

1 Complexity of Minimum-Size Arc-Inconsistency 2 Explanations

3 **Anonymous author**

4 Anonymous affiliation

5 — Abstract —

6 Explaining the outcome of programs has become one of the main concerns in AI research. In
7 constraint programming, a user may want the system to explain why a given variable assignment is
8 not feasible or how it came to the conclusion that the problem does not have any solution. One
9 solution to the latter is to return to the user a sequence of simple reasoning steps that lead to
10 inconsistency. Arc consistency is a well-known form of reasoning that can be understood by a
11 human. We consider explanations as sequences of propagation steps of a constraint on a variable
12 (i.e. the ubiquitous revise function in arc consistency algorithms) that lead to inconsistency. We
13 characterize, on binary CSPs, cases for which providing a shortest such explanation is easy: when
14 domains are Boolean or when variables have maximum degree two. However, these polynomial cases
15 are tight. Providing a shortest explanation is NP-hard if the maximum degree is three, even if the
16 number of variables is bounded, or if domain size is bounded by three. It remains NP-hard on trees,
17 despite the fact that arc consistency is a decision procedure on trees. Finally, the problem is not
18 FPT-approximable unless the Gap-ETH conjecture is false.

19 **2012 ACM Subject Classification** Theory of computation → Constraint and logic programming;
20 Theory of computation → Fixed parameter tractability; Theory of computation → Problems,
21 reductions and completeness

22 **Keywords and phrases** Constraint programming, constraint propagation, minimum explanations,
23 complexity

24 **Digital Object Identifier** 10.4230/LIPIcs.CP.2022.

25 **Acknowledgements** Anonymous acknowledgements

26 **1** Introduction

27 Constraint Programming (CP) is a technology that allows the user to solve combinatorial
28 problems formulated as constraint networks. A constraint network is characterized by a set of
29 variables taking values in a finite domain that are subject to constraints. Constraints restrict
30 the combinations of values that specified subsets of variables can take. One of the advantage
31 of using CP is that in general constraint networks represent the problem to solve much more
32 compactly than would an integer linear program or a SAT formula. CP formulations are
33 not only compact but also easy to understand for the user thanks to the expressiveness
34 of constraints that allow to remain close to the original problem. However, nowadays, AI
35 becomes even more demanding in terms of *explainability*. A user may want to not only
36 understand the formulation of their problem as a constraint network but also to be provided
37 with explanations of why this assignment is the only solution, why that value is not feasible,
38 or why the problem does not have any solution.

39 An *abductive explanation* for a proposition is often defined as a *prime implicant* of that
40 proposition, i.e. an implicant that cannot be generalized further. For instance, an explanation
41 of a Machine Learning model's prediction is often defined as a minimal subset of features
42 that entails that prediction [16, 10]. Similarly, a *minimal unsatisfiable core* (irreducible
43 unsatisfiable subset of constraints) can be seen as an abductive explanation for unsatisfiability
44 since it is a sufficient and minimal reason for unsatisfiability. At least one term of an abductive



© Anonymous author(s);
licensed under Creative Commons License CC-BY 4.0

CP 2022.



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 explanation must be relaxed in order to change the outcome. This is the viewpoint adopted
 46 in many existing approaches. For instance by providing explanations in the form of minimal
 47 sets of choices of the user that lead to the given value removal (e.g., product configuration
 48 [1]), or explanations in the form of minimal sets of constraints that lead to an inconsistency
 49 [11]. The purpose of such approaches is to help the user to repair the inconsistency, not to
 50 let them understand why it is an inconsistency.

51 Intuitively, an explanation is more than a sufficient condition. In particular, if an
 52 abductive explanation answers the “*why*” question, it does not answer the “*how*” question.
 53 An intuitive definition of an explanation also covers the *demonstration* of how the considered
 54 cause has that consequence. For instance, when solving a logic puzzle, we may want to
 55 let the user understand why the zebra is necessarily in the middle house, not by providing
 56 a set of constraints of the problem that rule out all other positions for the zebra, but by
 57 displaying a sequence of simple reasoning steps that lead to that conclusion. This notion
 58 of *demonstrative explanation* can be related to proof systems and to the notion of formal
 59 proof. A formal proof better explains unsatisfiability by making every step explicit down to
 60 axiomatic definitions. For instance, a refutation proof log using the *reverse unit propagation*
 61 (RUP) system [8, 9] allows to formally verify the unsatisfiability of a formula, provided that
 62 one can “trust” the application of the unit propagation rule, i.e. trust that a given formula
 63 that is refutable via unit propagation is indeed unsatisfiable. This is valid in the context of
 64 formal proof verification where each unit propagation refutation can be checked efficiently.
 65 However, this may produce very long proofs in which each step might be too complex for an
 66 explanation to a non-expert.

67 We would therefore want to produce demonstrative explanations, allowing a trustworthy
 68 verification, however with minimal requirements on the recipient of the explanation. This is
 69 of course impossible in general. In [17, 2], the choice was made to provide explanations in
 70 the form of sequences of inferences performed by constraint propagation. We consider an
 71 even simpler, and also incomplete, proof system: Arc Consistency. Arc consistency has often
 72 been considered as a sufficiently strong inference technique on applications where the human
 73 is in the loop (configuration [12], logic puzzles [17]).

74 Our goal is to analyze the complexity of providing the *shortest* possible explanation of arc
 75 inconsistency of a problem. For simplicity of presentation, we restrict ourselves to normalized
 76 networks of binary constraints. We show that when variables have degree two or domains
 77 are Boolean, finding a shortest explanation of arc inconsistency is polynomial. However, the
 78 problem is NP-hard in general and the two polynomial cases above are tight. Finding a
 79 shortest explanation of arc inconsistency is NP-hard as soon as variables have degree three,
 80 even if the number of variables is bounded (even though the problem is obviously polynomial
 81 to solve). It is also NP-hard if domain size is bounded by three. Perhaps more surprisingly,
 82 it remains NP-hard on trees, where arc consistency is known to be a decision procedure.
 83 We also show that there is little hope that we can efficiently find short (if not shortest)
 84 explanations: the problem is not FPT-approximable unless the Gap-ETH conjecture is false.

85 **2** Background and Definitions

86 The *constraint satisfaction problem (CSP)* involves finding solutions to a constraint network.
 87 A *constraint network* (or CSP instance) is defined as a set of n variables $X = X_1, \dots, X_n$,
 88 a set of domains $D = \{D(X_1), \dots, D(X_n)\}$, where $D(X_i)$ is the finite set of values that X_i
 89 can take, and a set C of constraints. A binary constraint $c(X_i, X_j)$ is a binary relation that
 90 specifies which combinations of values (tuples) the variables (X_i, X_j) are allowed to take. A

91 CSP is *binary* when all the constraints are binary. A binary CSP is said to be *normalized* if
 92 there is at most one constraint per pair of variables. A *degree-2* CSP does not contain any
 93 variable involved in more than two constraints. *Arc consistency (AC)* is the basic form of
 94 inference reasoning on constraint networks. A tuple τ of values on (X_i, X_j) is called a *support*
 95 on constraint $c(X_i, X_j)$ for a value $v \in D(X_i)$ (and $\tau[X_j]$ its support in $D(X_j)$) if and only
 96 if $\tau[X_i] = v$, $\tau[X_j] \in D(X_j)$ and $\tau \in c(X_i, X_j)$. A value v in $D(X_i)$ is arc consistent if and
 97 only if v has a support on every constraint involving X_i . A network is arc consistent if all
 98 values in all domains are arc consistent. The operation $revise(X_i, c(X_i, X_j))$, often denoted
 99 by $X_i \stackrel{c}{\leftarrow} X_j$ in the following, removes from $D(X_i)$ all values that do not have any support
 100 on $c(X_i, X_j)$. If enforcing arc consistency on a network (that is, applying $revise()$ operations
 101 until a fix point is reached) leads to a domain wipe out (i.e. an empty domain), we say that
 102 the network is *arc inconsistent*.

