

Combining binary constraint networks in qualitative reasoning

Jason Jingshi Li^{1,2} and Tomasz Kowalski¹ and Jochen Renz^{1,2} and Sanjiang Li³

Abstract. Constraint networks in qualitative spatial and temporal reasoning are always complete graphs. When one adds an extra element to a given network, previously unknown constraints are derived by intersections and compositions of other constraints, and this may introduce inconsistency to the overall network. Likewise, when combining two consistent networks that share a common part, the combined network may become inconsistent.

In this paper, we analyse the problem of combining these binary constraint networks and develop certain conditions to ensure combining two networks will never introduce an inconsistency for a given spatial or temporal calculus. This enables us to maintain a consistent world-view while acquiring new information in relation with some part of it. In addition, our results enable us to prove other important properties of qualitative spatial and temporal calculi in areas such as representability and complexity.

1 INTRODUCTION

An important ability of intelligent systems is to handle spatial and temporal information. Qualitative calculi such as the Region Connection Calculus (RCC8) [10] or Allen's Interval Algebra (IA) [1] intend to capture such information by representing relationships between entities in space and time. Such calculi have different advantages compared to quantitative spatial and temporal representations such as coordinate systems. They are closer to everyday human cognition, deal well with incomplete knowledge, and can be computationally more efficient than, say, the full machinery of metric spaces.

Defining a qualitative calculus is very intuitive. What is required is a domain of spatial or temporal entities, a set of jointly exhaustive and pairwise disjoint (JEPD) relations between the entities of the domain, and (weak) composition between the relations. These properties are essential for enabling constraint-based reasoning techniques for qualitative calculi [13]. However, not all qualitative calculi that can be defined in this way are equally well suited for representing and reasoning about spatial and temporal information.

Consider two consistent sets Θ_1, Θ_2 of spatial or temporal information. It is clear that if both sets refer to different entities, then combining the two sets will also lead to a consistent set as there are no potentially conflicting constraints. If the two sets contain information about the same entities, then it is clear that combining the two sets might lead to inconsistencies, as your favourite crime story will amply demonstrate. Here we are interested in a particular kind of com-

bination of sets, namely, combining sets that share only a very small number of entities and where the relationships between the shared entities are identical in both sets. Assume, for example, that Θ_1 and Θ_2 contain consistent information about the spatial relationships of entities in two adjacent rooms, Θ_1 for room 1 and Θ_2 for room 2. Assume further that the two rooms are connected by n closed doors such that the relationships between the n doors are exactly the same in Θ_1 and Θ_2 , and the doors are the only entities contained in both sets. Without considering any additional information (e.g. that there is only one computer in total, but both room 1 and room 2 contain the computer according to Θ_1 and Θ_2), it is common sense that combining both sets Θ_1 and Θ_2 to $\Theta = \Theta_1 \cup \Theta_2$ cannot lead to an inconsistency. However, as several examples in the literature show [6], there are qualitative calculi where this property is not satisfied and where inconsistencies are introduced when combining two sets that share a small number of entities with identical relations. Such calculi are counterintuitive and it is questionable whether they should be used for spatial or temporal representation and reasoning at all, as they introduce inconsistencies where there shouldn't be any.

Apart from this problem, there are some practically very important advantages of using a qualitative calculus that allows the consistent combination of two consistent sets of information: (1) It opens up the possibility to use divide-and-conquer techniques and to split a large set of qualitative constraints into smaller sets that can be processed independently. This is an essential requirement for hierarchical reasoning and may also speed up reasoning. (2) It becomes possible to ignore or filter additional information if it is clear that it won't affect the information important to us.

Unfortunately, there is currently no general way of determining for which qualitative calculi consistent sets can be consistently combined and for which calculi unnecessary inconsistencies are introduced. Some initial results were obtained by Li and Wang [6], where a special case of this problem called *one-shot extensibility* was analysed. Li and Wang considered the case of consistently extending a consistent atomic set of RCC8 constraints by one additional entity, and showed manually by an extensive case analysis that this is always possible for RCC8. Li and Wang showed that one-shot extensibility is also an essential requirement for other important computational properties of a qualitative calculus.

