

Weak Composition for Qualitative Spatial and Temporal Reasoning

Jochen Renz¹ and Gérard Ligozat²

¹ National ICT Australia*, Knowledge Representation and Reasoning Group,
UNSW Sydney, NSW 2052, Australia
jochen.renz@nicta.com.au

² LIMSI-CNRS, Université Paris-Sud, 91403 Orsay, France
ligozat@limsi.fr

Abstract. It has now been clear for some time that for many qualitative spatial or temporal calculi, for instance the well-known RCC8 calculus, the operation of composition of relations which is used is actually only weak composition, which is defined as the strongest relation in the calculus that contains the real composition. An immediate consequence for qualitative calculi where weak composition is not equivalent to composition is that the well-known concept of path-consistency is not applicable anymore. In these cases we can only use algebraic closure which corresponds to applying the path-consistency algorithm with weak composition instead of composition.

In this paper we analyse the effects of having weak compositions. Starting with atomic CSPs, we show under which conditions algebraic closure can be used to decide consistency in a qualitative calculus, how weak consistency affects different important techniques for analysing qualitative calculi and under which conditions these techniques can be applied. For our analysis we introduce a new concept for qualitative relations, the “closure under constraints”. It turns out that the most important property of a qualitative calculus is not whether weak composition is equivalent to composition, but whether the relations are closed under constraints. All our results are general and can be applied to all existing and future qualitative spatial and temporal calculi. We close our paper with a road map of how qualitative calculi should be analysed. As a side effect it turns out that some results in the literature have to be reconsidered.

1 Introduction

The domain of qualitative temporal reasoning underwent a major change when Allen [1] proposed a new calculus which up to a degree resulted in embedding it in the general paradigm of constraint satisfaction problems (CSPs). CSPs have their well-established sets of questions and methods, and qualitative temporal reasoning, and more recently qualitative spatial reasoning, has profited significantly from developing tools and methods analogous to those of classical constraint satisfaction. In particular, a central question for classical constraint networks is the consistency problem: is the set of constraints

* National ICT Australia is funded through the Australian Government’s *Backing Australia’s Ability* initiative, in part through the Australian Research Council.

specified by the constraint network consistent, that is, can the variables be instantiated with values from the domains in such a way that all constraints are satisfied?

Part of the apparatus for solving the problem consists of filtering algorithms which are able to restrict the domains of the variables without changing the problem, while remaining reasonably efficient from a computational point of view. Various algorithms such as arc consistency, path consistency, and various notions of k -consistency have been extensively studied in that direction. Reasoning about temporal or spatial qualitative constraint networks on the same line as CSPs has proved a fruitful approach. Both domains indeed share a general paradigm. However, there is a fundamental difference between the two situations:

- Relations in classical CSPs are finite relations, so they can be explicitly manipulated as sets of tuples of elements of a finite domain. In other terms, relations are given and processed in an extensional way.
- By contrast, relations in (most) qualitative temporal and spatial reasoning formalisms are provided in intentional terms – or, to use a more down-to-earth expression, they are infinite relations, which means that there is no feasible way of dealing with them extensionally.

But is that such an important point? We think it is, although this was not apparent for Allen's calculus. The differences began to appear when it became obvious that new formalisms, such as for instance the RCC8 calculus [19], could behave in a significantly different way than Allen's calculus. The differences have to do with changes in the notion of composition, with the modified meaning of the the classical concept of path-consistency and its relationship to consistency, and with the inapplicability of familiar techniques for analysing qualitative calculi.

1.1 Composition

Constraint propagation mainly uses the operation of composition of two binary relations. In the finite case, there is only a finite number of binary relations. In Allen's case, although the domains are infinite, the compositions of the thirteen atomic relations are themselves unions of atomic relations. But this is not the case in general, where insisting on genuine composition could lead to considering an infinite number of relations, whereas the basic idea of qualitative reasoning is to deal with a finite number of relations. The way around the difficulty consists in using weak composition, which only approximates true composition.

1.2 Path Consistency and Other Qualitative Techniques

When only weak composition is used then some algorithms and techniques which require true composition can only use weak composition instead. This might lead to the inapplicability of their outcomes. Path-consistency, for example, relies on the fact that a constraint between two variables must be at least as restrictive as every path in the constraint network between the same two variables. The influence of the paths depends on composition of relations on the path. If we use algebraic closure instead of path-consistency, which is essentially path-consistency with weak composition, then

we might not detect restrictions imposed by composition and therefore the filtering effect of algebraic closure is weaker than that of path-consistency. As a consequence it might not be possible to use algebraic closure as a decision procedure for certain calculi. Likewise, commonly used reduction techniques lose their strength when using only weak composition and might not lead to valid reductions.

The main goal of this paper is to thoroughly analyse how the use of weak composition instead of composition affects the applicability of the common filtering algorithms and reduction techniques and to determine under which conditions their outcomes match that of their composition-based counterparts.

1.3 Related Work

The concepts of weak composition and algebraic closure are not new. Although there has not always been a unified terminology to describe these concepts, many authors have pointed out that composition tables do not necessarily correspond to the formal definition of composition [4,5,8,13]. Consequently, many researchers have been interested in finding criteria for (refutation) completeness of compositional reasoning, and Bennett et al. ([4,5]) posed this as a challenge and conjectured a possible solution. Later work focused on dealing with this problem for RCC8 [6,11]. In particular Li and Ying ([11]) showed that no RCC8 model can be interpreted extensionally, i.e., for RCC8 composition is always only a weak composition, which gives a negative answer to Bennett et al.'s conjecture. Our paper is the first to give a general account on the effects of having weak composition and a general and clear criterion for the relationship between algebraic closure and consistency. Therefore, the results of this paper are important for establishing the foundations of qualitative spatial and temporal reasoning and are a useful tool for investigating and developing qualitative calculi.