103 ► **Definition 1** (Arc Inconsistency Explanation). *An arc inconsistency explanation for a CSP*
 104 *instance is a sequence of $revise()$ operations such that one of the domains is wiped out by*
 105 *the execution of the sequence of $revise()$ operations.*

106 ► **Definition 2** (Shortest Arc Inconsistency Explanation). *The shortest arc inconsistency*
 107 *explanation problem consists in finding an arc inconsistency explanation of minimum length.*

108 ► **Example 3** (Explaining the Zebra puzzle). The Zebra puzzle, which may (or may not) be
 109 due to Lewis Carroll, has a well known CSP model whereby, for each of the 5 *house colors*,
 110 *nationalities*, *beverages*, *cigarette brands*, and *pets*, we have a variable whose value is the
 111 number of the corresponding house (e.g., X_{Zebra} stands for the house where the Zebra lives).
 112 The constraints are statements such as *The Englishman lives in the red house* or *The Old*
 113 *Gold smoker owns snails*. Moreover, each house has a unique colour, its owner has a unique
 114 nationality, drinks a unique beverage, smokes a unique brand, and has a unique pet.

115 Applying arc consistency on this CSP detects that “**the Kools smoker does not live**
 116 **in the 2nd house**”. A demonstrative explanation would be: **The Norwegian lives in**
 117 **the first house**. Since *the Norwegian lives next to the blue house*, then **the 2nd house is**
 118 **blue**. Since *the 2nd house has a single color*, then **it is not yellow**. Since *Kools are smoked*
 119 *in the yellow house*, then **the Kools smoker does not live in the 2nd house**.

120 Each step corresponds to the arc consistency *revision* of some domain knowledge (in bold)
 121 with respect to a constraint (in italic), that is, it corresponds in our framework to the following
 122 sequence of $revise()$ operations: $\langle X_{Blue} \leftarrow X_{Norwegian}, X_{Yellow} \stackrel{c}{\leftarrow} X_{Blue}, X_{Kools} \stackrel{c}{\leftarrow} X_{Yellow} \rangle$.

123 3 Complexity of Explaining Arc Inconsistency: Structure

124 We show that if all variables are involved in no more than two constraints, finding shortest
 125 arc inconsistency explanations is polynomial. We then show that this class is tight. As soon
 126 as we allow a variable to be in the scope of three constraints, the problem becomes NP-hard,
 127 even if the CSP has no more than four variables. Perhaps even more surprising, the problem
 128 is NP-hard on CSPs structured as trees, despite arc consistency being a decision procedure
 129 on trees.

130 3.1 Tractability on degree-2 CSPs

131 ► **Theorem 4.** *SHORTEST ARC INCONSISTENCY EXPLANATION is solvable in polynomial*
 132 *time when restricted to binary normalized networks with maximum degree two.*

XX:4 Complexity of Minimum-Size Arc-Inconsistency Explanations

133 **Proof.** A constraint network of maximum degree two is composed of unconnected cycles
134 and paths. A shortest arc inconsistency explanation clearly always concerns only one of the
135 connected components of the network. An exhaustive search over all connected components
136 only increases complexity by at most a linear factor. Since, furthermore a path can be viewed
137 as a degenerate cycle (a cycle in which one constraint disallows no tuples), it follows that we
138 only need consider the case of a single cycle.

139 Without loss of generality, we suppose that the cycle is X_1, \dots, X_n , with constraints
140 $c(X_i, X_{i+1})$, where here and in the rest of the proof addition within subscripts is understood to
141 be modulo n , so that for example X_{n+1} actually refers to X_1 . We say that $revise()$ operations
142 are clockwise (resp. anticlockwise) if they are of the form $X_{i+1} \leftarrow X_i$ (resp. $X_i \leftarrow X_{i+1}$). We
143 say that a pair of $revise()$ operations R_1, R_2 commute if the two sequences R_1R_2 and R_2R_1
144 produce the same result. It is easy to verify that the only $revise()$ operations that may not
145 commute are those in which the destination variable of one is the source variable of the other.
146 Furthermore, $revise()$ operations in opposite directions (clockwise and anticlockwise) always
147 commute, even $X_i \leftarrow X_{i+1}$ and $X_{i+1} \leftarrow X_i$. Thus the only pairs of $revise()$ operations that
148 do not commute are of the form $\{X_i \leftarrow X_{i+1}, X_{i+1} \leftarrow X_{i+2}\}$. What's more, if we have the
149 operations $X_i \leftarrow X_{i+1}, X_{i+1} \leftarrow X_{i+2}$ in this order, then the set of value-eliminations cannot
150 decrease if we inverse the order of these two operations.

151 In a shortest arc inconsistency explanation E , a $revise()$ operation must be useful: it must
152 eliminate a domain value whose elimination is essential for a subsequent $revise()$ operation
153 or for the final domain wipe-out. In the former case, the operation $X_i \leftarrow X_{i+1}$ must be
154 followed later in the sequence by $X_{i-1} \leftarrow X_i$. Let S be the sequence of $revise()$ operations
155 in E between the operation $X_i \leftarrow X_{i+1}$ and the next subsequent occurrence of $X_{i-1} \leftarrow X_i$.
156 By the above discussion on commutativity, we can shift the operation $X_i \leftarrow X_{i+1}$ just after
157 S without *decreasing* the set of value-eliminations since S does not contain $X_{i-1} \leftarrow X_i$. In
158 this way, we can group together all the anticlockwise $revise()$ operations to form a sequence
159 of anticlockwise operations on consecutive edges in the cycle. The same argument holds
160 for clockwise operations which can be grouped together to form a sequence of clockwise
161 operations on consecutive edges in the cycle.

162 An obvious observation is that a shortest arc inconsistency explanation is necessarily of
163 length bounded by nd , where d is the maximum domain size, since at least one elimination
164 must occur at each operation. Moreover, there are up to n possible starting points for the
165 sequence of clockwise (resp. anticlockwise) operations. Hence a shortest explanation can be
166 found in polynomial time, by exhaustive search over the starting points and lengths of the
167 clockwise/anticlockwise sequences). ◀

168 3.2 Intractability on CSPs with four variables

169 The result in Theorem 4 is tight. We show that as soon as we allow variables to have degree
170 3, finding a shortest explanation becomes NP-hard. This is true even if the number of
171 variables is bounded by four. (Observe that all binary normalized CSPs on three variables are
172 degree-2.) We use a reduction from CLIQUE, which is NP-complete [13], to prove hardness.

173 ► **Definition 5** (CLIQUE).

174 **Input:** An undirected graph $G = (V, E)$ and an integer k

175 **Question:** Is there $S \subseteq V$ such that $|S| \leq k$ and for all $i \neq j \in S$, $\{i, j\} \in E$?

176 It is noticeable that CSPs with a bounded number of variables have a constant number
177 of possible $revise()$ operations available at each step –only 12 in the case of four variables.
178 This is not sufficient to make the problem of finding a shortest explanation easy.

179 ► **Theorem 6.** SHORTEST ARC INCONSISTENCY EXPLANATION is NP-hard, even on binary
180 normalized networks with four variables.

181 ► **Lemma 7.** Deciding whether there exists an arc inconsistency explanation of length smaller
182 than or equal to k is NP-complete, even on binary normalized networks with four variables.

183 **Proof.** *Membership.* Given a sequence of $revise()$ operations, we decide whether this sequence
184 is an arc inconsistency explanation by executing each $revise()$ in the order of the sequence
185 and checking whether one of the domains is empty after these executions. As constraints
186 have bounded arity, executing a $revise()$ operation is polynomial, so the whole process is
187 polynomial.

188 *Completeness.* We reduce the CLIQUE problem to the problem of deciding whether there is
189 an arc inconsistency explanation of length at most $3n + 3$ for a CSP instance. Let $G = (V, E)$
190 be a graph with set of vertices $V = \{1, \dots, n\}$.

