

A Fast Algorithm for Consistency Checking Partially Ordered Time

Leif Eriksson and Victor Lagerkvist*

Department of Computer and Information Science, Linköping University, Linköping, Sweden

{leif.eriksson, victor.lagerkvist}@liu.se

Abstract

Partially ordered models of time occur naturally in applications where agents/processes cannot perfectly communicate with each other, and can be traced back to the seminal work of Lamport. In this paper we consider the problem of *deciding* if a (likely incomplete) description of a system of events is consistent, the *network consistency problem* for the *point algebra of partially ordered time* (POT). While the classical complexity of this problem has been fully settled, comparably little is known of the *fine-grained* complexity of POT except that it can be solved in $\mathcal{O}^*((0.368n)^n)$ time by enumerating *ordered partitions*. We construct a *much* faster algorithm with a run-time bounded by $\mathcal{O}^*((0.26n)^n)$. This is achieved by a sophisticated enumeration of structures similar to total orders, which are then greedily expanded towards a solution. While similar ideas have been explored earlier for related problems it turns out that the analysis for POT is non-trivial and requires significant new ideas.

1 Introduction

Qualitative reasoning is an important formalism in artificial intelligence where the objective is to reason about continuous properties given certain relations between the unknown entities. Two important subfields are *temporal reasoning*, e.g., the *point algebra for partially ordered time* (POT), *Allen’s interval algebra* (ALLEN), and the *point algebra for branching time*, and *spatial reasoning*, e.g., the *region connection calculus* (RCC), the *cardinal direction calculus*, and the *rectangle algebra*. There are numerous applications of all of these formalisms in AI, e.g., in knowledge representation [Forbus, 2019], linguistics [Allen, 1984; Song and Cohen, 1988], and planning [Allen, 1991; Allen *et al.*, 2014; Nogueira *et al.*, 1996; Song and Cohen, 1996]. For a broad overview with further applications and references we refer to the survey by Dylla *et al.* [2017].

In this paper we are interested in constructing fast (but superpolynomial) algorithms for NP-hard temporal reasoning problems, with a particular focus on the POT problem.

*Partially supported by the Swedish research council (VR) under grant 2019-03690.




Figure 1: A scenario with tree tasks T_1, T_2, T_3 where T_1 precedes T_3 , T_1 and T_2 are incomparable, and where either T_2 precedes T_3 , or T_3 precedes T_2 .

Here, the basic task is to check whether a given set of events and a set of possible relationships between them is *consistent* in the sense that there exists some partial ordering of the events which does not contradict any of the given relationships. Thus, this model of time is suitable in applications where we are working with agents/processes who cannot perfectly communicate with each other and where a global, totally ordered model of time is not possible, e.g., in distributed or concurrent systems [Anger, 1989; Anger *et al.*, 1999; Lamport, 1986]. There are many possible formulations of the basic computational problem and we consider the setup where the possible basic relations among two points in time are less than ($<$), greater than ($>$), incomparable to ($||$) and equivalent to ($=$), i.e., the basic operations of the well-known *point algebra* [Ladkin and Maddux, 1994]. To make it possible to encode complex relationship between events we follow Broxwall & Jonsson [2003] and allow disjunctions of the basic relations. Crucially, disjunctions make it possible to model incomplete relationships between tasks, e.g., we are given two tasks, where one of them started the other but we lack the knowledge of which one was first. See Figure 1 for a visualization of a constraint network with three tasks (where $\{<, >\}$ means that $<$ or $>$ is true). This network is satisfiable since the tasks can be ordered as $T_1 < T_3, T_2 < T_3$ where T_1 and T_2 are incomparable. Let us also remark that all of the aforementioned problems can be formulated as infinite-domain *constraint satisfaction problems* (CSPs) over ω -categorical *constraint languages* [Bodirsky and Jonsson, 2017]. In this framework one first fixes a set of binary *basic relations* \mathcal{B} and then consider $\text{CSP}(\mathcal{B}^{\vee=})$ where $\mathcal{B}^{\vee=} = \bigcup_{R \in \mathcal{B}} R$ is the union

of the basic relations. For example, the POT problem can then be formulated as a CSP problem where the basic relations are formed over the *random partial order* (cf. Chapter 2 in Bodirsky [2021]).

1.1 Related Work

Significant attention has been devoted to finding maximally tractable subclasses of these problems, typically accomplished by local consistency methods, and for e.g. RCC-5, RCC-8, ALLEN, and POT, all maximal, tractable classes have been identified [Dylla *et al.*, 2017]. If one extends to arbitrary first-order reducts of the basic relations the satisfiability problem of POT has a complexity dichotomy by Kompatscher & Pham [2017], while the corresponding problem for RCC is generally undecidable. But, naturally, we cannot be content with merely understanding the tractable fragments, since their expressive power is too restrictive to be able to model real-world problems. Hence, we need methods for solving NP-hard reasoning tasks as fast as possible. To expand our understanding of the NP-hard cases we would thus like to (1) construct algorithms faster than exhaustive search, and (2) prove that certain types of speedups are *not* possible, subject to stronger complexity theoretical assumptions than $P \neq NP$. Complexity questions like these, especially for a precise complexity parameter such as the number of variables, n , typically fall under the scope of *fine-grained complexity*. Thus, given a reasonable set of basic relations \mathcal{B} , how fast can we expect to solve $\text{CSP}(\mathcal{B}^{\vee=})$ (which might be NP-hard even if $\text{CSP}(\mathcal{B})$ is tractable)? Here, let us first remark that any $\text{CSP}(\mathcal{B}^{\vee=})$ problem is solvable by an exhaustive backtracking algorithm in $2^{\mathcal{O}(n^2)}$ time, under mild assumptions on the set of basic relations \mathcal{B} (e.g., that $\text{CSP}(\mathcal{B})$ is solvable in polynomial time). However, for several prominent problems in qualitative reasoning, including POT, it is possible to argue that the $2^{\mathcal{O}(n^2)}$ bound is too naive to be used as a baseline for improvement. Instead, these problems can be solved in $2^{\mathcal{O}(n \log n)}$ time by enumerating *ordered partitions* [Jonsson and Lagerkvist, 2017; Jonsson *et al.*, 2021], pushing down the running time to $\mathcal{O}^*((2n)^{2n})$ for ALLEN, $\mathcal{O}^*((0.531n)^n)$ for RCC-8 and $\mathcal{O}^*((0.368n)^n)$ for RCC-5 and POT¹. Thus, these problems can be solved by enumerating objects similar to assignments in finite-domain CSPs, and the question is then whether it is possible to solve the problem faster than exhaustively enumerating all orderings, similar to how it is a major open question whether CNF-SAT is solvable in $\mathcal{O}^*(c^n)$ time for *some* $c < 2$. This is indeed known to be possible for certain reasoning problems, e.g., ALLEN, which recently has been solved in $\mathcal{O}^*((1.0615n)^n)$ time [Eriksson and Lagerkvist, 2021], and if the problem is restricted to intervals of length one then it can even be solved in $2^{\mathcal{O}(n \log \log n)}$ time [Dabrowski *et al.*, 2020], and if no point occurs inside more than k intervals then it can be solved in $\mathcal{O}^*(k^n)$ time [Eriksson and Lagerkvist, 2022]. A faster $f(k)^n$ time algorithm, for some function f , is also known for the special case of POT where a solution with *effective width* of at most k is asked for [Eriksson and Lagerkvist, 2022]. However, despite these improvements, we are still far away from an unconditional single-exponential $\mathcal{O}^*(c^n)$ time