The structure of the paper is as follows: Section 2 introduces the main notions and terminology about constraint networks, various notions of consistency and discusses weak composition and algebraic closure. Section 3 provides a characterisation of those calculi for which algebraic closure decides consistency for atomic networks. Section 4 examines the conditions under which general techniques of reduction can be applied to a qualitative calculus. Finally, Section 5 draws general conclusions in terms of how qualitative calculi should be analysed, and shows that some existing results have to be revisited in consequence.

2 Background

2.1 Constraint Networks

Knowledge between different entities can be represented by using constraints. A *binary relation* R over a domain \mathcal{D} is a set of pairs of elements of \mathcal{D} , i.e., $R \subseteq \mathcal{D} \times \mathcal{D}$. A *binary constraint* xRy between two variables x and y restricts the possible instantiations of x and y to the pairs contained in the relation R . A *constraint satisfaction problem* (CSP) consists of a finite set of variables \mathcal{V} , a domain \mathcal{D} with possible instantiations for each variable $v_i \in \mathcal{V}$ and a finite set \mathcal{C} of constraints between the variables of \mathcal{V} . A *solution* of a CSP is an instantiation of each variable $v_i \in \mathcal{V}$ with a value $d_i \in \mathcal{D}$ such that all

constraints of \mathcal{C} are satisfied, i.e., for each constraint $v_i R v_j \in \mathcal{C}$ we have $(d_i, d_j) \in R$. If a CSP has a solution, it is called *consistent* or *satisfiable*. Several algebraic operations are defined on relations that carry over to constraints, the most important ones being *union* (\cup), *intersection* (\cap), and *complement* ($\bar{\cdot}$) of a relation, defined as the usual set-theoretic operators, as well as *converse* (\cdot^{-1}) defined as $R^{-1} = \{(a, b) \mid (b, a) \in R\}$ and *composition* (\circ) of two relations R and S which is the relation $R \circ S = \{(a, b) \mid \exists c : (a, c) \in R \text{ and } (c, b) \in S\}$.

2.2 Path-Consistency

Because of the high complexity of deciding consistency, different forms of local consistency and algorithms for achieving local consistency were introduced. Local consistency is used to prune the search space by eliminating local inconsistencies. In some cases local consistency is even enough for deciding consistency. Montanari [15] developed a form of local consistency which Mackworth [14] later called path-consistency. Montanari's notion of path-consistency considers all paths between two variables. Mackworth showed that it is equivalent to consider only paths of length two, so path-consistency can be defined as follows: a CSP is *path-consistent*, if for every instantiation of two variables $v_i, v_j \in \mathcal{V}$ that satisfies $v_i R_{ij} v_j \in \mathcal{C}$ there exists an instantiation of every third variable $v_k \in \mathcal{V}$ such that $v_i R_{ik} v_k \in \mathcal{C}$ and $v_k R_{kj} v_j \in \mathcal{C}$ are also satisfied. Formally, for every triple of variables $v_i, v_j, v_k \in \mathcal{V} : \forall d_i, d_j : [(d_i, d_j) \in R_{ij} \rightarrow \exists d_k : ((d_i, d_k) \in R_{ik} \wedge (d_k, d_j) \in R_{kj})]$. Montanari also developed an algorithm that makes a CSP path-consistent, which was later simplified and called *path-consistency algorithm* or *enforcing path-consistency*. A path-consistency algorithm eliminates locally inconsistent tuples from the relations between the variables by successively applying the following operation to all triples of variables $v_i, v_j, v_k \in \mathcal{V}$ until a fixpoint is reached: $R_{ij} := R_{ij} \cap (R_{ik} \circ R_{kj})$. If the empty relation occurs, then the CSP is inconsistent. Otherwise the resulting CSP is path-consistent.

2.3 Varieties of k-Consistency

Freuder [7] generalised path-consistency and the weaker notion of arc-consistency to k-consistency: A CSP is *k-consistent*, if for every subset $\mathcal{V}_k \subset \mathcal{V}$ of k variables the following holds: for every instantiation of $k - 1$ variables of \mathcal{V}_k that satisfies all constraints of \mathcal{C} that involve only these $k - 1$ variables, there is an instantiation of the remaining variable of \mathcal{V}_k such that all constraints involving only variables of \mathcal{V}_k are satisfied. So if a CSP is k-consistent, we know that each consistent instantiation of $k - 1$ variables can be extended to any k-th variable. A CSP is strongly k-consistent, if it is i-consistent for every $i \leq k$. If a CSP with n variables is strongly n -consistent (also called *globally consistent*) then a solution can be constructed incrementally without backtracking. 3-consistency is equivalent to path-consistency, 2-consistency is equivalent to arc-consistency.