191 We construct the CSP instance P_G with four variables $X = \{X_1, X_2, X_3, X_4\}$, all with
192 domain $\{(p, i) : p \in 0..n + 1, i \in 1..n\} \cup \{s_t : t \in 1..k + 1\}$.

We build the set of constraints

$$C = \{c_1(X_1, X_2), c_2(X_1, X_3), c_3(X_2, X_3), X_1 = X_4, X_2 = X_4, X_3 = X_4\}$$

193 with:

$$\begin{aligned} 194 \quad c_1(X_1, X_2) &= \{(p-1, i), (p, i) : p \in [0, n+1], \forall i \neq p \in [1, n]\} \\ 195 &\quad \cup \{(p-2, i), (p, i) : p \in [0, n+1], \forall i \in [1, n]\} \\ 196 &\quad \cup \{(s_t, s_t) : t \in [1, k+1]\} \\ 197 \quad c_2(X_1, X_3) &= \{(p-1, i), (p, i) : p \in [0, n+1], \forall i \in [1, n]\} \cup \{(s_{t-1}, s_t) : t \in [1, k+1]\} \\ 198 \quad c_3(X_2, X_3) &= (\{(p, i) : p \in [0, n+1], i \in [1, n]\}^2) \setminus \\ 199 &\quad \{((n+1, i), (n+1, j)) : i = j \vee \{i, j\} \in E\} \\ 200 &\quad \cup \{(s_t, s_t) : t \in [1, k+1]\}^2 \\ 201 \end{aligned}$$

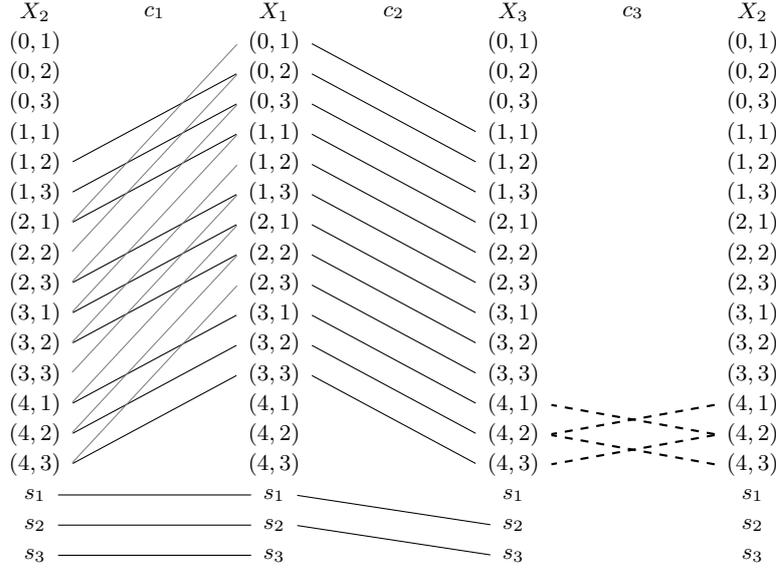
202 The constraint network for a graph with 3 vertices and the edges $\{1, 2\}$ and $\{2, 3\}$ is shown
203 in Figure 1.

204 We first show that if G contains a k -clique, then, there exists an arc inconsistency
205 explanation of length $3n + 3$ for P_G .

206 Assume that the set of vertices S is a k -clique. We build the sequence $R(S)$ of $revise()$
207 operations in the following way, and we will say that $R(S)$ encodes the set S , since there is a
208 one-to-one mapping between subsets $S \subseteq V$ and this type of explanation:

- 209 ■ If $p \notin S$, the $(3p - 2)$ th element in the sequence $R(S)$ is $X_2 \stackrel{\leftarrow 1}{\leftarrow} X_1$, the $(3p - 1)$ th element
210 is $X_4 \stackrel{\leftarrow}{\leftarrow} X_2$, and the $(3p)$ th element is $X_1 \stackrel{\leftarrow}{\leftarrow} X_4$.
- 211 ■ If $p \in S$, the $(3p - 2)$ th element in the sequence $R(S)$ is $X_3 \stackrel{\leftarrow 2}{\leftarrow} X_1$, the $(3p - 1)$ th element
212 is $X_4 \stackrel{\leftarrow}{\leftarrow} X_3$, and the $(3p)$ th element is $X_1 \stackrel{\leftarrow}{\leftarrow} X_4$.

213 Then the last three elements in the sequence $R(S)$ are $X_2 \stackrel{\leftarrow 1}{\leftarrow} X_1$, $X_3 \stackrel{\leftarrow 2}{\leftarrow} X_1$, and $X_2 \stackrel{\leftarrow 3}{\leftarrow} X_3$.
214 In the following, the subsequence composed of the $(3p - 2)$ th, the $(3p - 1)$ th, and the $(3p)$ th
215 operations (that is, $\langle X_2 \stackrel{\leftarrow 1}{\leftarrow} X_1, X_4 \stackrel{\leftarrow}{\leftarrow} X_2, X_1 \stackrel{\leftarrow}{\leftarrow} X_4 \rangle$ or $\langle X_3 \stackrel{\leftarrow 2}{\leftarrow} X_1, X_4 \stackrel{\leftarrow}{\leftarrow} X_3, X_1 \stackrel{\leftarrow}{\leftarrow} X_4 \rangle$), is
216 called the p th iteration.



■ **Figure 1** The CSP P_G , reduction of the graph $G = (\{1, 2, 3\}, \{(1, 2), (2, 3)\})$. Solid edges represent allowed tuples for c_1 and c_2 , while dashed edges stand for forbidden tuples of c_3 . The equality constraints are not represented. There are two explanations of Arc-Inconsistency of length 12. The first *encodes* the clique $\{2, 3\}$ with the *revise()* operations $X_2 \stackrel{c_1}{\leftarrow} X_1$, $X_3 \stackrel{c_2}{\leftarrow} X_1$, $X_3 \stackrel{c_3}{\leftarrow} X_1$ at positions 1, 4, and 7 in the sequence. The second *encodes* the clique $\{1, 2\}$ with the revision operations $X_3 \stackrel{c_2}{\leftarrow} X_1$, $X_3 \stackrel{c_3}{\leftarrow} X_1$, $X_2 \stackrel{c_1}{\leftarrow} X_1$ at positions 1, 4, and 7.

217 Before each iteration $p \in \{1, \dots, n\}$ of three domain revisions, the invariants are:

218
$$(q, i) \notin D(X_1) \quad \forall q < p - 1, \forall i \in [1, n] \tag{1}$$

219
$$s_j \in D(X_1) \iff k + 1 \geq j > |S \cap \{0, \dots, p - 1\}| \tag{2}$$

220
$$(p - 1, i) \in D(X_1) \iff i \in S \cup \{p, \dots, n\} \tag{3}$$

222 All invariants are verified before entering iteration $p = 1$. For each one, we show that if it
 223 is true before entering iteration $p \geq 1$ then it remains true before entering iteration $p + 1$.

224 Invariant 1: Notice that a value $(q, i) \in D(X_2)$ (resp. $D(X_3)$) is only supported by values
 225 $(q', i) \in D(X_1)$ such that $q' < q$. If Invariant 1 is true before iteration p , then when revising
 226 the domain of either X_2 or X_3 , $D(X_1)$ contains no value (q, i) with $q < p - 1$ and therefore
 227 all values $(p - 1, i)$ are removed from $D(X_2)$ (resp. $D(X_3)$). The revisions w.r.t. equality
 228 constraints make sure that this is propagated back to $D(X_1)$.

229 Invariant 2. Notice that a value $s_t \in D(X_3)$ is only supported by value $s_{t-1} \in D(X_1)$,
 230 whereas the tuple (s_t, s_t) is a support in all other constraints. If Invariant 2 is true before
 231 iteration p , then either $p \in S$ in which case the operation $X_3 \stackrel{c_2}{\leftarrow} X_1$ removes the value s_j (with
 232 $j = |S \cap \{0, \dots, p - 1\}| + 1$) from $D(X_3)$ since the value s_{j-1} was its only support and is not
 233 in $D(X_1)$; or $X_2 \stackrel{c_1}{\leftarrow} X_1$ removes no s value and $|S \cap \{0, \dots, p - 1\}|$ does not change.