¹The notation $\mathcal{O}^*(\cdot)$ suppresses polynomial factors.

algorithm and even further away from the best-known lower bounds which only rule out subexponential algorithms running in $2^{o(n)}$ time under the *exponential-time hypothesis* [Jonsson *et al.*, 2021]. Hence, cutting-edge research suggests that qualitative reasoning problems in many cases admit *significantly* improved algorithms even though general single-exponential running times seem to be out of reach with existing methods.

1.2 Our Contribution

In this paper we advance this frontier by describing a novel and significantly improved algorithm for the POT problem with a running time of $\mathcal{O}^*((0.26n)^n)$, which is *much* faster than the previously known baseline of $\mathcal{O}^*((0.368n)^n)$. Hence, our algorithm is not only a showcase that an improved algorithm is possible for POT but *significantly* beats the naive upper bound based on enumerating ordered partitions. We achieve this as follows: after introducing the necessary technical background (in Section 2) we start our work on POT in Section 3. We analyze the structures of potential solutions and use greedy choices to find a structure that is suitable for enumeration and which yields a significant improvement over enumerating ordered partitions. The basic idea is to group variables into pairs and then order these pairs relative to each other instead of all variables individually, and ordering the variables in each pair can thereafter be done greedily. Hence, the basic idea is not that complex, but actually proving soundness and completeness of our approach is non-trivial and requires novel techniques. Finally, we conclude our results in Section 4 and present a discussion over how these algorithms and the ideas behind them might be open to further improvements and what other problems these ideas may be applicable to. Notably, can the algorithm be adapted to solve RCC-5 or RCC-8, and how far can we push the upper bound with this algorithmic technique?

2 Preliminaries

Given a set of finitary relations Γ defined on a (potentially infinite) set D of values, we define the constraint satisfaction problem over Γ ($\text{CSP}(\Gamma)$) as follows.

CSP(Γ)

Instance: A tuple (V, C) , where V is a set of variables and C a set of constraints of the form $R(v_1, \dots, v_t)$, where t is the arity of $R \in \Gamma$ and $v_1, \dots, v_t \in V$.

Question: Is there a function $f: V \rightarrow D$ such that $(f(v_1), \dots, f(v_t)) \in R$ for every $R(v_1, \dots, v_t) \in C$?

The set Γ is referred to as a *constraint language* and the function f is sometimes called a *satisfying assignment* of an instance, I , or simply a *model* of I . We write $\|I\|$ for the number of bits required to represent an instance I of $\text{CSP}(\Gamma)$.

Definition 1. We define the following orders:

1. A pair (S, \leq) is a partial order if \leq is reflexive ($\forall x \in S$ then $x \leq x$), antisymmetric ($\forall x, y \in S$, if $x \leq y$ and $y \leq x$ then $x = y$), and transitive (if $x \leq y$ and $y \leq z$ then $x \leq z$).

2. A pair $(S, <)$ is a strict partial order if $<$ is irreflexive, asymmetric ($\forall x, y \in S$, if $x < y$ then $y < x$ does not hold), and transitive.
3. A pair (S, \leq) is a total order if \leq is reflexive, antisymmetric, transitive and strongly connected ($\forall x, y \in S$ then $x \leq y$ or $y \leq x$).

If $\odot \in \{<, >, ||, =\}$ and $P = (S, \leq_P)$ is a partial (or total) order then we write \odot_P for the relation induced by P : $x <_P y$ if $x \leq_P y$ and $y \leq_P x$ does not hold, conversely for $>_P$, $||_P$ if neither $x \leq_P y$ nor $y \leq_P x$, and $x =_P y$ if $x \leq_P y$ and $y \leq_P x$. We now define the main problem of the paper.

PARTIALLY ORDERED TIME

Instance: A set of variables V and a set of binary constraints C where $c \subseteq \{<, >, ||, =\}$ for each $c(x, y) \in C$.

Question: Is there a partial order $P = (S, \leq)$ with $|S| \leq |V|$ and a function $f: V \rightarrow S$ such that for every constraint $c(x, y) \in C$, $f(x) \odot_P f(y)$ for some $\odot \in c$?

Alternatively one can also use the *random partial order* P , i.e., the (unique) countable partial order which is *universal* (contains an isomorphic copy of every finite partial order) and is *homogeneous* (any isomorphism between finite substructures can be extended to an automorphism of P). Then, the POT problem can equivalently well be defined as $\text{CSP}(\mathcal{RP})$ where \mathcal{RP} is the closure of $\{<_P, >_P, ||_P, =_P\}$ under union (cf. [Bodirsky, 2021]).