2.4 Qualitative Spatial and Temporal Relations

The main difference of spatial or temporal CSPs to normal CSPs is that the domains of the spatial and temporal variables are usually infinite. For instance, there are infinitely

many time points or temporal intervals on the time line and infinitely many regions in a two or three dimensional space. Hence it is not feasible to represent relations as sets of tuples, nor is it feasible to apply algorithms that enumerate values of the domains. Instead, relations can be used as symbols and reasoning has to be done by manipulating symbols. This implies that the calculus, which deals with extensional relations in the finite case, becomes intensional in the sense that it manipulates symbols which stand for infinite relations. The usual way of dealing with relations in qualitative spatial and temporal reasoning is to have a finite (usually small) set \mathcal{A} of jointly exhaustive and pairwise disjoint (JEPD) relations, i.e., each possible tuple $(a, b) \in \mathcal{D} \times \mathcal{D}$ is contained in exactly one relation $R \in \mathcal{A}$. The relations of a JEPD set \mathcal{A} are called *atomic relations*. The full set of available relations is then the powerset $\mathcal{R} = 2^{\mathcal{A}}$ which enables us to represent indefinite knowledge, e.g., the constraint $x\{R_i, R_j, R_k\}y$ specifies that the relation between x and y is one of R_i, R_j or R_k , where R_i, R_j, R_k are atomic relations.

2.5 Composition and Weak Composition

Using these relations we can now represent qualitative spatial or temporal knowledge using CSPs and use constraint-based methods for deciding whether such a CSP is consistent, i.e., whether it has a solution. Since we are not dealing with explicit tuples anymore, we have to compute the algebraic operators for the relations. These operators are the only connection of the relation symbols to the tuples contained in the relations and they have to be computed depending on the tuples contained in the relations. Union, complement, converse, and intersection of relations are again the usual set-theoretic operators while composition is not as straightforward. Composition has to be computed only for pairs of atomic relations since composition of non-atomic relations is the union of the composition of the involved atomic relations. Nevertheless, according to the definition of composition, we would have to look at an infinite number of tuples in order to compute composition of atomic relations, which is clearly not feasible. Fortunately, many domains such as points or intervals on a time line are ordered or otherwise well-structured domains and composition can be computed using the formal definitions of the relations. However, for domains such as arbitrary spatial regions that are not well structured and where there is no common representation for the entities we consider, computing the true composition is not feasible and composition has to be approximated by using *weak composition* [6]. Weak composition (\diamond) of two relations S and T is defined as the strongest relation $R \in 2^{\mathcal{A}}$ which contains $S \circ T$, or formally, $S \diamond T = \{R_i \in \mathcal{A} \mid R_i \cap (S \circ T) \neq \emptyset\}$. The advantage of weak composition is that we stay within the given set of relations $\mathcal{R} = 2^{\mathcal{A}}$ while applying the algebraic operators, as \mathcal{R} is by definition closed under weak composition, union, intersection, and converse.

In cases where composition cannot be formally computed (e.g. RCC8 [19]), it is often very difficult to determine whether weak composition is equivalent to composition or not. Usually only non-equality can be shown by giving a counterexample, while it is very difficult to prove equality. However, weak composition has also been used in cases where composition could have been computed because the domain is well-structured and consists of pairs of ordered points, but where the authors did not seem to be aware that \mathcal{R} is not closed under composition (e.g. INDU, PDN, or PIDN [17,16,18]).

Example 1 (Region Connection Calculus RCC8 [19]). RCC8 is a topological constraint language based on eight atomic relations between extended regions of a topological space. Regions are regular subsets of a topological space, they can have holes and can consist of multiple disconnected pieces. The eight atomic relations *DC* (disconnected), *EC* (externally connected), *PO* (partial overlap), *EQ* (equal), *TPP* (tangential proper part), *NTPP* (non-tangential proper part) and their converse relations *TPPi*, *NTPPi* were originally defined in first-order logic. It was shown by Düntsch [6], that the composition of RCC8 is actually only a weak composition. Consider the consistent RCC8 constraints $B\{TPP\}A$, $B\{EC\}C$, $C\{TPP\}A$. If *A* is instantiated as a region with two disconnected pieces and *B* completely fills one piece, then *C* cannot be instantiated. So *TPP* is not a subset of $EC \circ TPP$ [11] and consequently RCC8 is not closed under composition.

2.6 Algebraic Closure

When weak composition differs from composition, we cannot apply the path-consistency algorithm as it requires composition and not just weak composition. We can, however, replace the composition operator in the path-consistency algorithm with the weak composition operator. The resulting algorithm is called the *algebraic closure algorithm* [12] which makes a network *algebraically closed* or *a-closed*.

If weak composition is equal to composition, then the two algorithms are also equivalent. But whenever we have only weak composition, an a-closed network is not necessarily path-consistent as there are relations *S* and *T* such that $S \circ T \subset S \diamond T$. So there are tuples $(u, v) \in S \diamond T$ for which there is no *w* with $(u, w) \in S$ and $(w, v) \in T$, i.e., for which $(u, v) \notin S \circ T$. This contradicts the path-consistency requirements given above.

Path-consistency has always been an important property when analysing qualitative calculi, in particular as a method for identifying tractability. When this method is not available, it is not clear what the consequences of this will be. Will it still be possible to find calculi for which a-closure decides consistency even if weak composition differs from composition? What effect does it have on techniques used for analysing qualitative calculi which require composition and not just weak composition? And what is very important, does it mean that some results in the literature have to be revised or is it enough to reformulate them? These and related questions will be answered in the remainder of the paper. As an immediate consequence, unless we have proven otherwise, we should for all qualitative spatial and temporal calculi always assume that we are dealing with weak composition and that it is not equivalent to composition.