234 Invariant 3. For any $i \in [1, n]$:

235 If $i > p$, then we have $(p - 1, i) \in D(X_1)$ which is a support for (p, i) w.r.t. c_1 and c_2
 236 hence (p, i) is not removed and the invariant holds because $i \in \{p + 1, \dots, n\}$.

237 If $i < p$, notice that by Invariant 1, the tuple $((p - 2, i), (p, i))$ cannot be a support
 238 for $(p, i) \in D(X_2)$. Therefore, both constraints c_1 and c_2 have the same unique potential
 239 support for the value (p, i) (in $D(X_2)$ and $D(X_3)$ respectively): $((p - 1, i), (p, i))$. So we have:

240 “ $(p-1, i) \in D(X_1)$ before iteration p ” iff “ $(p, i) \in D(X_1)$ before iteration $p+1$ ”. In addition,
 241 $i \in S \cup \{p, \dots, n\} \iff i \in S \cup \{p+1, \dots, n\}$ because $i < p$. Finally, by the induction
 242 hypothesis we have “ $(p-1, i) \in D(X_1)$ before iteration p ” iff $i \in S \cup \{p, \dots, n\}$, and hence
 243 by transitivity: “ $(p, i) \in D(X_1)$ before iteration $p+1$ ” iff $i \in S \cup \{p+1, \dots, n\}$.

244 If $i = p$, there are two cases: If $p \in S$, then the first operation at iteration p is $X_3 \stackrel{c_2}{\leftarrow} X_1$,
 245 (p, i) is not removed since it is supported by $(p-1, i)$, and the invariant is true at iteration
 246 $p+1$ since $i \in S$. If $p \notin S$, then the first operation at iteration p is $X_2 \stackrel{c_1}{\leftarrow} X_1$, (p, i) is
 247 removed, and the invariant is true at iteration $p+1$ since $i \notin S \cup \{p+1, \dots, n\}$.

248 After n iterations, the invariants hold for $p = n+1$ (i.e. after the $3n$ -th operation) and
 249 hence $D(X_1)$ is $\{(n, i) \forall i \in S\} \cup \{(n+1, i) \forall i\} \cup \{s_{k+1}\}$. The call to $X_2 \stackrel{c_1}{\leftarrow} X_1$ then yields
 250 $D(X_2) = \{(n+1, i) \forall i \in S\} \cup \{s_{k+1}\}$ and the call to $X_3 \stackrel{c_2}{\leftarrow} X_1$ yields $D(X_3) = \{(n+1, i) \forall i \in$
 251 $S\}$. Therefore, the last call to $X_2 \stackrel{c_3}{\leftarrow} X_3$ produces a wipe-out, since on layer $n+1$, the
 252 remaining vertices stand for a clique of G and the allowed tuples are non-edges of G .

253 We then prove that if G does not contain any k -clique, then the shortest arc inconsistency
 254 explanation for P_G is of length strictly greater than $3n+3$. We first show that the shortest
 255 explanation must use constraint c_3 , then we show that only explanations that encode a set
 256 $S \subseteq V$ (as defined above) such that S is a clique of size k of G can be the shortest.

257 Suppose first that the constraint c_3 does not appear in any $revise()$ of the explanation.
 258 By construction, the values (p, i) are organized in layers, where a layer q is the set of
 259 values $(q, i), \forall i$. Wiping out the domain of a variable requires removing the $n+2$ layers
 260 0 to $n+1$ from its domain. Moreover, removing a layer q from X_1 (resp. X_2 or X_3)
 261 requires having already removed layer $q+1$ (resp. $q-1$) from X_2 or X_3 (resp. X_1).
 262 Removing a layer q from X_4 requires having already removed layer q from X_1, X_2 , or
 263 X_3 . Hence, removing a layer q from a variable requires iteratively removing layers 0
 264 to $q-1$ or $n+1$ down to $q+1$ from other variables. The only way to do that is to
 265 execute a sequence of $revise()$ operations looping on a cycle of variables $\{X_1, X_2, X_4\}$,
 266 or on $\{X_1, X_3, X_4\}$, or both. Looping in the order $\langle X_1 \stackrel{c_1}{\leftarrow} X_2, X_4 \stackrel{c_2}{\leftarrow} X_1, X_2 \stackrel{c_3}{\leftarrow} X_4 \rangle$ or
 267 $\langle X_1 \stackrel{c_2}{\leftarrow} X_3, X_4 \stackrel{c_1}{\leftarrow} X_1, X_3 \stackrel{c_3}{\leftarrow} X_4 \rangle$ removes layers from $n+1$ down to q , whereas looping in the
 268 order $\langle X_2 \stackrel{c_1}{\leftarrow} X_1, X_4 \stackrel{c_2}{\leftarrow} X_2, X_1 \stackrel{c_3}{\leftarrow} X_4 \rangle$ or $\langle X_3 \stackrel{c_2}{\leftarrow} X_1, X_4 \stackrel{c_1}{\leftarrow} X_3, X_1 \stackrel{c_3}{\leftarrow} X_4 \rangle$ removes layers
 269 from 0 up to q . We can then compute the number of $revise()$ operations necessary to remove
 270 a layer q from a variable given the order in which we loop. If we execute $revise()$ operations
 271 in the orders $\langle X_1 \stackrel{c_1}{\leftarrow} X_2, X_4 \stackrel{c_2}{\leftarrow} X_1, X_2 \stackrel{c_3}{\leftarrow} X_4 \rangle$ or $\langle X_1 \stackrel{c_2}{\leftarrow} X_3, X_4 \stackrel{c_1}{\leftarrow} X_1, X_3 \stackrel{c_3}{\leftarrow} X_4 \rangle$, layer q
 272 is removed from the domain of X_1 (resp. X_2/X_3 , or X_4) in $3(n+1-q)+1$ operations
 273 (resp. $3(n+1-q)+3$, or $3(n+1-q)+2$ operations). If we execute $revise()$ operations
 274 in the orders $\langle X_2 \stackrel{c_1}{\leftarrow} X_1, X_4 \stackrel{c_2}{\leftarrow} X_2, X_1 \stackrel{c_3}{\leftarrow} X_4 \rangle$ or $\langle X_3 \stackrel{c_2}{\leftarrow} X_1, X_4 \stackrel{c_1}{\leftarrow} X_3, X_1 \stackrel{c_3}{\leftarrow} X_4 \rangle$, layer q is
 275 removed from the domain of X_1 (resp. X_2/X_3 , or X_4) in $3q+3$ operations (resp. $3q+1$, or
 276 $3q+2$ operations). As wiping out a domain requires, given a value q , to remove layers 0 to q
 277 from below and layers $n+1$ down to $q+1$ from above, we conclude that a domain wipe out,
 278 on either X_1, X_2, X_3 , or X_4 , requires at least $3n+4$ $revise()$ operations. This means that
 279 there does not exist any arc inconsistency explanation for P_G of length smaller than or equal
 280 to $3n+3$ if we do not use c_3 in the explanation.