3 Partially Ordered Time

Our approach to beat the naive $O^*((0.368n)^n)$ algorithm for POT involves exploring a carefully selected group of partial orders. The algorithm, in particular, organizes variables into pairs where we only have to consider a relative ordering with $\frac{n}{2}$ other variables. This scheme leads to a runtime that is dominated by $n!/2^{\frac{n}{2}}$. Demonstrating the correctness of this strategy is a nontrivial task, and the analysis itself is arguably as interesting as the precise bound we attain.

3.1 Definitions

We start by introducing the concepts necessary for the soundness and completeness proofs of the main algorithm.

Definition 2. If $P = (S, \leq_P)$ and $P' = (S', \leq_{P'})$ are two partial orders then P is a *stub* of P' if $S \subseteq S'$ and $\leq_P \subseteq \leq_{P'}$.

We chose to represent our instances as (multi-)relational networks rather than as sets of constraints. This will give us more flexibility when adding additional restrictions to our instances, since it allows us to limit ourselves to sets of partial orders under some restrictions.

Definition 3. For an arbitrary POT instance $I = (V, C)$ we define two different variants of relational networks:

1. A function $f: V^2 \rightarrow \{<, >, ||, =\}$ is a relational network (over V), if for every constraint $c(x, y) \in C$ then $f(x, y) \in c(x, y)$. We also say that f is a relational network for I .
2. A function $f: V^2 \rightarrow \mathcal{P}(\{<, >, ||, =\})$ is a multi relational network if for every constraint $c(x, y) \in C$ then

$f(x, y) \subseteq c(x, y)$. We also say that f is a (multi) relational network for I .

If I is a 'yes'-instance, i.e. there exists a partial order (S, \leq_P) such that for all $c(x, y) \in C$ then $x \odot_P y$ with $\odot \in c(x, y)$, we say that f is a consistent (multi) relational network.

We will occasionally view (multi) relational networks as sets of constraints in the obvious way. For example, if f is a multi relational network of an instance $I = (V, C)$, then $f \cup \{x\{<, >\}y\}$ is equivalent to the multi relational network for the instance $(V, C \cup \{x\{<, >\}y\})$.

Definition 4. Consider the set of all multi relational networks over a fixed set V .

1. For two distinct multi relational networks f and f' over V we write $f \preceq f'$ if $f(x, y) \subseteq f'(x, y)$ for all $x, y \in V$.
2. We write $f \prec f'$ for the corresponding irreflexive order.
3. A multi relational network f is said to be maximally general if it is a maximal element in \preceq , i.e., there does not exist f' such that $f \prec f'$.

A solution for a POT instance I with relational network f can now be represented by a relational network $g \preceq f$. This will be convenient since it eliminates the need to refer to fixed values in the context of a solution.

We also need a local consistency definition for our multi relational networks.

Definition 5. A multi relational network f over V is (locally) consistent over $s \subseteq V$ if $f(x, y) = \bigcup g_i(x, y)$, $x, y \in s$ where $\{g_1, \dots, g_n\}$ is the set of all consistent relational networks $g_i \preceq f$ over s .

Next, we define the central concept of composing a partial order with a multi relational network, roughly meaning that the partial order is used to simplify the multi relational network as much as possible.

Definition 6. Given a POT instance $I = (V, C)$ with multi relational network f and a partial order $P = (V, \leq_P)$, we define the POT instance

$$P \circ f = \begin{cases} f(x, y) \setminus \{>\}, & \text{if } x <_P y, \\ f(x, y) \setminus \{<\}, & \text{if } y <_P x, \\ f(x, y) \cap \{=\}, & \text{if } y =_P x, \\ f(x, y), & \text{if } y ||_P x. \end{cases}$$

Note that $P \circ f$ may have cases where $f(x, y) = \emptyset$, meaning that there is no partial order $P' = (V, \leq_{P'})$ satisfying I for which P is a stub.

Lemma 7. Let $I = (V, C)$ be a POT instance with multi relational network f . If there exists a partial order $P = (V, \leq_P)$ such that $P \circ f$ is a 'yes'-instance then I is a 'yes'-instance.

Proof. By Definition 6, $(P \circ f) \preceq f$ and hence we have that $(P \circ f)(x, y) \subseteq c(x, y)$ for all $c(x, y) \in C$. So, if there is an assignment satisfying $P \circ f$, the same assignment must also satisfy I . \square

We are now ready to present the limited set of partial orders that we are interested in enumerating.

Definition 8. A total ordering of pairs (TOP) is a partial order $P = (V, \leq_P)$ such that for every $x \in V$ there is at most one $y \in V \setminus \{x\}$ such that $x \parallel_P y$. We say that (x, y) is a pair in P if $x \parallel_P y$.

We will use the notation $(a_1, b_1) <_P \dots <_P (a_i, b_i)$ for writing (part of) a TOP P . Here (a_j, b_j) are our pairs for $j \in \{1, \dots, i\}$. Note that we do not make a difference between (a_j, b_j) and (b_j, a_j) .

Now we are ready to start applying TOPs on multi relational networks, see what important structures can occur and, most importantly, investigate why this is beneficial and determine for which cases the output is still difficult to solve.

Definition 9. Given a TOP P and a multi relational network f of a POT instance, a link in $P \circ f$ is a non-empty sequence of pairs $(a_1, b_1) <_P \dots <_P (a_n, b_n)$ such that

1. $(P \circ f)(a_i, b_i) \setminus \{=\} = \{<, >\}$,
2. $(P \circ f)(a_i, a_j) \setminus \{=\} = \{<\}$, $1 \leq i < j \leq n$,
3. $(P \circ f)(b_i, b_j) \setminus \{=\} = \{<\}$, $1 \leq i < j \leq n$, and
4. $\parallel \in (P \circ f)(a_i, b_j)$, $i \neq j$.