In this paper we analyse combinations where two sets share at most two entities and identify a method for automatically testing if this is always possible for a given qualitative calculus. This case is particularly important for different reasons: (1) It provides a purely algebraic and very general proof for one-shot extensibility [6]. (2) It (partially) solves some fundamental questions related to algebraic closure, consistency and (weak) composition, and (3) it provides a purely symbolic test for when a relation algebra is representable.

¹ RSISE, The Australian National University, Canberra ACT 0200, Australia, email: jason.li|tomasz.kowalski|jochen.renz@anu.edu.au

² NICTA Canberra Research Laboratory, Canberra ACT 2601, Australia

³ State Key Laboratory of Intelligent Technology and Systems, Department of Computer Science and Technology, Tsinghua University, Beijing 100084, P.R. China, email: lisanjiang@tsinghua.edu.cn

2 PRELIMINARIES

A qualitative calculus such as RCC8 or the Interval Algebra defines relationships over a given set of spatial or temporal entities, the domain \mathcal{D} . The *basic relations* \mathcal{B} form a partition of $\mathcal{D} \times \mathcal{D}$ which is jointly exhaustive and pairwise disjoint, i.e., between any two elements of the domain exactly one basic relation holds [7]. RCC8 for example uses a topological space of extended regions as the domain and defines eight basic relations DC, EC, PO, EQ, TPP, NTPP, TPPI, NTPPi which are verbalised as disconnected, externally connected, partial overlap, equal, tangential proper part, non-tangential proper part and the converse relations of the latter two [10]. In this paper we intensively use the precise mathematical definitions of relations, algebras and different algebraic operators which we summarize in the following. A more detailed overview can be found in [3, 7, 13].

A *nonassociative relation algebra* (NA) is an algebraic structure $\mathbf{A} = (A, \wedge, \vee, ;, -, \smile, 1', 0, 1)$, such that

- $(A, \wedge, \vee, -, 0, 1)$ is a Boolean algebra
- $(A, ;, \smile, 1')$ is an involutive groupoid with unit, that is, a groupoid satisfying the following equations

$$(a) x ; 1' = 1' ; x \quad (b) x^{\smile} = x \quad (c) (x ; y)^{\smile} = y^{\smile} ; x^{\smile}$$

- the operations $;$ and \smile are normal operators, that is, they satisfy the following equations
 - $x ; 0 = 0 = 0 ; x$
 - $0^{\smile} = 0$
 - $x ; (y \vee z) = (x ; y) \vee (x ; z)$
 - $(x \vee y) ; z = (x ; z) \vee (y ; z)$
 - $(x \vee y)^{\smile} = x^{\smile} \vee y^{\smile}$
- the following equivalences hold

$$(x ; y) \wedge z = 0 \text{ iff } (z ; y^{\smile}) \wedge x = 0 \text{ iff } (x^{\smile} ; z) \wedge y = 0$$

These are called *Peircean laws* or *triangle identities*.

A nonassociative relation algebra is a *relation algebra* (RA) if the multiplication operation $(;)$ is associative. For more on relation algebras and nonassociative relation algebras see [4] and [8].

Let \mathbf{A} be a NA. For any set U , called a *domain*, let $\mathcal{R}(U)$ be the algebra $(\wp(U \times U); \cup, \cap, \circ, -, ^{-1}, \Delta, \emptyset, U \times U)$, where the operations are union, intersection, composition, complement, converse, the identity relation, the empty relation and the universal relation (all with their standard set-theoretical meaning). Notice that since \circ is associative, $\mathcal{R}(U)$ is a RA. We say that \mathbf{A} is *weakly represented over* U if there is a map $\mu: A \rightarrow \wp(U \times U)$ such that μ commutes with all operations except $;$ for which we require only

$$\mu(a ; b) \supseteq \mu(a) \circ \mu(b)$$

This property of weak representation gives rise to a notion of *weak composition* of relations, namely, for $R, S \in \mu[A]$, we define $R \diamond S$ to be the smallest relation $Q \in \mu[A]$ containing $R \circ S$. Every NA has a weak representation, for example a trivial one, with $U = \emptyset$. Of course, interesting weak representations are nontrivial, and typically injective. A weak representation is a *representation* if μ is injective and the inclusion above is in fact equality, that is, if μ is an embedding of relation algebras. In such a case weak composition equals composition [12], and that is expressed by saying that weak composition is *extensional*. Not every NA, indeed not every RA is *representable*.