3 Weak Composition and Algebraic Closure

For analysing the effects of weak composition, we will mainly focus on its effects on the most commonly studied reasoning problem, the consistency problem, i.e., whether a given set Θ of spatial or temporal constraints has a solution. Recall that consistency means that there is at least one instantiation for each variable of Θ with a value from its domain which satisfies all constraints. This is different from global consistency which

Table 1. Does a-closure decide atomic CSPs depending on whether weak composition differs from composition?

	a-closure sufficient	a-closure not sufficient
weak composition = composition	Interval Algebra [1] rectangle algebra [9] block algebra [3]	STAR calculus [21] containment algebra [10] cyclic algebra [2]
weak composition \neq composition	RCC8 [19], discrete IA	INDU [17], PDN [16], PIDN [18]

requires strong k -consistency for all k . Global consistency cannot be obtained when we have only weak composition as we have no method for even determining 3-consistency. For the mere purpose of deciding consistency it actually seems overly strong to require any form of k -consistency as we are not interested in whether *any* consistent instantiation of k variables can be extended to $k + 1$ variables, but only whether there exists *at least one* consistent instantiation. Therefore it might not be too weak for deciding consistency to have only algebraic closure instead of path-consistency.

In the following we restrict ourselves to atomic CSPs, i.e., CSPs where all constraints are restricted to be atomic relations. If a-closure does not even decide atomic CSPs, it will not decide more general CSPs. We will later see how the results for atomic CSPs can be extended to less restricted CSPs. Let us first analyse for some existing calculi how the two properties whether a-closure decides atomic CSPs and whether weak composition differs from composition relate. We listed the results in Table 1 and they are significant:

Proposition 1. *Let \mathcal{R} be a finite set of qualitative relations. Whether a-closure decides consistency for atomic CSPs over \mathcal{R} is independent of whether weak composition differs from composition for relations in \mathcal{R} .*

This observation shows us that whether or not a-closure decides atomic CSPs does not depend on whether weak composition is equivalent to composition or not. Instead we will have to find another criterion for when a-closure decides atomic CSPs. In order to find such a criterion we will look at some examples where a-closure does not decide atomic CSPs and see if we can derive some commonalities.

Example 2 (STAR calculus [21]). Directions between two-dimensional points are distinguished by specifying an arbitrary number of angles which separate direction sectors. The atomic relations are the sectors as well as the lines that separate the sectors (see Figure 1 left). The domain is ordered so it is possible to compute composition. The relations are closed under composition. If more than two angles are given, then by using constraint configurations involving four or more variables, it is possible to refine the atomic relations that correspond to sectors to different particular angles (see Figure 1 right). By combining configurations that refine the same atomic relation to different angles, inconsistencies can be constructed that cannot be detected by a-closure. In this example we can see that even true composition can be too weak. Although we know the composition and all relations are closed under composition, it is possible to refine atomic relations using networks with more than three nodes.




Fig. 1. A STAR calculus with 3 angles resulting in 13 atomic relations (left). The right picture shows an atomic CSP whose constraints enforce that D must be 45 degrees to the left of B , i.e., the constraint $B\{11\}D$ is refined by the other constraints to the line orthogonal to relation 2. Therefore, the atomic relation 11 can be refined to a subatomic relation using the given constraints.

Example 3 (INDU calculus [17]). Allen’s 13 interval relations are combined with relative duration of intervals given in the form of a point algebra, i.e., INDU relations are of the form $R = I_\delta$ where I is an interval relation (precedes p , meets m , during d , starts s , overlaps o , finishes f , equal $=$, and the converse relations fi,oi,si,di,mi,pi) and δ a duration relation ($<$, $>$, $=$). This leads to only 25 atomic relations as some combinations are impossible, e.g., $a\{s\}b$ enforces that the duration of a must be less than that of b . Only weak composition is used, as for example the triple $a\{s_<\}b, a\{m_<\}c, c\{f_<\}b$ enforces that $a < 0.5*b$ and $c > 0.5*b$. So an instantiation where $a = 0.5*b$ cannot be extended to a consistent instantiation of c . In the same way it is possible to generate any metric duration constraint of the form $duration(x) R \alpha * duration(b)$ where $R \in \{<, >, =\}$ and α is a rational number. Consequently, it is possible to construct inconsistent atomic CSPs which are a-closed.

In both examples it is possible to refine atomic relations to subatomic relations that have no tuples in common, i.e., which do not overlap. This can be used to construct inconsistent examples which are still a-closed. Note that in the case of the interval algebra over integers it is possible to refine atomic relations to subatomic relations, e.g., $a\{p\}b, b\{p\}c$ leads to $a\{p+2\}c$, where $p+2$ indicates that a must precede c by at least 2 more integers than is required by the precedes relation. But since these new subatomic relations always overlap, it is not possible to construct inconsistencies which are a-closed. Let us formally define these terms.

Definition 1 (refinement to a subatomic relation). Let Θ be a consistent atomic CSP over a set A and $xRy \in \Theta$ a constraint. Let R' be the union of all tuples $(u, v) \in R$ that can be instantiated to x and y as part of a solution of Θ . If $R' \subset R$, then Θ refines R to the subatomic relation R' .