Two links overlap if they share any pairs.

Note that a single link does not need to contain more than a single pair. Also, two links sharing the same pairs, e.g. (a, b) and (b, a) , are technically the same link. This is just two different representations, or *directions*, of the same link, which will be relevant for the following important concept.

Definition 10. Given a TOP P and multi relational network f for some POT instance, a chain in $P \circ f$ is two variables $x <_P y$ with $(P \circ f)(x, y) \setminus \{=\} = \{\parallel\}$ and a non-empty sequence of m links

$$(a_{1,1}, b_{1,1}) <_P \dots <_P (a_{1,n_1}, b_{1,n_1}) <_P \dots <_P (a_{m,1}, b_{m,1}) <_P \dots <_P (a_{m,n_m}, b_{m,n_m})$$

such that

1. either $x = b_{1,1}$ or $(P \circ f)(x, b_{1,1}) \subseteq \{<, =\}$,
2. either $y = a_{m,n'}$ or $(P \circ f)(a_{m,n'}, y) \subseteq \{<, =\}$,
3. $(P \circ f)(a_{i,n_i}, b_{i+1,1}) \setminus \{=\} = \{<\}$,
4. $(P \circ f)(a_{i,i'}, a_{j,j'}) \setminus \{=\} = \{<, \parallel\}$, $i < j$, $i' \in \{1, \dots, n_i\}$, $j' \in \{1, \dots, n_j\}$ and
5. $(P \circ f)(b_{i,i'}, b_{j,j'}) \setminus \{=\} = \{<, \parallel\}$, $i < j$, $i' \in \{1, \dots, n_i\}$, $j' \in \{1, \dots, n_j\}$.

We say that a chain is broken if $(P \circ f)(a_{i,i'}, b_{i,i'}) = \{<\}$ for all i' for some i . The length of a chain is the number of links m in the chain. We say that x is the head of the chain and y the tail.

While we here differ between (a, b) and (b, a) in contrast to how we did for TOPs, it should be clear by the context when it is of importance or not, i.e. when we are speaking of chains and when we are only discussing (parts of) a TOP.

Example 11. Take the TOP

$$(a_1, b_1) <_P (a_2, b_2) <_P (a_3, b_3) <_P (a_4, b_4) <_P (a_5, b_5).$$

Let $f(b_1, b_3) = \{<\}$, $f(a_2, a_3) = \{<\}$, $f(b_2, b_3) = \{<\}$, $f(a_3, b_4) = \{<\}$, $f(b_3, b_5) = \{<\}$, $f(a_1, a_4) = \{\parallel\}$, $f(a_2, a_5) = \{\parallel\}$ and for every $i \in \{1, \dots, 5\}$,




Figure 2: Graphic representation of two chains overlapping in one link (a_3, b_3) , as described in Example 11. Variables/pairs to the left are ordered before those to the right in our TOP and hence we assume either a $<$ - or a \parallel -relation going from left to right.

$f(a_i, b_i) = \{<, >\}$. We now have five links: (a_1, b_1) , $(a_2, b_2) <_P (a_3, b_3)$, (b_3, a_3) , (b_4, a_4) and (b_5, a_5) . We also have two chains: $(a_1, b_1) <_P (b_3, a_3) <_P (b_4, a_4)$ and $(a_2, b_2) <_P (a_3, b_3) <_P (b_5, a_5)$. Here $(a_2, b_2) <_P (a_3, b_3)$ and (b_3, a_3) contain the same pair in a_3, b_3 , but their directions differ, and hence these links overlap in the opposite directions.

This example is also visualized in Figure 2.

Before we prove that $T \circ f$ is solvable in polynomial time for any total order T we introduce the following comparability property of partial orders.

Definition 12. Given two relational networks f and g over the same variable set V , f is \parallel -larger than g if, for any variables $x, y \in V$:

1. if $g(x, y) = \{\parallel\}$ then $f(x, y) = \{\parallel\}$,
2. if $g(x, y) = \{<\}$ then $f(x, y) \in \{<, \parallel\}$,
3. if $g(x, y) = \{>\}$ then $f(x, y) \in \{>, \parallel\}$,
4. if $g(x, y) = \{=\}$ then $f(x, y) \in \{<, >, \parallel, =\}$,

and there exists $x, y \in V$ such that $g(x, y) \neq f(x, y)$.

Furthermore, we say that f is \parallel -maximal if there does not exist any g which is \parallel -larger than f .

3.2 The Algorithm

We begin by showing how to solve $T \circ f$ in polynomial time for a total order T .

Lemma 13. Let $I = (V, C)$ be a POT instance with multi relational network f and let $T = (V, <_T)$ be a total order. Then $T \circ f$ is solvable in polynomial time.

Proof. Solving $T \circ f$ can be done by enforcing consistency for each triple $x, y, z \in V$, e.g., if $f(x, y) = \{<\}$ and $f(y, z) = \{<\}$ then $f(x, z) = \{<\}$. We repeat this until no more changes occur. If any $f(x, y) = \emptyset$ then $T \circ f$ must be a 'no'-instance, otherwise it is a 'yes'-instance. Soundness for this approach follows naturally. To see that this approach is complete take the \parallel -maximal relational network g which can give T when topologically sorted. Since I given T is a

'yes'-instance, g must exist. After the first round of local consistency propagation the following must be true for all triples $x <_T y <_T z$ for our f :

1. if $= \in f(x, z)$ then $= \in f(x, y)$ and $= \in f(y, z)$, or $|| \in f(x, y)$ and $|| \in f(y, z)$,
2. if $< \in f(x, z)$ then $f(x, y) = f(y, z)$ cannot equal $=$, and
3. if $|| \in f(x, z)$ then $|| \in f(x, y)$ or $|| = f(y, z)$.