Although weak representations are not as interesting as representations, curiously, it is the former that gave rise to a notion of *qualitative calculus*, which is a triple (\mathbf{A}, U, μ) where \mathbf{A} is a NA, U is a set and $\mu: A \rightarrow \wp(U \times U)$ is a weak representation of \mathbf{A} over U . Since (\mathbf{A}, U, μ) is notationally cumbersome, we will later write \mathbf{A} for both a NA and a corresponding calculus (\mathbf{A}, U, μ) , if U and μ are clear from context or their precise form is not important. A calculus (\mathbf{A}, U, μ) is *extensional* if μ is a representation of \mathbf{A} . Notice that if (\mathbf{A}, U, μ) is extensional, then \mathbf{A} is a RA, indeed a representable one. The converse need not hold, as the example of RCC8 demonstrates.

All NAs considered in this paper are assumed to be finite (hence atomic) and such that $1'$ is an atom. These are severe restrictions on the class of NAs, but natural from a qualitative calculi point of view.

A network N over a NA \mathbf{A} is a pair (V, ℓ) where V is a set of vertices (nodes) and $\ell: V^2 \rightarrow \mathbf{A}$ is any function. Thus, a network is a complete directed graph labelled by elements of \mathbf{A} . Abusing notation a little we will often write N for the set of vertices of N , if it does not cause confusion. Where double precision is important, we will write V_N and ℓ_N for the set of vertices of N and its labelling function, respectively. For convenience we assume that the set V of nodes is always a set of natural numbers. We will also frequently refer to the label on the edge (i, j) as R_{ij} . A network M is a *subnetwork* of N , if all nodes and labels of M are also nodes and labels of N . We will write $M \leq N$ in such case. A network M is a *refinement* of N if $V_M = V_N$ and $\ell_M(i, j) \leq \ell_N(i, j)$, for any $i, j \in V$ (where \leq is the natural ordering among the labels as elements of \mathbf{A}). A network is *atomic* if all the labels are basic relations (atoms) of \mathbf{A} . To indicate atomicity we will sometimes use lower case labels r_{ij} . A network N is *algebraically-closed* (a-closed) if the following hold

1. R_{ii} is the equality relation (identity element of \mathbf{A})
2. $R_{ij} \diamond R_{jk} \geq R_{ik}$ for any $i, j, k \in N$

Networks may be viewed as approximations to (weak) representations, indeed, if μ is a weak representation of \mathbf{A} over a domain U , then $\mu[A]$ is an a-closed network over \mathbf{A} . An arbitrary network N over \mathbf{A} is *consistent* with respect to a weak representation μ , if N is a subnetwork of $\mu[A]$.

This paper is mostly concerned with combining a-closed networks without introducing inconsistencies. Let N_0, N_1, N_2 be a-closed networks, such that $N_0 \leq N_1$ and $N_0 \leq N_2$. The triple (N_0, N_1, N_2) is called a *V-formation*. A V-formation (N_0, N_1, N_2) can be *amalgamated* if there is an a-closed network M such that $N_1 \leq M$ and $N_2 \leq M$. Such an M is called an *amalgam* of N_1 and N_2 over N_0 or just an *amalgam* if the rest is clear from the context. Notice that we do not formally require $V_M = V_{N_1} \cup V_{N_2}$. However, if an amalgam M exists, its restriction to $M' \leq M$ with $V_{M'} = V_{N_1} \cup V_{N_2}$ is an amalgam as well, so we can always assume that M only has nodes from N_1 and N_2 .

Definition 1 (Network Amalgamation Property) Let \mathbf{A} be a qualitative calculus (NA). \mathbf{A} has Network Amalgamation Property (NAP), if any V-formation (N_0, N_1, N_2) of networks over \mathbf{A} can be amalgamated by a network M over \mathbf{A} .

Clearly NAP is a hard property to come by, so some restrictions are necessary. One such restriction calls for the common subnetwork N_0 to be small in the following sense.