Definition 2 (closure under constraints). Let A be a set of atomic relations. A is closed under constraints if no relation $R \in A$ can be refined to non-overlapping subatomic relations, i.e., if for each $R \in A$ all subatomic relations $R' \subset R$ to which R can be refined to have a nonempty intersection.

In the following theorem we show that the observation made in these examples holds in general and we can prove in which cases a-closure decides atomic CSPs, which is

independent of whether weak composition differs from composition and only depends on whether the atomic relations are closed under constraints. Therefore, the new concept of closure under constraints turns out to be a very important property of qualitative reasoning.

Theorem 1. *Let \mathcal{A} be a finite set of atomic relations. Then a-closure decides consistency of CSPs over \mathcal{A} if and only if \mathcal{A} is closed under constraints.*

Proof Sketch. \Rightarrow : Given a set of atomic relations $\mathcal{A} = \{R_1, \dots, R_n\}$. We have to prove that if \mathcal{A} is not closed under constraints, then a-closure does not decide consistency over \mathcal{A} . \mathcal{A} is not closed under constraints means that there is an atomic relation $R_k \in \mathcal{A}$ which can be refined to non-overlapping subatomic relations using atomic sets of constraints over \mathcal{A} . We will prove this by constructing an a-closed but inconsistent set of constraints over \mathcal{A} for those cases where \mathcal{A} is not closed under constraints. We assume without loss of generality that if \mathcal{A} is not closed under constraints, there are at least two non-overlapping subatomic relations R_k^1, R_k^2 of R_k which can be obtained using the atomic sets of constraints Θ^1, Θ^2 (both are a-closed and consistent and contain the constraint $xR_k y$). We combine all tuples of R_k not contained in R_k^1 or R_k^2 to R_k^m and have that $R_k^1 \cup R_k^2 \cup R_k^m = R_k$ and that R_k^1, R_k^2, R_k^m are pairwise disjoint.

We can now form a new set of atomic relations \mathcal{A}' where R_k is replaced with R_k^1, R_k^2, R_k^m (analogous for R_k^{-1}). All the other relations are the same as in \mathcal{A} . The weak composition table of \mathcal{A}' differs from that of \mathcal{A} for the entries that contain R_k or R_k^{-1} . Since R_k^1 and R_k^2 can be obtained by atomic sets of constraints over \mathcal{A} , the entries in the weak composition table of \mathcal{A}' cannot be the same for R_k^1 and for R_k^2 . Therefore, there must be a relation $R_l \in \mathcal{A}$ for which the entries of $R_l \diamond R_k^1$ and of $R_l \diamond R_k^2$ differ. We assume that $R_l \diamond R_k = S$ and that $R_l \diamond R_k^1 = S \setminus S_1$ and $R_l \diamond R_k^2 = S \setminus S_2$, with $S, S_1, S_2 \in 2^A$ and $S_1 \neq S_2$. We chose a non-empty one, say S_1 , and can now obtain an inconsistent triple $xR_k^1 y, zR_l x, zS_1 y$ for which the corresponding triple $xR_k y, zR_l x, zS_1 y$ is consistent. Note that we use \mathcal{A}' only for identifying R_l and S_1 .

If we now consider the set of constraints $\Theta = \Theta^1 \cup \{zR_l x, zS_1 y\}$ (where z is a fresh variable not contained in Θ^1), then Θ is clearly inconsistent since Θ^1 refines $xR_k y$ to $xR_k^1 y$ and since $R_l \diamond R_k^1 = S \setminus S_1$. However, applying the a-closure algorithm to Θ (resulting in Θ') using the weak composition table of \mathcal{A} does not result in an inconsistency, since a-closure does not see the implicit refinement of $xR_k y$ to $xR_k^1 y$.

\Leftarrow : Proof by induction over the size n of Θ . Induction hypothesis: $P(n) = \{\text{For sets } \Theta \text{ of atomic constraints of size } n, \text{ if it is not possible to refine atomic relations to non-overlapping subatomic relations, then a-closure decides consistency for } \Theta.\}$ This is clear for $n \leq 3$. Now take an a-closed atomic CSP Θ of size $n+1$ over \mathcal{A} and assume that $P(n)$ is true. For every variable $x \in \Theta$ let Θ_x be the atomic CSP that results from Θ by removing all constraints that involve x . Because of $P(n)$, Θ_x is consistent for all $x \in \Theta$. Let R_x be the subatomic relation to which R is refined to in Θ_x and let R' be the intersection of R_x for all $x \in \Theta$. If R' is non-empty for every $R \in \mathcal{A}$, i.e., if it is not possible to refine R to non-overlapping subatomic relations, then we can choose a consistent instantiation of Θ_x which contains for every relation R only tuples of R' . Since no relation R of Θ_x can be refined beyond R' by adding constraints of Θ that involve x , it is clear that we can then also find a consistent instantiation for x , and thereby obtain a consistent instantiation of Θ . ■

This theorem is not constructive in the sense that it does not help us to prove that a-closure decides consistency for a particular calculus. But such a general constructive theorem would not be possible as it depends on the semantics of the relations and on the domains whether a-closure decides consistency. This has to be formally proven in a different way for each new calculus and for each new domain. What our theorem gives us, however, is a simple explanation why a-closure is independent of whether weak composition differs from composition: It makes no difference whatsoever whether non-overlapping subatomic relations are obtained via triples of constraints or via larger constellations (as in Example 2). In both cases a-closure cannot detect all inconsistencies. Our theorem also gives us both, a simple method for determining when a-closure does not decide consistency, and a very good heuristic for approximating when it does. Consider the following heuristic:

Does the considered domain enable more distinctions than those made by the atomic relations, and if so, can these distinctions be enforced by a set of constraints over existing relations?