Similarly, $g(x, y) \in f(x, y)$. By then removing $=$ -relations from any $|f(x, y)| > 1$ the relations $g(x, y) \in f(x, y)$ must still be true and hence the network is still a 'yes'-instance. Again, we enforce consistency for triples until no changes occur. For any triple $x <_T y <_T z$, where $g(x, y)$, $g(x, z)$ and $g(y, z)$ are not equal to $\{=\}$, it then still holds that if $|| \in f(x, z)$ then $|| \in f(x, y)$ or $|| \in f(y, z)$. Hence, for every $|| \in f(x, y)$ we can assume $f(x, y) = \{||\}$ and still have a 'yes'-instance. In fact, $f = g$ since g is $||$ -maximal. Hence, if I is a 'yes'-instance and if there is any relational network for I that topologically sorts to T , we will answer 'yes'.

For the complexity, we only need to propagate complexity for variable triples, and each propagation cycle must remove at least one relation from f , else no change is made. As there are at most four possible atomic relations between each variable pair, the number of cycles is polynomially bounded by $O(|V|^2)$. Hence, this approach can be done in polynomial time. \square

Why it is actually interesting to solve $T \circ f$ in polynomial time is shown by the following lemma.

Lemma 14. *If a POT instance I with multi relational network f is a 'yes'-instance then there is a total order T such that $T \circ f$ is a 'yes'-instance.*

Proof. We remind the reader that we can represent a solution for I with a relation network. Take any relational network g satisfying I and the partial ordering P described by this relational network. Topologically sort P to a total order T . From the definition of $T \circ f$ we know that $T \circ f(x, y) = f(x, y) \setminus \{>\}$ if $x <_T y$. As $x \not<_T y$ if $g(x, y) = \{>\}$, g must be consistent with $T \circ f$ and hence g is also satisfies $T \circ f$, proving the lemma. \square

As a sanity check for making sure that an instance does not have small subinstances that are 'no'-instances, i.e. the instance is locally consistent, we prove the following lemma. Recall that $||f||$ is the number of bits needed to represent f .

Lemma 15. *For any multi relational network f , integer $k > 0$ and a set of variables $s \subseteq V$ with $|s| \leq k$, a maximally general multi relational network $f' \preceq f$ which is locally consistent with s can be computed in $h(k) \cdot ||f||^{O(1)}$ time for some computable function $h(k)$.*

Proof. Let $\{g_1, \dots, g_m\}$ be the set of all consistent relational networks $g \preceq f$ over s . For each s , there are roughly $16^{|s|^2}$ potential g_i s as we have four different relations that can either be allowed, or not allowed, between each pair of variables in s . Each potential g_i can be compared to f to check for consistency in $||f||^{O(1)}$ time as $||g_i|| \leq ||f||$. Now, for each

pair $x, y \in s$ let $f'(x, y) = \bigcup g_i(x, y)$ and for all pairs where either $x \notin s$ or $y \notin s$ let $f'(x, y) = f(x, y)$. Since $|s| \leq k$ the number of possible $16^{|s|^2}$ and the complexity of constructing f' is bounded by $h(k)$ for some computable function $h: \mathbb{N} \rightarrow \mathbb{N}$. Clearly, $f' \preceq f$, and since we enumerated all possible relations between all pairs in s , f' must also be the largest such multi relational network. \square

Now, we are ready to introduce the first of the two main gadgets necessary for our later proofs. For a partial order $P = (V, \leq)$ and distinct $x, y \in V$ we write $P^{x < y}$ for $(V, \leq \cup \{(x, y)\})$. While this operation may technically produce structures that are not partial orders (e.g., $x <_P y <_P x$), we in the forthcoming definitions will only use it on pairs in TOPs where the result is always guaranteed to be a partial order.

Definition 16. *For a TOP $P = (V, \leq)$ and a multi relational network f over V we describe a partial function R_{par} defined according to the following rules.*

1. $R_{\text{par}}(P, f) = (P, f)$ if P is a total order,
2. $R_{\text{par}}(P, f) = R(P^{x < y}, f)$ if there exists distinct $x, y \in V$ with $x ||_P y$ and $< \notin (P \circ f)(x, y)$,
3. $R_{\text{par}}(P, f) = R(P^{a_i < b_i})$ if there exists a link $(a_1, b_1) < \dots < (a_k, b_k)$ such that there is no other link $(c_1, d_1) < \dots < (c_{k'}, d_{k'})$ for which $a_i = d_{i'}$ and $b_i = c_{i'}$ for any i and i' ,
4. $R_{\text{par}}(P, f) = R_{\text{par}}(P, f')$ if there exists $a, b, c, d \in V$ and a maximally general multi relational network $f' \neq f$, $f' \preceq f$ and locally consistent with $\{a, b, c, d\}$.

It is easy to see that if $R_{\text{par}}(P, f) = (T, f)$ is defined then it returns a total order T such that P is a stub of T . We then consider the following extension of R_{par} which is guaranteed to be totally defined due to the second rule.

Definition 17. *For a TOP $P = (V, \leq)$ and a multi relational network f over V we describe a total function R_{tot} defined according to the following rules.*

1. $R_{\text{tot}}(P, f) = R_{\text{par}}(P, f)$ if $R_{\text{par}}(P, f)$ is defined, and
2. $R_{\text{tot}}(P, f) = R_{\text{tot}}(P^{x < y}, f)$ if there exists $x, y \in V$ where $\{<, >\} \subseteq (P \circ f)(x, y)$.

Note that R_{tot} will always return a total order, since otherwise the second rule could be applied. However, the function is not sound and can given a 'yes'-instance $P \circ f$ return a 'no'-instance $T \circ f$. Even so, this function will be the one we later use for solving POT instances.

As a second extension of R_{par} , we introduce the total function R_{corr} . This function will computationally be more expensive than R_{tot} , but with the trade-off that it in the second step makes sure extending the given TOP with $x < y$ is actually a reasonable choice that does not get the function stuck with a 'no'-instance. Thus, if extending the partial order by $x < y$ would yield a 'no'-instance, R_{corr} (in contrast to R_{tot}) simply stops and returns (P, f) .