Definition 2 (k -Amalgamation Property) Let \mathbf{A} be a qualitative calculus (NA). \mathbf{A} has k Amalgamation Property (k -AP), if any V-formation (N_0, N_1, N_2) of networks over \mathbf{A} , such that $|N_0| \leq k$, can be amalgamated by a network M over \mathbf{A} .




Figure 1. (a) 3-extensibility and, (b) 4-extensibility. Both amalgamate over the edge (1,2). The dashed arrows represent the new edges.

It is obvious that n -AP implies m -AP for $n \geq m$. The smallest interesting case for a qualitative calculus is that of 2-AP. We will approach it step by step, beginning with $|N_1| = |N_2| = 3$, i.e., amalgamation of two triangles over a common edge. We will show that this follows from the associativity of \mathbf{A} . The next case, namely, $|N_1| = 4$ and $|N_2| = 3$ (adding a triangle to a square) turns out to be crucial. We will analyse it in some detail and then show that certain strong version of this case implies 2-AP for atomic networks.

3 EXTENSIBILITY

In this section we deal with 2-AP for the case with $|N_2| = 3$, which can be seen as extending an a-closed network N_1 by a triangle N_2 over a common edge. We refer to this as a *one-shot extension* [6].

Definition 3 ((generic) k -extensibility) Let \mathbf{A} be a qualitative calculus (NA) and k a natural number. \mathbf{A} is k -extensible if any atomic V-formation (N_0, N_1, N_2) of networks over \mathbf{A} , such that $|N_0| = 2$, $|N_1| = k$ and $|N_2| = 3$, can be amalgamated by a network $|M|$ over \mathbf{A} . If N_i ($i \in \{0, 1, 2\}$) are non-atomic, then \mathbf{A} is generically k -extensible (see Figure 1).

Lemma 1 Let \mathbf{A} be a RA. If \mathbf{A} is associative, then \mathbf{A} is generically 3-extensible.

Proof sketch. Let $N_0 = \{1, 2\}$, $N_1 = \{0, 1, 2\}$ and $N_2 = \{1, 2, 3\}$. Put $R_{03} = R_{01} \diamond R_{13} \cap R_{02} \diamond R_{23}$. By associativity, $R_{12} \neq \emptyset$. We need to show that the triangles $\{0, 1, 3\}$ and $\{0, 2, 3\}$ are a-closed. By symmetry it suffices to prove it for $\{0, 1, 3\}$, so we need to show three inclusions:

$$(R_{01} \diamond R_{13}) \cap (R_{02} \diamond R_{23}) \leq R_{01} \diamond R_{13} \quad (1)$$

$$R_{13} \leq R_{10} \diamond [(R_{01} \diamond R_{13}) \cap (R_{02} \diamond R_{23})] \quad (2)$$

$$R_{01} \leq [(R_{01} \diamond R_{13}) \cap (R_{02} \diamond R_{23})] \diamond R_{31} \quad (3)$$

The first of these is trivial, the two others follow from relation algebra identities. To show 3-extensibility, put $R_{03} = r_{01} \diamond r_{13} \cap r_{02} \diamond r_{23}$, where r_{ij} are atoms. Then any refinement of R_{03} satisfies the inclusions above, so any atomic refinement r_{03} satisfies them as well. ■

Since algebras that fail associativity are somewhat pathological, the above lemma is widely applicable.

Unlike 3-extensibility, 4-extensibility may fail in associative algebras, indeed even in representable ones. Consider the group \mathbb{Z}_7 (the integers under addition modulo 7) and for $x, y \in \mathbb{Z}_7$ define




Figure 2. (a) The RA \mathbf{B}_9 and, (b) The network \mathbb{S} that is not 4-extensible

- xIy if $x = y$
- xGy if $x = y \pm 1 \pmod{7}$
- xBy if $x = y \pm 2 \pmod{7}$
- xRy if $x = y \pm 3 \pmod{7}$

Then, $\{I, R, G, B\}$ are atoms of a representable relation algebra. Its representation using red for R , green for G and blue for B is shown in Figure 2. This algebra is known as \mathbf{B}_9 (cf. [5]). Consider the network $\mathbb{S} = \{0, 1, 2, 3\}$ with $\ell(0, 1) = R = \ell(2, 3)$, $\ell(0, 3) = B = \ell(1, 2)$, $\ell(0, 2) = G = \ell(1, 3)$, and $\ell(i, i) = I$, $\ell(i, j) = \ell(j, i)$. Verifying that \mathbb{S} is a-closed but not extensible is left to the reader as an exercise. Since \mathbb{S} is atomic, \mathbf{B}_9 is not 4-extensible. We will return to \mathbb{S} twice more in this paper, hence the fancy font.