This works for the three examples we already mentioned. It also works for any other calculus that we looked at. Take for instance the containment algebra which is basically the interval algebra without distinguishing directions [10]. So having directions would be a natural distinction and it is easy to show that we can distinguish relative directions by giving constraints: If a is disjoint from b and c touches b but is disjoint from a , then c must be on the same side of a as b . This can be used to construct a-closed inconsistent configurations. For RCC8, the domain offers plenty of other distinctions, but none of them can be enforced by giving a set of RCC8 constraints. This gives a good indication that a-closure decides consistency (which has been proven in [22]). If we restrict the domain of RCC8, e.g., to two-dimensional discs of the same size, then we can find distinctions which can be enforced by giving constraints.

When defining a new qualitative calculus by defining a set of atomic relations, it is desirable that algebraic closure decides consistency of atomic CSPs. Therefore, we recommend to test the above given heuristic when defining a new qualitative calculus and to make sure that the new atomic relations are closed under constraints. In section 5 we discuss the consequences of having a set of relations which is not closed under constraints.

4 Effects on Qualitative Reduction Techniques

In the analysis of qualitative calculi it is usually tried to transfer properties such as tractability or applicability of the a-closure algorithm for deciding consistency to larger sets of relations and ideally find the maximal sets that have these properties. Such general techniques involve composition of relations in one way or another and it is not clear whether they can still be applied if only weak composition is known and if they have been properly applied in the literature. It might be that replacing composition with weak composition and path-consistency with a-closure is sufficient, but it might also be that existing results turn out to be wrong or not applicable. In this section we look at two important general techniques for extending properties to larger sets of relations.




Fig. 2. (1) A consistent INDU network which becomes inconsistent when replacing $b\{s<, d<\}a$ with (2). From (1) we get $b > 0.5 * a$ and from (2) we get $b < 0.5 * a$.

The first technique is very widely used and is based on the fact that a set of relations $\mathcal{S} \subseteq 2^A$ and the closure $\widehat{\mathcal{S}}$ of \mathcal{S} under composition, intersection, and converse have the same complexity. This results from a proof that the consistency problem for $\widehat{\mathcal{S}}$ (written as $\text{CSPSAT}(\widehat{\mathcal{S}})$) can be polynomially reduced to $\text{CSPSAT}(\mathcal{S})$ by inductively replacing each constraint xRy over a relation $R \in \widehat{\mathcal{S}} \setminus \mathcal{S}$ by either $xSy \wedge xTy$ or by $xSz \diamond zTy$ for $S, T \in \mathcal{S}$ [22]. If we have only weak composition, then we have two problems. First, we can only look at the closure of \mathcal{S} under intersection, converse, and weak composition (we will denote this *weak closure* by $\widehat{\mathcal{S}}_w$). And, second, we can replace a constraint xRy over a relation $R \in \widehat{\mathcal{S}}_w \setminus \mathcal{S}$ only by $xSy \wedge xTy$ or by $xSz \diamond zTy$ for $S, T \in \mathcal{S}$. For $xSz \diamond zTy$ we know that it might not be a consistent replacement for xRy . In Figure 2 we give an example for a consistent set of INDU constraints which becomes inconsistent if we replace a non-atomic constraint by an intersection of two weak compositions of other INDU relations.

So it is clear that this widely used technique does not apply in all cases where we have only weak composition. In the following theorem we show when it can still be applied.

Theorem 2. *Let \mathcal{R} be a finite set of qualitative relations and $\mathcal{S} \subseteq \mathcal{R}$ a set of relations. Then $\text{CSPSAT}(\widehat{\mathcal{S}}_w)$ can be polynomially reduced to $\text{CSPSAT}(\mathcal{S})$ if a-closure decides consistency for atomic CSPs over \mathcal{R} .*

Proof Sketch. Consider an a-closed set Θ of constraints over $\widehat{\mathcal{S}}_w$. When inductively replacing constraints over $\widehat{\mathcal{S}}_w$ with constraints over \mathcal{S} , i.e., when replacing xRy where $R \in \widehat{\mathcal{S}}_w$ with xSz and zTy where $S \diamond T = R$ and $S, T \in \mathcal{S}$ and z is a fresh variable, then potential solutions are lost. However, all these triples of relations (R, S, T) are minimal, i.e., every atomic relation of R can be part of a solution of the triple. No solutions are lost when replacing constraints with the intersection of two other constraints or by a converse constraint. Let Θ' be the set obtained from Θ after inductively replacing all constraints over $\widehat{\mathcal{S}}_w$ with constraints over \mathcal{S} . Since potential solutions are lost in the transformation, the only problematic case is where Θ is consistent but Θ' is inconsistent. If Θ is consistent, then there must be a refinement of Θ to a consistent atomic CSP Θ_a . For each constraint xRy of Θ which is replaced, all the resulting triples are minimal and are not related to any other variable in Θ . Note that due to the inductive replacement, some constraints will be replaced by stacks of minimal triples. Therefore, each R can be replaced with any of its atomic relations without making the

resulting stacks inconsistent. Intersecting Θ' with Θ_a followed by computing a-closure will always result in an a-closed set. Since the stacks contain only minimal triples, it is clear that they can be subsequently refined to atomic relations. The relations between the fresh variables and the variables of Θ can also be refined to atomic relations as they were unrelated before applying a-closure. The resulting atomic CSP will always be a-closed, so Θ' must be consistent if a-closure decides atomic CSPs. ■