281 Hence, we must use c_3 . However, by construction of c_3 , every value in $D(X_2)$ (resp.
 282 $D(X_3)$) is supported as long as at least one value (p, i) with $p \in [0, n]$, and any value s_t is
 283 in the domain of $D(X_3)$ (resp. $D(X_2)$). In other words, to remove a layer with a $revise$
 284 on c_3 , the domains of X_2 and X_3 must only contain (p, i) values from layer $n+1$. This
 285 requires to remove all layers from 0 to $n-1$ from X_1 by executing n loops by a sequence of
 286 $revise()$ operations $\langle X_2/X_3 \leftarrow X_1, X_4 \leftarrow X_2/X_3, X_1 \leftarrow X_4 \rangle$ for a cost of $3n$ operations, plus

287 a $X_2 \stackrel{c_1}{\leftarrow} X_1$ and a $X_3 \stackrel{c_2}{\leftarrow} X_1$ to remove layer n from X_2 and X_3 . In other words, it must be
 288 a sequence of *revise()* operations that *encodes* a set, i.e., $R(S)$ for some set $S \subseteq \{1, \dots, n\}$.
 289 Now, suppose that S is not a clique and let i_1 and i_2 be two non-adjacent vertices in S . By
 290 Invariant 3, at iteration $n + 1$, we have $(n, i_1) \in D(X_1)$ and $(n, i_2) \in D(X_1)$ and hence after
 291 operations $X_2 \stackrel{c_1}{\leftarrow} X_1$ and $X_3 \stackrel{c_2}{\leftarrow} X_1$, we have $(n + 1, i_1) \in D(X_2)$ and $(n + 1, i_2) \in D(X_3)$.
 292 Therefore, neither $X_2 \stackrel{c_3}{\leftarrow} X_3$ nor $X_3 \stackrel{c_3}{\leftarrow} X_2$ would fail, and at least one more operation is
 293 necessary. Finally, suppose that $|S| < k$. Then by Invariant 2, at iteration $n + 1$, we have
 294 $s_k \in D(X_1)$ and hence after operations $X_2 \stackrel{c_1}{\leftarrow} X_1$ and $X_3 \stackrel{c_2}{\leftarrow} X_1$, we have $s_{k+1} \in D(X_2)$
 295 and $s_{k+1} \in D(X_3)$. Therefore, at least one more operation is necessary. Consequently, the
 296 number of operations can be equal to $3n + 3$ only if S is a clique of size k of G . ◀

297 3.3 Intractability and inapproximability on trees

298 We have seen in Section 3.2 that SHORTEST ARC-INCONSISTENCY EXPLANATION is already
 299 NP-hard on networks with four variables. This result does not completely settle the intract-
 300 ability of the problem. For example, it is still possible that a polynomial-time algorithm
 301 exists for some broad generalization of degree-2 networks that does not contain 4-cliques (for
 302 instance, networks of treewidth 2). We show that it is not the case. We use a simple reduc-
 303 tion from DOMINATING SET, which is NP-complete [7], to derive NP-hardness of SHORTEST
 304 ARC-INCONSISTENCY EXPLANATION, even on trees.

305 ▶ **Definition 8** (DOMINATING SET).

306 **Input:** An undirected graph $G = (V, E)$ and an integer k

307 **Question:** Is there $S \subseteq V$ such that $|S| \leq k$ and for all $i \in V$, there is $j \in S$ with $\{i, j\} \in E$?

308 The NP-hardness of SHORTEST ARC-INCONSISTENCY EXPLANATION on trees circum-
 309 scribes even more tightly the degree-2 tractability island of Section 3.1. However, these
 310 NP-hardness results do not rule out efficient approximation algorithms nor fixed-parameter
 311 tractable algorithms, which could be satisfactory for applications where only short explan-
 312 ations are worth computing and optimality is not strictly necessary. We again show that
 313 such desirable scenarios are not possible. We show that our reduction from DOMINATING
 314 SET can be used to derive (conditional) fixed-parameter inapproximability of SHORTEST
 315 ARC-INCONSISTENCY EXPLANATION.

316 We must briefly introduce some terminology before we can formally present the result.
 317 A minimization problem \mathcal{P} is *fpt-approximable* [4] if there exist computable functions $f, \rho :$
 318 $\mathbb{N} \rightarrow \mathbb{R}_{\geq 1}$ such that $n \cdot \rho(n)$ is nondecreasing and an algorithm A that, given as input a
 319 non-negative integer k and an instance I of \mathcal{P} that has a solution of cost at most k , computes
 320 a solution to I of cost at most $k \cdot \rho(k)$ in time $f(k) \cdot |I|^{O(1)}$. Here, ρ is the approximation
 321 ratio and f is possibly exponential. Note that if a problem is not FPT-approximable, then no
 322 such algorithm A exists for *any* computable functions f and ρ ; such problems are sometimes
 323 called *completely inapproximable* [15].

324 Our FPT-inapproximability result is conditional on a complexity hypothesis known as the
 325 *Gap-ETH* [6, 14], which states that there exists a constant $\epsilon > 0$ such that no algorithm with
 326 runtime $2^{o(n)}$ can distinguish satisfiable 3-SAT instances from those in which no assignment
 327 satisfies a $(1 - \epsilon)$ fraction of the clauses. It has been shown recently [3] that the MINIMUM
 328 DOMINATING SET problem (which consists in finding the smallest dominating set in a graph)
 329 is not FPT-approximable unless the Gap-ETH is false.

330 ► **Lemma 9.** *Deciding whether there exists an arc inconsistency explanation of length smaller*
 331 *than or equal to k is NP-complete, even on binary tree-structured normalized constraint*
 332 *networks.*

333 **Proof.** *Membership.* As in Lemma 7.

334 *Completeness.* We reduce the DOMINATING SET problem to the problem of deciding
 335 whether there is an arc inconsistency explanation of length at most k for a CSP instance.

336 Let $G = (V, E)$ be a graph, $V = \{v_1, \dots, v_n\}$. We construct a constraint network P_G as
 337 follows: the set of variables is $\{Y, X_1, \dots, X_n\}$, where the domain of Y is $\{v_1, \dots, v_n\}$ and
 338 the domain of each X_i is $\{v_i\}$, and P_G contains a constraint $c(Y, X_i) = \{(v_j, v_i) : \{v_i, v_j\} \notin$
 339 $E \text{ and } v_i \neq v_j\}$ for all $i \geq 1$. An example of this reduction is shown in Figure 2. We claim
 340 that G has a dominating set of size k if and only if P_G has an arc-inconsistency explanation
 341 of length k .

342 If G has a dominating set S of size k , then let R be a sequence containing every operation
 343 $Y \leftarrow X_i$ such that v_i belongs to S . Since every $v_j \in V$ is dominated by some $v_k \in S$ (which
 344 is either v_j itself or one of its neighbours), by construction v_j is removed from $D(Y)$ by
 345 $Y \leftarrow X_k$. Therefore $D(Y)$ is empty at the end of the sequence and R is an arc-inconsistency
 346 explanation of length k .

347 Conversely, if R is a minimal arc-inconsistency explanation of P_G of length k then we can
 348 assume that it is a sequence of operations of the form $Y \leftarrow X_i$. (Since each $D(X_i)$ contains
 349 a single value, only the last operation could be $X_i \leftarrow Y$ for some i , and in that case it can be
 350 replaced with $Y \leftarrow X_i$.) Then, the set $S = \{v_i : Y \leftarrow X_i \text{ occurs in } R\}$ must be a dominating
 351 set of size k : at the end of R every $v_j \in D(Y)$ has been pruned by some operation $Y \leftarrow X_k$,
 352 and every value removed at this step is by construction dominated by v_k in G .

