Problems with Local Consistency for Qualitative Calculi

Gerard Ligozat †, Jochen Renz ‡

Abstract. Qualitative spatial and temporal reasoning problems are usually expressed in terms of constraint satisfaction problems, with determining consistency as the main reasoning problem. Because of the high complexity of determining consistency, several notions of local consistency, such as path-consistency, k-consistency and corresponding algorithms have been introduced in the constraint community and adopted for qualitative spatial and temporal reasoning. Since most of these notions of local consistency are equivalent for Allen's Interval Algebra, the first and best known calculus of this kind, it is believed by many that these notions are equivalent in general—which they are not! In this paper we discuss these various notions of consistency and give examples showing their different behaviours in qualitative reasoning. We argue that algebraic closure, which can be enforced by applying a path-consistency algorithm, is the only feasible algebraic method for deciding consistency, and give a heuristic about when algebraic closure decides consistency.

1 Qualitative spatial and temporal calculi

Qualitative spatial or temporal knowledge can be represented using constraints over a given set of relations, the main reasoning problem is deciding whether a network of spatial or temporal constraints is consistent. What distinguishes spatial and temporal CSPs from standard ones is mainly that the domains of the variables are usually infinite, as there is an infinite number of spatial or temporal entities. Solving these CSPs is in general NP-hard, since it is not possible to effectively enumerate the domains. As shown by Ladkin and Maddux [4], it is still possible to apply constraint based reasoning methods such as path-consistency for temporal (and spatial) CSPs with infinite domains, provided that constraint propagation is done over the relation labels instead of the domain values. So constraint propagation is essentially applying relational composition. This worked well for the Interval Algebra [1], the first and best known temporal algebra, which represents different topological relationships between convex intervals on a continuous time line: the usual properties of local consistency were preserved and the same methods could be applied. It is therefore believed by many that this carries over to other spatial and temporal calculi–which it does not! We will analyse in the following which properties of qualitative spatial and temporal calculi are responsible for this unexpected behaviour and try to find out what can be done about it.

1.1 Constraint networks and local consistency

We first give some background and recap some important notions.

Constraint networks Knowledge between different entities can be represent by using constraints. A binary *constraint* xRy between two variables x and y restricts the possible instantiations of x and y to the tuples 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* ($\cdot^{\smile}$) defined as $R^{\smile} = \{(a,b) | (b,a) \in R\}$ and *composition* ($\circ$) of two relations R and S which is the relation $R \circ S = \{(a,b) \mid \exists c : (a,c) \in R \text{ and } (c,b) \in S\}$.

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 [7] developed a form of local consistency which Mackworth [6] 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.

We emphasise at this point that these are actually two different definitions of path-consistency: one involves the formal definition using instantiations of variables; the second one defines path-consistency as the result of the path-consistency algorithm. Since both definitions are equivalent under the above defined preconditions, they are often used interchangeably in the literature. Sometimes the first definition is called 3-consistency, while the second definition is called path-consistency. In this paper, we use the term *algebraically closed*, or *a-closed*, for CSP's satisfying the second definition.

k-consistency Freuder [3] 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.

2 Local consistency for qualitative calculi

The definition of k-consistency relies upon properties that must hold for all values of the domains involved. For CSPs with finite domains these properties can be verified, but for qualitative calculi over indefinite domains, this definition seems to be overly strong and it seems impossible to verify the necessary properties in all cases. For 3-consistency, this does not seem to be a problem as it is directly related

[1] †LIMSI-CNRS, Paris-Sud University, F-91403 Orsay, France,
‡National ICT Australia, UNSW Sydney, 2052, NSW, Australia

to compositions of relations. And, as shown by the best known qualitative temporal and spatial calculi, the Interval Algebra and RCC8 [8] the path-consistency algorithm decides consistency for these calculi, so the definitions seem to work also for qualitative calculi over infinite domains. In the following we will show that this assumption is wrong in general and show what we can do about it.

2.1 Encounter with the unexpected

As first pointed out by Düntsch [2], the composition table of RCC8 does not reflect the real composition according to the definition of composition, but only corresponds to the weak composition $(R \diamond S)$ of R and S which is defined as the union of all atomic relations intersecting $R \circ S$. Therefore, the path-consistency algorithm operates on weak composition instead of composition, and consequently the equivalence between algebraic closure (the result of the path-consistency algorithm) and 3-consistency does not hold anymore. This is shown in the following example which is an algebraically closed but not 3-consistent atomic CSP over RCC8: $B\{TPP\}A$, $B\{EC\}C$, $C\{TPP\}A$. If A is instantiated as a region with two pieces and B completely fills one piece, then C cannot be instantiated. Proving whether the specified composition is actually a real composition can be extremely difficult if not impossible when having infinite domains and often becomes apparent only if a counterexample can be found. It also depends on the domain which is used. For the interval algebra, for example, algebraic closure is equivalent to 3-consistency if a rational domain is used, while it is not equivalent if integer domains are used.