We did not find any equations that would imply 4-extensibility in a manner similar to the role of associativity in Lemma 1. Checking generic 4-extensibility exhaustively takes too long for even a relatively small calculus such as RCC8. However, we could construct RCC8 networks for which generic 4-extensibility fails. Interestingly, all such networks we managed to construct contained relations that are known to be NP-hard (cf. [11]). On the other hand, 4-extensibility can be exhaustively tested by a program that performs checks on all atomic a-closed networks with four nodes.

Theorem 1 If a qualitative calculus (\mathbf{A}, U, μ) is extensional and \mathbf{A} is not 4-extensible, then a-closure does not decide consistency for networks of atomic relations.

Proof sketch. In an extensional calculus, consistent networks can always be extended by one-shot [6]. However, if \mathbf{A} is not 4-extensible, then there exists an atomic network N on four nodes that has no a-closed one-shot extension. Therefore N is not consistent. ■

One example of such an algebra is \mathbf{B}_9 , with \mathbb{S} in place of N .

3.1 Strong 4-extensibility

4-extensibility allows two networks of size 3 and 4 respectively to be combined over one edge without introducing inconsistencies. In this section, we show a special case of 4-extensibility that allows us to combine any two atomic networks of arbitrary size over one edge.

Definition 4 (Strong 4-extensibility) Let \mathbf{A} be a qualitative calculus (NA). \mathbf{A} is strongly 4-extensible if any V-formation (N_0, N_1, N_2) of atomic networks over \mathbf{A} , with $N_0 = \{1, 2\}$, $N_1 = \{0, 1, 2\}$ and $N_2 = \{1, 2, 3, 4\}$, can be amalgamated by a network $|M|$ over \mathbf{A} , such that for all $i \in N_2 \setminus N_0$

$$R_{i0} = (r_{i1} \diamond r_{10}) \cap (r_{i2} \diamond r_{20})$$

It follows easily by triangle identities that strong 4-extensibility implies 4-extensibility. The beauty of strong 4-extensibility is that for a given one-shot extension, labels for new edges are precisely the

intersections of compositions of labels on existing edges. This property is in fact possessed by both RCC8 and IA and can be checked even more efficiently than simple 4-extensibility.

Theorem 2 *If a NA \mathbf{A} is strongly 4-extensible, then \mathbf{A} has 2-Amalgamation Property if N_1, N_2 are atomic.*

Proof sketch. Let (N_0, N_1, N_2) be a V-formation of atomic networks, with $N_0 = \{0, 1\}$. Let $M = N_1 \cup N_2$ be the network retaining all the labels from N_1 and N_2 and with the new labels for edges (x, y) with $x \in N_i \setminus N_j$ and $y \in N_j \setminus N_i$ ($\{i, j\} = \{1, 2\}$) defined by $\ell(x, y) = r_{x0} \diamond r_{0y} \cap r_{x1} \diamond r_{1y}$. We will show that M is a-closed. Suppose the contrary. Then, there is a triangle in M with edges labelled by A, B, C , such that $C \not\leq A \diamond B$. Now, A, B and C cannot all be edges from N_i ($i \in \{1, 2\}$), for N_i is a-closed. So at least one of A, B, C is of the form $\ell(x, y)$ with $x \in N_i \setminus N_j$ and $y \in N_j \setminus N_i$ ($\{i, j\} = \{1, 2\}$). Notice also that there at most two of A, B, C can be such (three such edges do not form a triangle). We have then two cases. If there is exactly one such edge among A, B, C , it violates the assumption of 3-extensibility; if there are exactly two such edges, then it violates the assumption of strong 4-extensibility. Thus, M is a-closed as claimed. ■

The above theorem showed that if the calculus is strong 4-extensible, then we can amalgamate any two atomic networks over one edge. In the following we will show additional benefits of strong 4-extensibility for a qualitative calculus or relation algebra.