Definition 18. *For a TOP $P = (V, \leq)$ and a multi relational network f over V we describe a total function R_{corr} defined according to the following rules.*

1. $R_{\text{corr}}(P, f) = R_{\text{par}}(P, f)$ if $R_{\text{par}}(P, f)$ is defined.
2. $R_{\text{corr}}(P, f) = R_{\text{corr}}(P^{x < y}, f)$ if there exists $x, y \in V$ with $\{<, >\} \subseteq (P \circ f)(x, y)$ and where $P \circ (f \cup \{x \{<, ||, =\} y\})$ is a 'yes'-instances, and
3. $R_{\text{corr}}((P, f)) = (P, f)$ otherwise.

Note that R_{corr} differs from R_{tot} in the sense that if R_{corr} (via the second rule) is given a 'yes'-instance it also guarantees that the output is a 'yes'-instance. The precise relation between R_{tot} and R_{corr} will be of much interest to us. In fact, we will show that if f is a 'yes'-instance, then there must exist a TOP such that $R_{\text{tot}} = R_{\text{corr}}$. We begin with the following lemma.

Lemma 19. *For an arbitrary POT instance with multi relational network f and some arbitrary TOP $P = (V, \leq_P)$, $R_{\text{tot}}(P, f)$ can be computed in $\text{poly}(|I|)$ time and space.*

Proof. Step 1 and 2 of R_{par} and step 2 of R_{tot} are quite clearly polynomial. Since there are $|V|^4$ sets of four variables, step 4 of R_{par} is polynomial by Lemma 15. Last, step 3 of R_{par} is polynomial since there are at most $|V|/2$ pairs, and finding if they are part of some chain (and in which directions) can be done in polynomial time. Hence, all steps of R_{tot} are doable in polynomial time. \square

Using $R_{\text{tot}}(P, f) = (T, f')$ we are now left with $T \circ f'$, which we can solve in polynomial time according to Lemma 13, and if this approach returns *true* then I must be a 'yes'-instance (via Lemma 7).

Returning to the relationship between R_{tot} and R_{corr} we are now ready to describe the property needed for $R_{\text{tot}}(P, f) \neq R_{\text{corr}}(P, f)$ to occur.

Lemma 20. *Let $P = (V, \leq_P)$ be a TOP and f a multi relational network over V such that $P \circ f$ is a 'yes'-instance. Then $R_{\text{tot}}(P, f) \neq R_{\text{corr}}(P, f)$ only if $P \circ f$ contains a non-empty set S of chains of length at least two and such that every link of every chain in the set overlaps with some other link in the opposite direction.*

Proof. Assume $R_{\text{tot}} \neq R_{\text{corr}}$ and $S = \emptyset$. The only step where R_{tot} and R_{corr} differ in output is in step 2. Further, assume we have a topological sorting T of some solution for $P \circ f$. Whenever R_{tot} reaches step 2 and for a pair $x ||_P y$ chooses $x \{<, =, ||\} y$, without knowing if this is a good choice or not, we have four different cases that could theoretically occur and are worth considering.

1. If $x = y$ or $x ||_P y$ in our solution for $P \circ f$, then $x <_T y$ and $y <_T x$ are both valid and will yield the same solution.
2. If $x < y$ in our solution, but $y < x$ in some other, but otherwise identical solution, then $x <_T y$ and $y <_T x$ are both valid and yield 'yes'-instances.
3. If $x < y$ in our solution and there is no solution with $y < x$ but which is otherwise identical, then there must be two variables u, v such that $x \leq u, v \leq y$ and $u ||_P v$. This describes a link and a chain. Then, either this chain has length one, in which point Step 4 of R_{par} will be applicable as local consistency over $\{x, y, u, v\}$ would not allow the case of $y < x$, or this chain has length longer than two, but contains a link not overlapping with

any other link in the opposite direction. Since there is here a link that does not overlap with any other in the opposite direction, step 3 of R_{par} would be applicable and we would set $x < y$ or $y < x$ depending on which one breaks the chain.

4. There are chains of length two or more, but no chain contains any link that does not overlap with some other link in the opposite direction. This matches the definition of chains in S , and hence S is non-empty.

In the first two of these cases R_{tot} and R_{corr} behave identically, and hence yield the same output. For the third one, neither R_{tot} nor R_{corr} reaches their respective step 2, and hence they behave identically. For the fourth case, S is non-empty. So neither of these four cases satisfies our assumption. But our four cases are exhaustive: they cover all relations between x and y in solutions for $T \circ f$ and all cases for all these relations. So the initial assumptions must be false, meaning that either $R_{\text{tot}} = R_{\text{corr}}$ or $S \neq \emptyset$, completing the proof. \square

Before we introduce the final piece of the puzzle, we give two short definitions: one for a sub-class of TOPs that are easier to enumerate and one for our notation for $R_{\text{tot}}(P, f) = R_{\text{corr}}(P, f)$.

Definition 21. *A proper total ordering of pairs (PTOP) is a TOP $P = (V, \leq_P)$ such that for every $x \in V$ either there is exactly one $y \in V \setminus \{x\}$ such that $x ||_P y$ or $y <_P x$ for all $y \in V \setminus \{x\}$.*

Definition 22. *If for some arbitrary TOP $P = (V, \leq_P)$ and relational network f over V , $R_{\text{tot}}(P, f) = R_{\text{corr}}(P, f)$ we say that P is reducible for f .*

And now, finally, we can show, and prove, why R_{tot} , R_{corr} , reducibility and (P)TOPs are interesting for multi relational networks for POT instances.

Lemma 23. *For every 'yes'-instance I of POT with multi relational network f , there is a PTOPT reducible for f .*

Proof. Since I is a 'yes'-instance we have a relational network g satisfying I , and by Lemma 14 we know there is a total order T such that $T \circ f$ is also satisfied by g .