353 P_G is a tree-structured constraint network and can be constructed in polynomial time
 354 from G . Therefore, SHORTEST ARC-INCONSISTENCY EXPLANATION is NP-hard on such
 355 networks. ◀

356 ► **Theorem 10.** SHORTEST ARC INCONSISTENCY EXPLANATION *is NP-hard and not FPT-*
 357 *approximable unless the Gap-ETH is false, even on binary tree-structured normalized con-*
 358 *straint networks.*

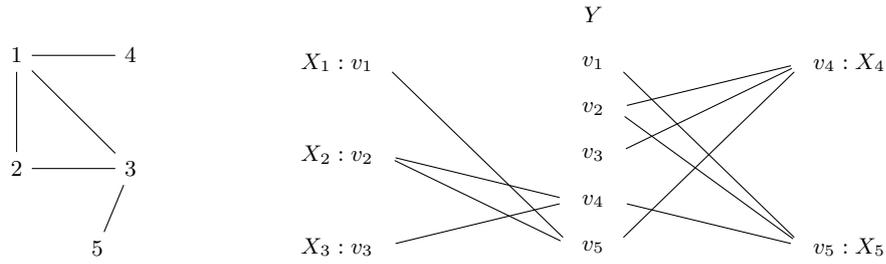
359 **Proof.** In the reduction of the proof of Lemma 9, the size- k dominating sets of G are in one-
 360 to-one correspondence with arc-inconsistency explanations of P_G of length k . Furthermore,
 361 the dominating set corresponding to an explanation can be computed in polynomial time,
 362 so any FPT-approximation algorithm for SHORTEST ARC-INCONSISTENCY EXPLANATION
 363 translates into one for MINIMUM DOMINATING SET. By the results of [3], this would imply
 364 that the Gap-ETH is false. ◀

365 As a final remark, we note that the same inapproximability result can be established
 366 under the weaker (and more conventional) complexity hypothesis $\text{FPT} \neq \text{W}[2]$. However,
 367 the proof is significantly more involved and has been left out for the sake of brevity. It will
 368 be made available in the full version of this paper.

369 4 Complexity of Explaining Arc Inconsistency: Domain Size

370 We show that finding shortest arc inconsistency explanations is polynomial on binary
 371 normalized CSPs with Boolean domains. Again, this class is tight: As soon as we allow three
 372 values per domain, the problem becomes NP-hard.

XX:10 Complexity of Minimum-Size Arc-Inconsistency Explanations



■ **Figure 2** Left: a graph G . Right: the constraint network P_G in the proof of Lemma 9.

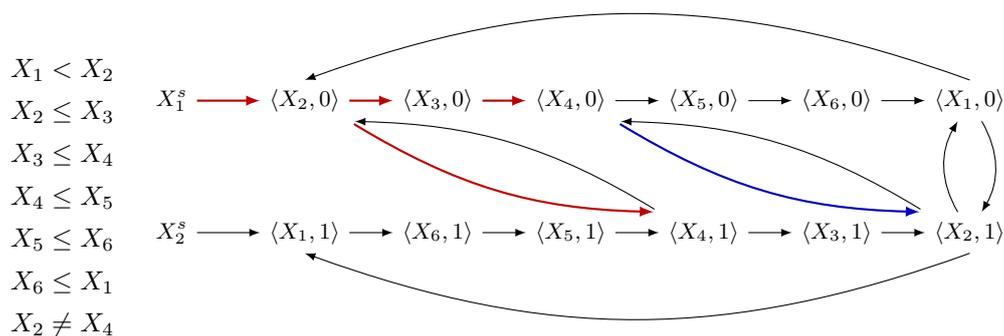
373 4.1 Tractability on Boolean domains

374 ► **Theorem 11.** SHORTEST ARC INCONSISTENCY EXPLANATION *is solvable in polynomial*
 375 *time when restricted to binary normalized networks with all domains of size at most two.*

376 **Proof.** Let $P = (X, D, C)$ be a binary CSP with domain size at most two. We assume,
 377 without loss of generality, that all domains $D(X_i)$ are non-empty subsets of $\{0, 1\}$ and
 378 that no constraint relation is empty. Let X_r be the variable at which a domain wipe-out
 379 occurs in a shortest arc inconsistency explanation. Complexity is only multiplied by n
 380 if we perform an exhaustive search over all possible variables X_r , so in the following we
 381 consider X_r to be fixed. We construct a directed causal graph G_P in which shortest arc
 382 inconsistency explanations correspond to particularly simple subgraphs. In G_P there are
 383 two types of vertices: source-variable vertices X_i^s ($i = 1, \dots, n$), and variable-value vertices
 384 $\langle X_i, a \rangle$ ($i = 1, \dots, n, a \in \{0, 1\}$). G_P has the following directed edges: $(X_i^s, \langle X_j, b \rangle)$ (for all
 385 i, j, b such that $b \in D(X_j)$ has no support in $D(X_i)$), and $(\langle X_i, a \rangle, \langle X_j, b \rangle)$ (for all i, j, a, b
 386 such that $a \in D(X_i)$ is the only support of $b \in D(X_j)$). Each arc corresponds to a possible
 387 revise operation: $(X_i^s, \langle X_j, b \rangle)$ corresponds to the elimination of b from $D(X_j)$ since it has
 388 no support in $D(X_i)$, and $(\langle X_i, a \rangle, \langle X_j, b \rangle)$ corresponds to the elimination of b from $D(X_j)$
 389 when its unique support $a \in D(X_i)$ has been eliminated. An example of the causal for a
 390 simple CSP is shown in Figure 3.

391 Let R be a shortest arc inconsistency explanation, and let X_r be the variable at which
 392 a wipe-out occurs. By minimality of R , each revise operation in R eliminates a value from
 393 a domain. Indeed, each operation, except possibly the last, eliminates exactly one value
 394 otherwise there would be a domain wipe-out before the end of R . Furthermore, the only way
 395 that the final revise operation $X_r \leftarrow X_i$ of R can cause the simultaneous elimination of both
 396 0 and 1 from $D(X_r)$ (without there already being a wipe-out at $D(X_i)$) is that (1) some
 397 value $b \in D(X_r)$ never had any support at X_i and (2) the other value $1-b$ lost its unique
 398 support a at X_i by a previous operation in R . We can deduce from (1) and (2) that just
 399 before the execution of $X_r \leftarrow X_i$, the value $1-a$ in $D(X_i)$ has no support at X_r . This implies
 400 that we can replace the last operation $X_r \leftarrow X_i$ of R by its inverse operation $X_i \leftarrow X_r$ to
 401 produce an arc inconsistency explanation of the same length as R but in which the final
 402 operation eliminates a single value (namely $1-a$ from $D(X_i)$ leading to a wipe-out at X_i).

403 For any revise operation in R , eliminating b from $D(X_j)$, there is a corresponding arc
 404 (u, v) in G_P where v is the vertex $\langle X_j, b \rangle$ and u is the cause of the elimination of b from
 405 $D(X_j)$. By the above argument, we can assume that each revise operation in R corresponds
 406 to a single elimination and hence a single arc in G_P . Let G_R be the subgraph of G_P consisting
 407 of the arcs corresponding to the operations of R . Let X_r be again the variable at which a
 408 wipe-out occurs at the end of R . For each $a \in D(X_r)$, in G_R there must be a directed path



■ **Figure 3** Left: a Boolean binary CSP P (the domain of every variable is $\{0, 1\}$). Right: the causal graph G_P of the proof of Theorem 11. The shortest explanation involves the two paths in red originating from X_1^s and corresponds to the sequence $\langle X_2 \leftarrow X_1, X_3 \leftarrow X_2, X_4 \leftarrow X_3, X_4 \leftarrow X_2 \rangle$.

409 P_a from a source-variable vertex to $\langle X_r, a \rangle$. By minimality, the set of arcs of G_R is the union
 410 of the set of arcs of P_a ($a \in D(X_r)$). Since each elimination has a unique cause (given by
 411 the arc corresponding to the revise operation in R producing the elimination), the in-degree
 412 of each vertex in G_R is at most one. Furthermore, source-variable vertices have in-degree 0.
 413 It follows that P_0 and P_1 can only possibly share arcs along an initial common subpath.