Where do we go from here? In most cases, qualitative spatial and temporal calculi are based on weak composition only [5]. But if we have only weak composition, then the result of applying a path-consistency algorithm will not necessarily be a path-consistent/3-consistent CSP. Does this mean that all local consistency methods are useless for qualitative spatial and temporal calculi, since we cannot be sure in many cases whether we are dealing with composition or only with weak composition?

2.2 Some answers and possible solutions

Our answer to this question is a clear "no"! What is important in qualitative spatial and temporal reasoning is mainly to have a (efficient) procedure for deciding whether a given instance is consistent or not. It is only of secondary interest whether an instance is or can be made k-consistent. In our opinion k-consistency is much too strong, as it requires that some properties hold for all possible instantiations (of an infinite domain!), while for consistency only one instantiation is necessary. We also want to stress that for the same reason it is not bad if only weak composition is used in a qualitative calculus.

So how can we make use of local consistency methods when maybe only weak composition is known? As we can see from the RCC8 case, the path-consistency algorithm is a very powerful local consistency method. Contrary to what its name suggests, it does not necessarily make a given CSP over an infinite domain path-consistent, but it always computes its algebraic closure (it should be called the *algebraic closure algorithm* instead). A non-empty algebraic closure is often equivalent to a CSP being consistent. Actually the algebraic closure method is the strongest available purely algebraic method, i.e., only based on relational operators, so in the following we will analyse when the path-consistency algorithm decides consistency and when it does not. In general this has to be proven for each calculus anew, so the question is whether it is possible to find some indications by which we can tell for a given calculus whether the path-consistency algorithm is likely to decide consistency or not. We analysed atomic CSPs over different spatial and temporal calculi in order to see if we can derive some general properties.

First we looked at those calculi where the path-consistency algorithm does not decide consistency, among which are the INDU calculus, the cyclic interval calculus, the Star calculus, the pentagonal algebra and the containment algebra. For the Star calculus, e.g., it is possible to refine atomic relations to subatomic relations by specifying a particular arrangement of constraints, independent of the instantiation of the variables [9]. Using a different arrangement, it is possible to refine the same atomic relation to a different subatomic relation. It is clear that when we combine these two arrangements such that the constraint is refined to different subatomic relations, then the CSP is inconsistent, which cannot be detected by the path-consistency algorithm. From this example we can derive a conjecture which we will test with the other calculi.

Conjecture 1 *The path-consistency algorithm decides consistency for a calculus $\mathcal{A}$ if and only if it is not possible to refine an atomic relation to subatomic relations whose intersection is empty.*

For all the calculi we studied where the path-consistency algorithm does not decide consistency, we found an arrangement of constraints by which an atomic relation can be refined to sub-atomic relations that do not intersect. It can sometimes be quite hard to find such an arrangement. As a simple heuristic one should try to see if the domain enables more distinctions than those made by the atomic relations and if any of these finer distinctions can be enforced by an arrangement of constraints over atomic relations. For those calculi where the path-consistency algorithm decides consistency, we could not find a way of refining atomic relations. This shows us an important point: How can we prove that there is no arrangement of constraints such that an atomic relation can be refined? This might be as hard as showing that path-consistency decides consistency.

3 Conclusions

We showed in this paper that local consistency for qualitative CSPs behaves in quite different ways than in the classical finite case, and tried to clarify the confusion which apparently exists in the spatial and temporal reasoning community regarding different notions of local consistency. In many cases the various notions are not equivalent, in particular k-consistency can often not be enforced by k-consistency algorithms. For the algebras we usually consider in qualitative spatial and temporal reasoning and for their infinite domains we often only have weak compositions. However, since we are mainly interested in determining consistency, the lack of real composition is not a big loss. Instead it is more important to find out when the path-consistency algorithm, which is actually the strongest general and purely algebraic method we have at hand, decides consistency. By analysing several spatial and temporal calculi we have identified a property which seems to be very helpful in answering this question.

REFERENCES

[1] J. F. Allen. Maintaining knowledge about temporal intervals. *Comm. of the ACM*, 26(11):832–843, 1983.

[2] I. Düntsch, H. Wang, and S. McCloskey. A relation algebraic approach to the Region Connection Calculus. *TCS*, 255:63-83, 2001.

[3] E.C. Freuder. Synthesizing constraint expressions. *Communications of the ACM*, 21(11):958–965, 1992.

[4] P. B. Ladkin and R. D. Maddux. On binary constraint problems. *J.ACM*, 41(3):435–469, 1994.

[5] G. Ligozat and J. Renz. What is a qualitative calculus? A general framework. In *Proc. PRICAI'04*, 2004.

[6] A.K. Mackworth. Consistency in networks of relations. *AIJ*, 8:99–118, 1977.

[7] U. Montanari. Networks of constraints: Fundamental properties and applications to picture processing. *Information Sciences*, 7:95–132, 1974.

[8] D.A. Randell, Z. Cui, and A.G. Cohn. A Spatial Logic Based on Regions and Connection. In *Proc. KR-92*, pages 165–176, 1992.

[9] J. Renz and D. Mitra. Qualitative Direction Calculi with arbitrary granularity. In *Proc. PRICAI'04*, 2004.