We start the process of finding a TOP reducible for f by generating (from T) an arbitrary PTOPT P such that P is a stub of T . If P is reducible for f , then return P . Otherwise if P is not reducible for f we assume that $R_{\text{corr}}(P, f) = (P, f)$. Since R_{corr} can remove chains of length one by local consistency over quads, we know by Lemma 20 that $P \circ f$ contains chains longer than two, and such that their links overlap with others in the opposite direction. I.e., we have a link

$$(a_{1,1}, b_{1,1}) <_P \dots <_P (a_{1,n}, b_{1,n})$$

and a second link

$$(u_{1,1}, v_{1,1}) <_P \dots <_P (u_{1,m}, v_{1,m})$$

such that there is some $a_{1,i} = v_{2,i'}$ and $b_{1,i} = u_{2,i'}$. Furthermore, we also have a link

$$(a_{2,1}, b_{2,1}) <_P \dots <_P (a_{2,n'}, b_{2,n'})$$

that is part of the same chain as

$$(a_{1,1}, b_{1,1}) <_P \dots <_P (a_{1,n}, b_{1,n})$$

and some link

$$(u_{2,1}, v_{2,1}) <_P \dots <_P (u_{2,m'}, v_{2,m'})$$

that is part of the same chain as

$$(u_{1,1}, v_{1,1}) <_P \dots <_P (u_{1,m}, v_{1,m}).$$

From all such overlaps, take the overlapping pair with the lowest index in P and call it (x, y) . Without loss of generality, assume that $g(x, y) = <$ in our relational network g satisfying both I and $T \circ f$. Take the variable $y <_T z$ such that $g(x, z) = g(z, y) = ||$, and such that there is no other variable $z' <_T z$ fulfilling the same conditions. To see that such a variable z actually exists, take the tail x' and y' of the chains overlapping in opposite directions in (x, y) . Since the chain which includes x and x' must be broken, and similarly for y and y' , we must have that $g(x, x') = ||$ and $g(y, y') = ||$. Assume that $g(y, x') = <$. But since $g(x, y) = g(y, x') = <$ then $g(x, x') = <$ while we assumed $g(x, x') = ||$, so we have a contradiction. Hence, we have $g(x, x') = g(x', y) = ||$ and so at least one variable with this property exists and so does our z . In T there is now a sequence $x <_T y <_T x_1 <_T \dots <_T x_j <_T z$. Since z is incomparable to both x and y in g and as z is the least indexed such variable in this sequence, then for all x_i in the sequence we have $g(z, x_i) = ||$. Hence, if we construct a new total order T' such that $x <_{T'} z <_{T'} y <_{T'} x_1 <_{T'} \dots <_{T'} x_j$ but which is otherwise identical to T , then g also satisfies $T' \circ f$. If we now also construct a new arbitrary TOP P' from T' our two original links are no longer links for any chains. We repeat this process until the resulting TOP P' is reducible for f .

The question is now if the above construction of a reducible PTOP for f halts. To answer this we make the following observation: for each recursion we choose some index i in P such that the pair at i in $P \circ f$ is changed from a link to a non-link in $P' \circ f$, without making any new links out of pairs of index less than i that have previously been chosen. Hence the number of pairs at index i or less that cannot become links by choosing a new index $j > i$ has increased by one. Call the pairs at these indexes selected *protected*. Note that a pair at index i loses its protected status and can become a link again if $j < i$ is chosen, but then the same logic applies to j instead: the number of protected pairs at index j or less has increased. So, in each iteration, either i decreases or the number of protected pairs indexes less than i increases. Hence, the function must reach a point where no pair can be selected and P is reducible for f . \square

One may ask what happens to chains in $P \circ f$ that are satisfied by equality in the solution and are hence technically never broken. For all of the links in these chains, the local consistency check in Step 4 of R_{par} will keep the equality relation, while Step 2 in R_{tot} will always produce a 'yes'-instance given a 'yes'-instance when working on the pairs in these links. Hence, these are a non-issue and will be satisfied by our greedy approach (as long the variables are equal in some solution).

We now have everything we need to present the main result for this section:

Theorem 24. *Any arbitrary POT instance $I = (V, C)$ with $|V| = n$ is solvable in $\mathcal{O}^*(n!/2^{\frac{n}{2}}) \subseteq \mathcal{O}^*((0.2601n)^n)$ time and $\text{poly}(|I|)$ space.*

Proof. By Lemma 7 we know that if we find any partial order P such that $P \circ f$ is a 'yes'-instance, then I must be a 'yes'-instance. From Lemma 23 we know there exists a PTOP P such that $P \circ f$ is a 'yes'-instance if I is and such that we can solve $P \circ f$ in polynomial time by applying Lemma 13 to the output of $R_{\text{tot}}(P, f)$, and by Lemma 19 $R_{\text{tot}}(P, f)$ can be calculated in polynomial time. There are $n!/2^{\frac{n}{2}}$ such PTOPs, and we can, in a polynomial factor on n , enumerate over all PTOPs over V until we find our P . If no 'yes'-instance $P \circ f$ is found in this way then I must be a 'no'-instance, again by Lemma 23, and we can safely answer 'no'. Via Stirling's approximation of $n!$ we obtain $\mathcal{O}^*((0.2601n)^n)$. \square

With Theorem 24 in place, we have now seen the entire chain of reasoning of how enumerating PTOPs is enough to solve any POT instance. While providing a new state of the art.

4 Conclusion and Discussion

In this paper we used structural properties on potential solutions for POT instances to achieve a new and significantly improved upper bound. These results, and the techniques used, raise new questions.

A major open question is whether our algorithm for POT can be adapted to other qualitative reasoning tasks such as RCC-5 or RCC-8. The issue here seems to be that the constraints *partially overlapping* and *disjoint* are harder to handle than just incomparability. Combined, however, these relations behaves identical to incomparability, so the idea is by no means far-fetched. Partially overlapping enforces a form of upwards transitivity in that every set containing a set A must also be at least partially overlapping with some set B if A and B are partially overlapping. Similarly, disjoint enforces a form of downwards transitivity in that every set contained in C must be disjoint to D if C is disjoint to D . Effectively this introduces some different forms of chains than for POT that require some novel ideas to handle. As a first step, finding a method to handle these chains faster than 2^n for n links would be enough to improve the current state-of-the-art.

