) are an extension to CP-nets in which some priority between variables is taken into account, but this is not strong enough to represent lexicographic preferences (Wilson, 2004). (Wilson, 2004)'s own approach can handle such preferences, but does not allow to use auxiliary variables and knowledge as described above.

Third, in a CP-net, every variable occurs exactly once. In a QPS, some variables may not occur in any criterion, and some variables may occur in multiple criteria, e.g. if the preference on its values is different from different perspectives, or if the preferences of multiple people are combined.

4 CONCLUSIONS

We introduced Qualitative Preference Systems, a new framework for representing multi-criteria preferences. QPSs combine different features for compactly expressing preferences. These features include the well-known lexicographic rule which combines basic preferences over variables, and a cardinality-based rule which counts criteria that are satisfied. In addition, QPSs provide a tool for expressing feasibility constraints as well as abstractions (concept definitions). Finally, such systems support a layered structure for representing preference orderings.

This combination of features provides a very expressive preference representation framework which at the same time allows for a compact representation of preference orderings. We have shown that the Logical Preference Descriptions introduced in (Brewka, 2004) can be embedded in the QPS framework, with the exception of the disjunction operator which is not very natural. The 'logical' operators of (Brewka, 2004) translate to structural features of QPSs. We have also shown that QPSs are able to express conditional preferences by providing an order preserving embedding of acyclic CP-nets into QPSs. Last but not least, these embeddings are size preserving, i.e. the resulting QPSs provide a representation that is as succinct as the LPD or CP-net representation. This fact indicates that various problems such as dominance testing for QPSs have an associated computational complexity that is at most as difficult as these alternative frameworks for preference representation.

ACKNOWLEDGEMENTS

This research is supported by the Dutch Technology Foundation STW, applied science division of NWO and the Technology Program of the Ministry of Economic Affairs. It is part of the Pocket Negotiator project with grant number VICI-project 08075.

REFERENCES

Andréka, H., Ryan, M., and Schobbens, P.-Y. (2002). Operators and laws for combining preference relations. *J. Logic Comput.*, 12(1):13–53.

Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., and Poole, D. (2004). CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res.*, 21:135–191.

Brafman, R. I. and Domshlak, C. (2002). Introducing variable importance tradeoffs into CP-nets. In *Proc. UAI*, pages 69–76.

Brewka, G. (2004). A rank based description language for qualitative preferences. In *Proc. ECAI*, pages 303–307.

Visser, W., Hindriks, K. V., and Jonker, C. M. (2011). Interest-based preference reasoning. In *Proc. ICAART*, pages 79–88.

Wilson, N. (2004). Extending CP-nets with stronger conditional preference statements. In *Proc. AAAI*, pages 735–741.