414 If $D(X_r)$ is a singleton $\{a\}$, then G_R must be a shortest path in G_P from a source-variable
 415 vertex to $\langle X_r, a \rangle$ and hence can be found in polynomial time by a standard shortest-path
 416 algorithm. So now suppose that $D(X_r) = \{0, 1\}$. If the set of edges of P_0 and P_1 are disjoint
 417 then P_0 and P_1 must both be shortest paths in G_P from source-variable vertices to $\langle X_r, 0 \rangle$
 418 and $\langle X_r, 1 \rangle$, respectively. If P_0 and P_1 have an initial common subpath, then they must
 419 diverge at some vertex v of G_P , the common initial subpath is a shortest path in G_P from a
 420 source-variable vertex to v and the remaining divergent paths P'_0 and P'_1 are shortest paths
 421 from v to $\langle X_r, 0 \rangle$ and $\langle X_r, 1 \rangle$, respectively. By an exhaustive search over the $O(n)$ vertices v
 422 of G_P , we can determine the paths P_0 and P_1 in polynomial time. ◀

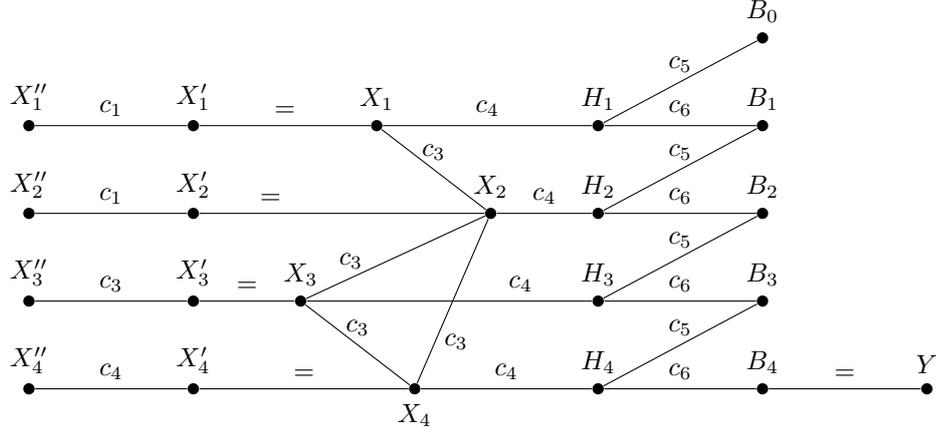
423 It is interesting to note that in the proof of Theorem 11, one of the paths P'_0, P'_1 may
 424 actually be empty. In this case, G_R is a path (either P_0 or P_1). This occurs if the elimination
 425 of a from $D(X_r)$ triggers a sequence of revise operations that leads to the elimination of $1-a$
 426 from $D(X_r)$. Another interesting point is that if P'_0, P'_1 are both non-empty, then the revise
 427 operations corresponding to P'_1 can all be inverted (i.e. each $X_i \leftarrow X_j$ becomes $X_j \leftarrow X_i$)
 428 and their order reversed in R to produce an alternative shortest arc inconsistency proof \tilde{R}
 429 which ends in a wipe-out at the variable X_k at which P'_0 and P'_1 diverged. For instance,
 430 the sequence $\langle X_2 \leftarrow X_1, X_3 \leftarrow X_2, X_4 \leftarrow X_3, X_2 \leftarrow X_4 \rangle$ is also a shortest explanation in
 431 the example of Figure 3. In this case, $G_{\tilde{R}}$ is a path (obtained in the example by using the
 432 edge in blue ($\langle X_4, 0 \rangle, \langle X_2, 1 \rangle$) instead of ($\langle X_2, 0 \rangle, \langle X_4, 1 \rangle$)). Hence, we can optimise since the
 433 exhaustive search over vertices v is unnecessary.

434 4.2 Intractability on domains with three values

435 ▶ **Theorem 12.** SHORTEST ARC INCONSISTENCY EXPLANATION is NP-hard, even on binary
 436 normalized networks with all domains of size at most three.

437 ▶ **Lemma 13.** Deciding whether there exists an arc inconsistency explanation of length
 438 smaller than or equal to k is NP-complete, even on binary normalized networks with all
 439 domains of size at most three.

XX:12 Complexity of Minimum-Size Arc-Inconsistency Explanations



■ **Figure 4** The constraint network P_G in the proof of Lemma 13 when looking for a dominating set in the graph $G = (\{1, 2, 3, 4\}, \{(1, 2), (2, 4), (2, 3), (3, 4)\})$

440 **Proof. Membership.** As in Lemma 7.

441 *Completeness.* We reduce the DOMINATING SET problem (whether a graph G has a
 442 dominating set of size at most k) to the problem of deciding whether there is an arc
 443 inconsistency explanation of length at most $4n + k + 1$ for a CSP instance. Let $G = (V, E)$
 444 be a graph with $V = \{1, \dots, n\}$.

We construct the CSP instance P_G with $5n + 2$ variables

$$X = \{X_1, \dots, X_n, X'_1, \dots, X'_n, X''_1, \dots, X''_n, H_1, \dots, H_n, B_0, \dots, B_n, Y\}$$

445 all with domain $\{0, 1, 2\}$ except B_0 whose domain is $\{0\}$ and Y whose domain is $\{2\}$.

446 We build the set of constraints

$$\begin{aligned} 447 \quad C = & \{c_1(X''_i, X'_i) : i \in [1, n]\} \cup \{c_2(X'_i, X_i) : i \in [1, n]\} \\ 448 & \cup \{c_3(X_i, X_j) : \{i, j\} \in E\} \cup \{c_4(X_i, H_i) : i \in [1, n]\} \\ 449 & \cup \{c_5(B_{i-1}, H_i) : i \in [1, n]\} \cup \{c_6(H_i, B_i) : i \in [1, n]\} \cup \{B_n = Y\} \end{aligned}$$

where

$$\begin{aligned} c_1(X''_i, X'_i) &= \{(0, 0)\} \\ c_2(X'_i, X_i) &= \{(0, 0), (1, 1), (2, 2)\} \\ c_3(X_i, X_j) &= \{0, 1, 2\} \times \{0, 1, 2\} \setminus \{(0, 2), (2, 0)\} \\ c_4(X_i, H_i) &= \{0, 1, 2\} \times \{0, 1, 2\} \setminus \{(0, 1), (1, 1)\} \\ c_5(B_{i-1}, H_i) &= \{0, 1, 2\} \times \{0, 1, 2\} \setminus \{(0, 2), (1, 2)\} \\ c_6(H_i, B_i) &= \{0, 1, 2\} \times \{0, 1, 2\} \setminus \{(0, 2)\} \end{aligned}$$

450 The constraint network is shown in Figure 4 for a graph with $n = 4$ vertices and 4 edges.

451 We first prove that if G contains a k -dominating set, then there exists an arc inconsistency
 452 explanation of length $4n + k + 1$ for P_G . Assume that the set of vertices S is a k -dominating
 453 set. We build the sequence R of *revise()* operations in the following way. The first k elements
 454 in R are $X'_i \stackrel{c_1}{\leftarrow} X''_i$ for each vertex i in S . The k next elements in R are $X_i \stackrel{c_2}{\leftarrow} X'_i$, again
 455 for vertices i in S . After those $2k$ *revise()* operations, for all i in S , $D(X_i) = \{0\}$. Then,

456 for each vertex j in $V \setminus S$, R contains $X_j \stackrel{c_3}{\leftarrow} X_i$, where $i \in S$ and $\{i, j\} \in E$. We know
 457 such a vertex i exists for each j because S is a dominating set. After those additional
 458 $n - k$ *revise()* operations, for all i not in S , $D(X_i) = \{0, 1\}$. The n next elements in R are
 459 $H_i \stackrel{c_4}{\leftarrow} X_i$, removing value 1 from $D(H_i)$ because $2 \notin D(X_i)$. The $2n$ next elements in R are
 460 $\langle H_i \stackrel{c_5}{\leftarrow} B_{i-1}, B_i \stackrel{c_6}{\leftarrow} H_i \rangle$ in increasing order of i from 1 to n . Each $H_i \stackrel{c_5}{\leftarrow} B_{i-1}$ removes value
 461 2 from $D(H_i)$ if $2 \notin D(B_{i-1})$ and $B_i \stackrel{c_6}{\leftarrow} H_i$ removes value 2 from $D(B_i)$ if $1, 2 \notin D(H_i)$.
 462 As $B_0 = 0$ and value 1 has already been removed from all H_i 's domains, those $2n$ *revise()*
 463 remove value 2 from the domain of all B_i . Finally, after these $2k + (n - k) + n + 2n = 4n + k$
 464 *revise()* operations, the last element in R , $Y \stackrel{c_7}{\leftarrow} B_n$, wipes out the domain of Y and proves
 465 arc inconsistency.