It is also natural to ask if it is possible to avoid enumerating certain orderings of pairs and thus push down the runtime even further. For example, can we find certain orderings of pairs that will never satisfy our instances? And can we do this fast enough, and often enough, that it yields a significant improvement to the overall runtime? Another promising idea is to partition variables into triples instead of pairs, or quads that are later further restricted to pairs in an intelligent manner. This requires non-trivial changes to our algorithm but does not seem impossible. If this approach can be pushed further, it would likely hit a limit of $\mathcal{O}^*(\sqrt{cn}^n)$ where the local consistency checks would start bottlenecking. Even so, this would be the first $o(n)^n$ complexity result for POT and a major advance in solving NP-hard qualitative reasoning problems.

Acknowledgements

We thank the anonymous reviewers for several insightful comments. The first author is partially supported by the *National Graduate School in Computer Science* (CUGS), Sweden. The second author is partially supported by the Swedish Research Council (VR) under grant 2019-03690.

References

- [Allen *et al.*, 2014] James F. Allen, Henry A. Kautz, Richard N. Pelavin, and Josh D. Tenenber. *Reasoning About Plans*. Elsevier Science, 2014.
- [Allen, 1984] James F. Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23(2):123–154, 1984.
- [Allen, 1991] James F. Allen. Planning as temporal reasoning. In James F. Allen, Richard Fikes, and Erik Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR-1991)*. Cambridge, MA, USA, April 22-25, 1991, pages 3–14. Morgan Kaufmann, 1991.
- [Anger *et al.*, 1999] Frank D. Anger, Debasis Mitra, and Rita V. Rodríguez. Satisfiability in nonlinear time: Algorithms and complexity. In *Proceedings of the 12th International Florida Artificial Intelligence Research Society Conference (FLAIRS-1999)*, pages 406–411, 1999.
- [Anger, 1989] Frank D. Anger. On lamport’s interprocessor communication model. *ACM Transactions on Programming Languages and Systems*, 11(3):404–417, 7 1989.
- [Bodirsky and Jonsson, 2017] Manuel Bodirsky and Peter Jonsson. A model-theoretic view on qualitative constraint reasoning. *Journal of Artificial Intelligence Research (JAIR)*, 58:339–385, 2017.
- [Bodirsky, 2021] Manuel Bodirsky. *Complexity of Infinite-Domain Constraint Satisfaction*. Cambridge University Press, 2021. Preprint available from <https://www.math.tu-dresden.de/~bodirsky/Book.pdf>.
- [Broxvall and Jonsson, 2003] Mathias Broxvall and Peter Jonsson. Point algebras for temporal reasoning: Algorithms and complexity. *Artificial Intelligence*, 149(2):179–220, 2003.
- [Dabrowski *et al.*, 2020] Konrad K. Dabrowski, Peter Jonsson, Sebastian Ordyniak, and George Osipov. Fine-grained complexity of temporal problems. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR-2020)*, pages 284–293, 2020.
- [Dylla *et al.*, 2017] Frank Dylla, Jae H. Lee, Till Mossakowski, Thomas Schneider, André V. Delden, Jasper V. D. Ven, and Diedrich Wolter. A survey of qualitative spatial and temporal calculi: Algebraic and computational properties. *ACM Computing Surveys (CSUR)*, 50(1):7:1–7:39, April 2017.
- [Eriksson and Lagerkvist, 2021] Leif Eriksson and Victor Lagerkvist. Improved algorithms for Allen’s interval algebra: a dynamic programming approach. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI-2021)*, pages 1873–1879, 2021.
- [Eriksson and Lagerkvist, 2022] Leif Eriksson and Victor Lagerkvist. A multivariate complexity analysis of qualitative reasoning problems. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 1804–1810. International Joint Conferences on Artificial Intelligence Organization, 7 2022. Main Track.
- [Forbus, 2019] K.D. Forbus. *Qualitative Representations: How People Reason and Learn about the Continuous World*. MIT Press, 2019.
- [Jonsson and Lagerkvist, 2017] Peter Jonsson and Victor Lagerkvist. An initial study of time complexity in infinite-domain constraint satisfaction. *Artificial Intelligence*, 245:115–133, 2017.
- [Jonsson *et al.*, 2021] Peter Jonsson, Victor Lagerkvist, and George Osipov. Acyclic orders, partition schemes and cpsps: Unified hardness proofs and improved algorithms. *Artificial Intelligence*, 296:103505, 2021.
- [Kompatscher and Pham, 2017] Michael Kompatscher and Van T. Pham. A complexity dichotomy for poset constraint satisfaction. In *Proceedings of the 34th Symposium on Theoretical Aspects of Computer Science, (STACS-2017)*, volume 66 of *LIPICs*, pages 47:1–47:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [Ladkin and Maddux, 1994] Peter B. Ladkin and Roger D. Maddux. On binary constraint problems. *Journal of the ACM*, 41(3):435–469, 1994.
- [Lamport, 1986] Leslie Lamport. The mutual exclusion problem: Part i—a theory of interprocess communication. *J. ACM*, 33(2):313–326, 4 1986.
- [Nogueira *et al.*, 1996] José Helano Matos Nogueira, Antonio L. Furtado, and José de Jesús Pérez Alcázar. A hybrid formal theory of plan recognition and its implementation. In *Proceedings of the 13th Brazilian Symposium on Artificial Intelligence, (SBIA-1996)*, volume 1159 of *Lecture Notes in Computer Science*, pages 31–40. Springer, 1996.
- [Song and Cohen, 1988] Fei Song and Robin Cohen. The interpretation of temporal relations in narrative. In *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI-1998)*, pages 745–750, 01 1988.
- [Song and Cohen, 1996] Fei Song and Robin Cohen. A strengthened algorithm for temporal reasoning about plans. *Computational Intelligence*, 12:331–356, 1996.