This covers all the cases in the middle column of Table 1 such as RCC8, but does not cover those cases in the bottom right cell. This result is very important for all existing and future calculi where only weak composition is used. We know now that for all calculi where a-closure decides atomic CSPs, complexity results can be transferred between a set of relations and its closure, independent of whether we are using weak composition or composition. This also resolves all doubts (Düntsch, personal communication) about applying this technique to RCC8. On the other hand, we cannot use this popular method of transferring complexity results in cases where we have only weak composition and a-closure does not decide atomic CSPs. For all existing calculi that fall into this category, we should reconsider the complexity analysis. In the following section we will have a look at the complexity results of INDU and PIDN and it turns out that some of the complexity results in the literature are wrong.

The second general technique which is very useful for analysing computational properties and identifying large tractable subsets is the refinement method [20]. It gives a simple algorithm for showing if a set $\mathcal{S} \subseteq 2^A$ can be refined to a set $\mathcal{T} \subseteq 2^A$ in the sense that for every path-consistent set of constraints Θ over \mathcal{S} and every relation $S \in \mathcal{S}$ we can always refine S to a subrelation $T \subseteq S$ with $T \in \mathcal{T}$. If path-consistency decides consistency for \mathcal{T} then it must also decide consistency for \mathcal{S} .

Theorem 3. *Let \mathcal{R} be a finite set of qualitative relations for which a-closure decides atomic CSPs. The refinement method also works for weak composition by using the a-closure algorithm instead of the path-consistency algorithm.*

Proof Sketch. Any a-closed triple of variables is minimal. So if a relation S can be refined to T in any a-closed triple that contains S , then the refinement can be made in any a-closed network without making the resulting network not a-closed. If a-closure decides the resulting network, then it also decides the original network. ■

Note that the refinement method only makes sense if a-closure decides atomic CSPs as the whole purpose of the refinement method is to transfer applicability of a-closure for deciding consistency from one subset of \mathcal{R} to another.

5 A Road Map for Analysing Qualitative Calculi

Using the results of our paper we can now analyse new and revisit existing qualitative spatial and temporal calculi. When defining a new set of atomic relations and the domain is not ordered, we have to assume that we have only weak composition unless we can prove the contrary. The most important step is to prove whether a-closure decides atomic CSPs for our new calculus. It is possible to use the heuristic given in the previous section, but if a-closure decides atomic CSPs, then this has to be proven using

the semantics of the relations. If it turns out that a-closure decides atomic CSPs then we can proceed by applying the techniques we discussed in the previous section, i.e., we can identify larger tractable subsets by using the refinement method and by computing the closure of known tractable subsets under intersection, converse and (weak) composition. But what if it does not?

5.1 When A-Closure Does Not Decide Atomic CSPs

This is the case for many calculi in the literature (see e.g. Table 1) and will probably be the case for many future calculi. As shown in Theorem 1 this means that it is possible to enforce non-overlapping subatomic relations. If we only get finitely many non-overlapping subatomic relations, as, e.g., for the containment algebra, then it is best to study the calculus obtained by the finitely many new atomic relations and treat the original calculus as a subcalculus of the new calculus. If we do get infinitely many non-overlapping subatomic relations, however, then we suggest to proceed in one of two different ways. Let us first reflect what it means to have infinitely many non-overlapping subatomic relations: An important property of a qualitative calculus is to have only finitely many distinctions. So if we have to make infinitely many distinctions, then we do not have a qualitative calculus anymore! Therefore we cannot expect that qualitative methods and techniques that are only based on (weak) compositions help us in any way. This is also the reason why we analysed the techniques in the previous section only for cases where a-closure decides atomic CSPs, i.e., where we do have qualitative calculi.¹

One way of dealing with these calculi is to acknowledge that we do not have a qualitative calculus anymore and to use algorithms that deal with quantitative calculi instead. It might be that consistency can still be decided in polynomial time using these algorithms. Another way is to find the source that makes the calculus quantitative and to eliminate this source in such a way that it has no effect anymore, e.g., by combining atomic relations to form coarser atomic relations. Both of these ways were considered for the STAR calculus [21]. A third way, which is sometimes chosen, but which we discourage everyone from taking, is to look at 4-consistency.

5.2 Problems with Using 4-Consistency

We sometimes see results in the literature of the form “4-consistency decides consistency for a set of relations $\mathcal{S} \subseteq 2^A$ and therefore \mathcal{S} is tractable.” What we have not seen so far is a proper 4-consistency algorithm. For infinite domains where we only manipulate relation symbols, a 4-consistency algorithm must be based on composition of real ternary relations. The question then is how can we show that the composition of the ternary relations is not just a weak composition. Just like computing composition for binary relations, we might have to check an infinite number of domain values. Consequently, there could be no 4-consistent configurations at all or it could be NP hard to show whether a configuration is 4-consistent. This makes these results rather

¹ It is unlikely to find a version of Theorem 2 for cases where a-closure does not decide atomic CSPs. As a heuristic, the following property could be considered: xRy can only be replaced with xSz, zTy if for all weak compositions $R_i \diamond R_j$ that contain R the intersection of all real compositions $R_i \circ R_j$ is nonempty.

useless from a practical point of view and certainly does not allow the conclusion that these sets are tractable. We illustrate this using an example from the literature where 4-consistency was wrongly used for proving that certain subsets of INDU or PIDN [17,18] are tractable.