466 (Sketch.) We then prove that if there exists an arc inconsistency explanation for P_G
 467 of length $4n + k + 1$, then G contains a k -dominating set. We first observe that if we
 468 remove $c_5(B_0, H_1)$ or $B_n = Y$ from P_G , the instance becomes satisfiable. (B_0 is necessary
 469 to trigger removals of value 2 from the H_i s and Y to trigger removals of value 0.) Hence,
 470 no wipe out can occur without executing $2n + 1$ *revise()* operations on the path from B_0
 471 to Y . Furthermore, if a single variable H_i still has value 1 in its domain, the propagation
 472 of removals stops. As a result, value 2 needs to be removed from all X_i s and a *revise()*
 473 needs to be executed on the n constraints c_4 . We then have $n + k$ remaining available
 474 operations to remove value 2 from all X_i s. If we do these removals thanks to the sequence
 475 $\langle X_i' \stackrel{c_1}{\leftarrow} X_i'', X_i \stackrel{c_2}{\leftarrow} X_i' \rangle$, it costs $2n$ operations, which is more than $n + k$. To reach $n + k$, we
 476 need to remove value 2 in a single operation for at least $n - k$ variables. The only way to
 477 do that is through a $X_j \stackrel{c_3}{\leftarrow} X_i$ for $n - k$ variables X_j . Now, $X_j \stackrel{c_3}{\leftarrow} X_i$ removes value 2 from
 478 $D(X_j)$ only if $D(X_i) = \{0\}$ and $c_3(X_i, X_j) \in C$. $D(X_i)$ is equal to $\{0\}$ only if X_i is one of
 479 the k variables on which $\langle X_i' \stackrel{c_1}{\leftarrow} X_i'', X_i \stackrel{c_2}{\leftarrow} X_i' \rangle$ has been executed. $c_3(X_i, X_j)$ belongs to C
 480 only if $\{i, j\} \in E$. As a result, the set of k vertices i corresponding to the k variables with
 481 $D(X_i) = \{0\}$ is a dominating set. ◀

482 5 Conclusion

483 We have investigated the complexity of finding a shortest proof of inconsistency of a binary
 484 CSP in the form of a sequence of arc consistency operations. Our characterisation in terms
 485 of structure or domain size shows that this problem is polynomial when variables have degree
 486 two or domains are Boolean. The problem is NP-hard if the CSP has four variables of degree
 487 three or if the domain size is bounded by three. It is also NP-hard on trees. In addition,
 488 the problem is not FPT-approximable unless the Gap-ETH conjecture is false. Although
 489 our initial motivation was to provide short explanations for human users, there are other
 490 possible applications. Virtual Arc Consistency (VAC) algorithms for cost-function networks
 491 use arc-inconsistency explanations in the CSP of zero-cost tuples in order to update cost
 492 functions [5]. Our NP-hardness results can be seen as a justification for the use of minimal
 493 rather than minimum-cardinality arc-inconsistency explanations by VAC algorithms. On a
 494 final positive note, the polynomial-time algorithm for the special case of size-2 domains may
 495 prove an inspiration for heuristic methods to improve minimal arc inconsistency explanations
 496 via the search for shortest paths in the causal graph described in the proof of Theorem 11.

497 References

- 498 1 Jérôme Amilhastre, Hélène Fargier, and Pierre Marquis. Consistency restoration and ex-
 499 planations in dynamic CSPs application to configuration. *Artif. Intell.*, 135(1-2):199–234,
 500 2002.

XX:14 Complexity of Minimum-Size Arc-Inconsistency Explanations

- 501 2 Bart Bogaerts, Emilio Gamba, and Tias Guns. A framework for step-wise explaining how to
502 solve constraint satisfaction problems. *Artif. Intell.*, 300:103550, 2021.
- 503 3 Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi,
504 Danupon Nanongkai, and Luca Trevisan. From Gap-ETH to FPT-inapproximability: Clique,
505 dominating set, and more. In Chris Umans, editor, *Proceedings of the 58th IEEE Annual*
506 *Symposium on Foundations of Computer Science (FOCS'17)*, pages 743–754. IEEE Computer
507 Society, 2017.
- 508 4 Yijia Chen, Martin Grohe, and Magdalena Grüber. On parameterized approximability. In
509 Hans L. Bodlaender and Michael A. Langston, editors, *Parameterized and Exact Computation,*
510 *Second International Workshop, IWPEC*, volume 4169 of *Lecture Notes in Computer Science*,
511 pages 109–120. Springer, 2006. doi:10.1007/11847250_10.
- 512 5 Martin C. Cooper, Simon de Givry, Martí Sánchez-Fibla, Thomas Schiex, and Matthias
513 Zytnicki. Virtual arc consistency for weighted CSP. In Dieter Fox and Carla P. Gomes, editors,
514 *AAAI 2008*, pages 253–258. AAAI Press, 2008.
- 515 6 Irit Dinur. Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover.
516 *Electronic Colloquium on Computational Complexity*, page 128, 2016.
- 517 7 M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of*
518 *NP-Completeness (Series of Books in the Mathematical Sciences)*. W. H. Freeman, first edition
519 edition, 1979.
- 520 8 Allen Van Gelder. Verifying RUP proofs of propositional unsatisfiability. In *ISAIM*, 2008.
- 521 9 E. Goldberg and Y. Novikov. Verification of proofs of unsatisfiability for CNF formulas. In
522 *2003 Design, Automation and Test in Europe Conference and Exhibition*, pages 886–891, 2003.
523 doi:10.1109/DATE.2003.1253718.
- 524 10 Alexey Ignatiev, Nina Narodytska, and Joao Marques-Silva. Abduction-based explanations
525 for machine learning models. *Proceedings of the AAAI Conference on Artificial Intelligence*,
526 33(01):1511–1519, 2019. doi:10.1609/aaai.v33i01.33011511.
- 527 11 Ulrich Junker. QUICKXPLAIN: preferred explanations and relaxations for over-constrained
528 problems. In Deborah L. McGuinness and George Ferguson, editors, *Proceedings of the*
529 *Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative*
530 *Applications of Artificial Intelligence*, pages 167–172. AAAI Press / The MIT Press, 2004.
- 531 12 Ulrich Junker. Configuration. In Francesca Rossi, Peter van Beek, and Toby Walsh, editors,
532 *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*, pages
533 837–873. Elsevier, 2006.
- 534 13 R. Karp. Reducibility among combinatorial problems. In R. Miller and J. Thatcher, editors,
535 *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- 536 14 Pasin Manurangsi and Prasad Raghavendra. A birthday repetition theorem and complexity of
537 approximating dense CSPs. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca
538 Muscholl, editors, *Proceedings of the 44th International Colloquium on Automata, Languages,*
539 *and Programming (ICALP'17)*, volume 80 of *LIPICs*, pages 78:1–78:15. Schloss Dagstuhl -
540 Leibniz-Zentrum für Informatik, 2017.
- 541 15 Dániel Marx. Completely inapproximable monotone and antimonotone parameterized problems.
542 In *Proceedings of the 25th IEEE Annual Conference on Computational Complexity*, pages
543 181–187, 2010. doi:10.1109/CCC.2010.25.
- 544 16 Andy Shih, Arthur Choi, and Adnan Darwiche. A symbolic approach to explaining bayesian
545 network classifiers. *IJCAI'18*, page 5103–5111. AAAI Press, 2018.
- 546 17 Mohammed H. Sqalli and Eugene C. Freuder. Inference-based constraint satisfaction supports
547 explanation. In William J. Clancey and Daniel S. Weld, editors, *AAAI 96, IAAI 96, Volume*
548 *1*, pages 318–325. AAAI Press / The MIT Press, 1996.