Definition 5 (One-Shot Extensibility [6]) *A qualitative calculus (\mathbf{A}, U, μ) is one-shot extensible if any consistent atomic V-formation (N_0, N_1, N_2) with $|N_0| = 2$ and $|N_2| = 3$, can be amalgamated by a consistent atomic network M .*

Corollary 1 *If a qualitative calculus \mathbf{A} is strongly 4-extensible, and a-closure decides consistency for networks of atomic and universal relations, then \mathbf{A} is one-shot extensible.*

One-shot extensibility was used in [6] to prove (for certain \mathbf{A}) that tractability of a set of relations S is equivalent to tractability of its closure \bar{S} under weak composition, intersection and converse. The method from [6] involves numerous manual calculations in the semantics of \mathbf{A} . However, if we know that a-closure happens to decide consistency for networks of atomic and universal relations in a qualitative calculus, as it for example does in RCC8 [2], then a simple check on the composition table for strong 4-extensibility is sufficient to prove one-shot extensibility.

Definition 6 (One-Shot Proto-Extensibility) *A qualitative calculus (NA) \mathbf{A} is one-shot proto-extensible if any atomic V-formation (N_0, N_1, N_2) with $|N_0| = 2$ and $|N_2| = 3$, can be amalgamated by an atomic network M .*

One-shot proto-extensibility ensures that the amalgam has an a-closed atomic refinement. Its advantage over one-shot extensibility is that it is a syntactic notion that is independent to any (weak) representation. Any representable algebra is trivially one-shot extensible relative to its representation.

Theorem 3 *Any one-shot proto-extensible RA is representable.*

Proof sketch. Let \mathbf{A} be a RA with the required property. We build a representation of \mathbf{A} inductively, beginning with any atomic a-closed triangle. At any given stage i , we have constructed an atomic a-closed network N_i . By one-shot proto-extensibility, we can pick any

atomic a-closed triangle T and add it to N_i , in effect amalgamating N_i and T over an edge that they share, obtaining an atomic a-closed network N_{i+1} . Let $N = \bigcup_{i \in \omega} N_i$. Define $\mu: A \rightarrow N$ putting $\mu(a) = \{(x, y) : \ell_N(x, y) = a\}$ for an atom $a \in A$. By finiteness of \mathbf{A} , each $u \in A$ is a join of finitely many atoms. Thus, we can extend μ onto the whole universe of A setting $\mu(u) = \mu(a_1) \cup \dots \cup \mu(a_n)$, where a_1, \dots, a_n are atoms with $u = a_1 \vee \dots \vee a_n$. It can be verified that the so defined μ is a representation of \mathbf{A} . ■

It is not the case that one-shot extensibility implies one-shot proto-extensibility, even for representable algebras. This is connected to the existence of atomic a-closed networks that are not consistent. A counterexample is again provided by \mathbf{B}_9 , which is representable, hence one-shot extensible, but not one-shot proto-extensible, as the network \mathbb{S} in Figure 2 shows.

4 ATOMIC REFINEMENT OF AMALGAMATED NETWORKS

In the previous section we showed that a NA \mathbf{A} has the 2-Amalgamation Property over atomic networks if it is strongly 4-extensible. Then if a calculus (\mathbf{A}, U, μ) has the property that a-closure decides consistency for networks of atomic and universal relations, there is always an atomic amalgam of the two networks, hence the calculus is one-shot proto-extensible. However, this is not a satisfactory result, as one-shot proto-extensibility is a purely syntactic concept based on the relation algebra, and we want to be able to prove it without resorting to the semantics of the qualitative calculus. We want a procedure that ensures the amalgam always has an a-closed atomic refinement. Such a procedure would provide a purely syntactic way to prove one-shot proto-extensibility, and hence representability.