1. 4-consistency decides consistency for $\mathcal{S} \subseteq 2^A$
2. Deciding consistency is NP-hard for $\mathcal{T} \subseteq \mathcal{S}$

The first result was proven for some subsets of INDU and PIDN [17,18]. We obtained the second result by a straightforward reduction of the NP-hard consistency problem of PDN [16] to INDU and PIDN. It is clear from this example that 4-consistency results cannot be used for proving tractability. Validity and applicability of similar results in the literature should be reconsidered as well.

6 Conclusions

We started with the well-known observation that in many cases in qualitative spatial and temporal reasoning only weak composition can be determined. This requires us to use a-closure instead of path-consistency. We thoroughly analysed the consequences of this fact and showed that the main difficulty is not whether weak composition differs from composition, but whether it is possible to generate non-overlapping subatomic relations, a property which we prove to be equivalent to whether a-closure decides atomic CSPs. Since this occurs also in cases where weak composition is equal to composition, our analysis does not only affect cases where only weak composition is known (which are most cases where the domains are not ordered) but qualitative spatial and temporal calculi in general. We also showed under which conditions some important techniques for analysing qualitative calculi can be applied and finally gave a roadmap for how qualitative calculi should be developed and analysed. As a side effect of our analysis we found that some results in the literature have to be reconsidered.

References

1. J. F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.
2. P. Balbiani and A. Osmani. A model for reasoning about topologic relations between cyclic intervals. In *Proceedings of KR-2000*, Breckenridge, Colorado, 2000.
3. P. Balbiani, J.-F. Condotta, and L. Farinas del Cerro. A tractable subclass of the block algebra: constraint propagation and preconvex relations. In *Proceedings of the 9th Portuguese Conference on Artificial Intelligence*, pages 75–89, 1999.
4. B. Bennett. Some observations and puzzles about composing spatial and temporal relations. In *Proceedings ECAI-94 Workshop on Spatial and Temporal Reasoning*, 1994.
5. B. Bennett, A. Isli, and A.G. Cohn. When does a composition table provide a complete and tractable proof procedure for a relational constraint language? In *Proceedings of the IJCAI-97 Workshop on Spatial and Temporal Reasoning*, Nagoya, Japan, 1997.
6. I. Düntsch, H. Wang, and S. McCloskey. A relation - algebraic approach to the region connection calculus. *Theoretical Computer Science*, 255(1-2):63–83, 2001.

7. E.C. Freuder. Synthesizing constraint expressions. *Communications of the ACM*, 21(11):958–965, 1992.
8. M. Grigni, D. Papadias, and C. Papadimitriou. Topological inference. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 901–906, Montreal, Canada, August 1995.
9. H. Guesgen. Spatial reasoning based on Allen’s temporal logic. Technical Report TR-89-049, ICSI, Berkeley, CA, 1989.
10. P. B. Ladkin and R. D. Maddux. On binary constraint problems. *Journal of the ACM*, 41(3):435–469, 1993.
11. S. Li and M. Ying. Region connection calculus: Its models and composition table. *Artificial Intelligence*, 145(1-2):121–146, 2003.
12. G. Ligozat and J. Renz. Qualitative calculi: a general framework. In *Proceedings of PRICAI’04*, pages 53–64, 2004.
13. G. Ligozat. When tables tell it all: Qualitative spatial and temporal reasoning based on linear orderings. In *Spatial Information Theory: Foundations of Geographic Information Science, International Conference, COSIT 2001*, pages 60–75, 2001.
14. A. K. Mackworth. Consistency in networks of relations. *Artificial Intelligence*, 8:99–118, 1977.
15. U. Montanari. Networks of constraints: Fundamental properties and applications to picture processing. *Information Sciences*, 7:95–132, 1974.
16. I. Navarrete, A. Sattar, R. Wetprasit, and R. Marin. On point-duration networks for temporal reasoning. *Artificial Intelligence*, 140(1-2):39–70, 2002.
17. A. K. Pujari, G. Vijaya Kumari, and A. Sattar. INDU: An Interval and Duration Network. In *Australian Joint Conference on Artificial Intelligence*, pages 291–303, 1999.
18. A. K. Pujari and A. Sattar. A new framework for reasoning about points, intervals and durations. In Dean Thomas, editor, *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI’99)*, pages 1259–1267. Morgan Kaufmann Publishers, 1999.
19. D. Randell, Z. Cui, and T. Cohn. A spatial logic based on regions and connection. In *Proceedings of KR-92*, pages 165–176, San Mateo, CA, 1992. Morgan Kaufmann.
20. J. Renz. Maximal tractable fragments of the region connection calculus: A complete analysis. In *Proceedings of IJCAI’99*, pages 448–455, 1999.
21. J. Renz and D. Mitra. Qualitative direction calculi with arbitrary granularity. In *Proceedings of PRICAI’04*, pages 65–74, 2004.
22. J. Renz and B. Nebel. On the complexity of qualitative spatial reasoning: A maximal tractable fragment of the Region Connection Calculus. *Artificial Intelligence*, 108(1-2):69–123, 1999.