4.1 Flexibility Ordering

Under strong 4-extensibility, each non-atomic relation in the amalgam of two networks over a common edge is precisely the intersection of the two paths from nodes in one network to another. One way to ensure there is always an atomic refinement to these relations such that the entire network is a-closed is to have a flexible atom (cf. [9]). A relation algebra with a set of atoms \mathcal{B} has a flexible atom a if the following condition hold:

$$\exists a \in \mathcal{B} : \forall b, c \in \mathcal{B} \setminus \{1'\}, a \leq b \diamond c$$

A flexible atom is contained in any composition of two atomic relations, so to make an amalgam atomic and a-closed one would just need to replace all the non-atomic relations in it by the flexible atom.

However, requiring a flexible atom is a very strong condition, and we do not know of a qualitative calculus, whose associated algebra has this property. Instead, we propose to construct an ordering of atoms that will emulate this property when refining amalgams, given the algebra has strong 4-extensibility. That is, we create a sequence of atomic relations, such that for any non-atomic edge R in the amalgam, we can refine it to the first element in the sequence that is contained in R , and the network remains a-closed.

Formally, let \mathbf{A} be a relation algebra with a set of atoms \mathcal{B} and S be a sequence of its atoms. A *choice refinement* of a non-atomic relation R over S is the first member of S that is a refinement of R .

Definition 7 (Flexibility Ordering) *For a strongly 4-extensible relation algebra \mathbf{A} , its Flexibility Ordering is a sequence S of atomic relations, such that for any amalgam M of an atomic V-formation*

(N_0, N_1, N_2) with $|N_0| = 2$, the non-atomic relations from M can be replaced by their respective choice refinements over S and the resulting network is a-closed.

The idea is that we define a sequence S of atomic relations such that in any M , when we replace a non-atomic edge R by its choice refinement r over S , it will never be inconsistent with the atomic edges of M , or atomic edges which arise as choice refinements of other non-atomic relations in M that are prior or equal to r in S .

To construct such a sequence, we propose an automated procedure that consists of two parts: First, for a given sequence S , that may not cover all cases, we test if a new atomic relation r that is not in S to see if it is compatible with S . That is, for an amalgam M of any two atomic a-closed network $\{0, 1, 2\}$ and $\{1, 2, 3, 4\}$, in the case that no current member of S is contained in the new edge R_{03} but r is, we check whether the following hold:

1. If R_{04} is already atomic, then when we replace R_{03} with r , the triangle $\{0, 3, 4\}$ is a-closed.
2. Else if there exists a choice refinement r_{04} of R_{04} over S , then when we replace R_{03} with r , R_{04} with r_{04} , the triangle $\{0, 3, 4\}$ is a-closed.
3. Else if R_{04} contains r , then when we replace both R_{03} and R_{04} by r , the triangle $\{0, 3, 4\}$ is a-closed.

If the above hold for all such amalgams M , r and S is compatible.

The second part involves the construction of such a list. Starting from an empty list, we incrementally add atoms that pass the compatibility test with the list, and backtrack when no further candidates can pass the test. It is worth noting that each branch of the search tree may terminate early: e.g. if an atom a is not compatible with an empty ordering, then we do not have to test any entries with a at the head of the ordering.

Theorem 4 *If a NA is strongly 4-extensible, and it has a Flexibility Ordering, then it is one-shot proto-extensible.*

Proof sketch. From Theorem 2 we get a network M that is a-closed, but the new edges between N_1 and N_2 may not be atomic. However, with a Flexibility Ordering we can refine each of these edges to atomic relations, knowing that similar atomic refinements of other new edges will not introduce an inconsistent triple, since we have checked all possible cases in the construction of the Flexibility Ordering. Therefore the entire network is refined to be atomic and a-closed, thus the relation algebra is one-shot proto-extensible. ■

This general result, together with Theorem 3, allows us to prove representability of a RA \mathbf{A} from its composition table. This means that \mathbf{A} can be a part of an extensional qualitative calculus (\mathbf{A}, U, μ) . It also implies that consistency can be preserved when amalgamating two atomic a-closed networks over two nodes if we know that a-closure decides consistency for only atomic relations.

4.2 Empirical Evaluations of Flexibility Ordering on RCC8 and Interval Algebra

Both RCC8 and IA are prime candidates to test for Flexibility Orderings, as they are well known and non-trivial calculi in the spatial-temporal domain, and their respective relation algebras are both strongly 4-extensible. For RCC8, the procedure found the Flexibility Ordering: (DC, EC, PO, TPP, TPPI), whereas for IA, the procedure found $(<, di, o, s, oi, f)$. Hence we have proved from their composition table that their relation algebras are representable.

Computationally the worst case of the procedure is $O(|\mathcal{B}|!)$. However, this would be extremely rare, as most branches of the search tree will be terminated earlier than exhaustive search, thus trimming down a majority of potential search space. For IA, with 13 atoms, the procedure found an ordering in 4 seconds on a Intel Core2Duo 2.4GHz processor with 2GB RAM, and for RCC8 it found a solution in less than a second. Therefore, our procedure is widely applicable.

5 CONCLUSION AND FUTURE WORK

We provided sufficient conditions to amalgamate two atomic networks of any size over a common edge. Hence, for a calculus where a-closure decides consistency for networks that contain only atomic and universal relations, two atomic networks can always be consistently amalgamated. The property of strong 4-extensibility, together with other known results, also tell us when a-closure does not decide consistency for atomic networks. It provides an efficient computational test to check, for a non-extensional calculus, whether complexity results for a set of relations can be transferred to its closure.

More importantly, we have provided a procedure that proves the resulting amalgamated network has an a-closed atomic refinement, independent of any information about the domain of the calculus. This allows us to prove representability of a relation algebra from its composition table. It preserves consistency under amalgamation of two atomic networks over two nodes, if a-closure decides consistency for networks of atomic relations.

The first obvious future step is to see whether two atomic a-closed networks can be amalgamated over n nodes for $n > 2$. Then we take it to the non-atomic ones. It is also interesting to see under what conditions a calculus has the Network Amalgamation Property, that is, networks can be combined regardless of number of shared nodes.

Secondly, our proposed notion of one-shot proto-extensibility is a sufficient, but not necessary condition for representability of a relation algebra. There are other representable relation algebras which are not one-shot proto-extensible. It would be interesting to see if there are any connections between one-shot proto-extensibility and Hirsch-Hodkinson type games [4], and whether Hirsch-Hodkinson games can be interpreted as a sequence of one-shot extensions.

REFERENCES

- [1] J.F. Allen, Maintaining knowledge about temporal intervals, *Comm. ACM*, **26**, 832–843, 1983.
- [2] B. Bennett, Determining consistency of topological relations, *Constraints*, **3**, 213–225, 1998.
- [3] I. Düntsch, Relation algebras and their application in temporal and spatial reasoning, *Artif. Intell. Rev.*, **23-4**, 315–357, 2005.
- [4] R. Hirsch, I. Hodkinson, *Relation algebras by games*, Elsevier, 2002.
- [5] P. Jipsen, E. Lukács, Minimal relation algebras, *Algebra Universalis*, **32**, 189–203, 1994.
- [6] S. Li and H. Wang, RCC8 binary constraint network can be consistently extended, *Artificial Intelligence*, **170**, 1–18, (2006).
- [7] G. Ligozat, J. Renz, What is a qualitative calculus? A general framework, *Proc. PRICAI'04*, 53–64, 2004.
- [8] R. Maddux, *Relation algebras*, Elsevier, 2006.
- [9] R. Maddux, Some varieties containing relation algebras, *Trans. Amer. Math. Soc.*, **2**, 501–526, 1982.
- [10] D.A. Randell, Z. Cui, and A.G. Cohn, A spatial logic based on regions and connection. *Proc. KR'92*, 165–176, 1992.
- [11] J. Renz, Maximal tractable fragments of the Region Connection Calculus: A complete analysis, *Proc. IJCAI'99*, 448–455, 1999.
- [12] J. Renz, G. Ligozat, Weak composition for qualitative spatial and temporal reasoning, *Proc. CP'05*, 534–548, 2005.
- [13] J. Renz, B. Nebel, Qualitative spatial reasoning using constraint calculi, *Handbook of Spatial Logics*, Springer, 161–215, 2007.