

Causality Analysis for Concurrent Reactive Systems

Rayna Dimitrova¹, Rupak Majumdar², and Vinayak S. Prabhu³

¹ University of Leicester
rd307@leicester.ac.uk

² MPI-SWS

rupak@mpi-sws.org

³ Colorado State University
vinayak@mpi-sws.org

Abstract

We present a comprehensive language theoretic causality analysis framework in the setting of concurrent reactive systems. Our framework allows us to uniformly express a number of causality notions studied in the areas of artificial intelligence and formal methods, as well as define new ones that are of potential interest in these areas. Furthermore, our formalization provides means for reasoning about the relationships between individual notions which have mostly been considered independently in prior work; and allows us to judge appropriateness of the different definitions for various applications in system design. In particular, we consider causality analysis notions for debugging, error resilience, and for liability resolution in concurrent reactive systems. Finally, we derive automata based algorithms for computing various causal sets based on our language theoretic encoding, and derive the algorithmic complexities.

1 Introduction

Causality analysis, which investigates questions of the form “Does event e_1 cause event e_2 ?” plays an important role in many areas of science, medicine and law. In formal methods, causality analysis has been used to determine the *coverage of specifications* [5] (that is, which parts of the system under scrutiny are relevant for the satisfaction of a specification), to *explain counterexamples* [2] (identify points in a counterexample trace that are relevant for the failure of a temporal specification), to fault tree construction [18], and to *automatically refine system abstractions* [4]. In artificial intelligence, causality-based explanation finding has applications in natural language processing, automated medical diagnosis, vision processing, and planning. Resolving liabilities in a legal setting often relies on establishing the causal relations between potential causes and the occurred damage [3].

Causality definitions based on *counterfactuals*, which are alternative scenarios where the suspected cause e_1 of e_2 did not happen, date back to [13] and have been extensively studied in philosophy (cf. [19]). In computer science, the most prominent and widely used definition of causality is that of [12], in which the authors write “... while it is hard to argue that our definition (or any other definition, for that matter) is the right definition, we show that it deals with the difficulties that have plagued other approaches in the past ...”. Halpern and Pearl’s approach is based on *structural equations*, which describe causal dependencies between Boolean variables. We extend the Boolean study of causality to the *temporal* setting; specifically, we formalize notions of causality in *concurrent reactive systems* whose behaviors evolve over time. A concurrent reactive system is a composition of interacting components; the system behavior is determined by the *repeated* interaction between the components over time. We consider the setting where component implementations are not available for analysis and the designer can only rely on specifications of their expected behavior. Thus, when analyzing an *error trace* (an execution of the system that violates a desired system-level property), the only available information about the system is the components’ specifications and the observed trace.

Recently causality analysis of component-based systems has drawn a lot of attention [8, 7, 21, 6, 9, 23, 10]. The goal is to identify a subset of the components, which have violated their *local* specifications, that are actually responsible for the violation of the *system-level* property. This requires

integrating the temporal order of events [16, 17, 1] in the analysis of logical causality. The main challenge lies in defining the set of *counterfactual traces* for a given observed trace \mathbf{tr} . These are traces used to reason about hypothetical scenarios where a subset of the system components behave in a way that differs from the trace \mathbf{tr} . Different approaches differ in the way they account for the dependencies between the behaviors of different components, that is, how changing the behavior of one component affects the behavior of others. The available information, a single observed system trace and the components' specifications, is often insufficient to faithfully reconstruct these alternative behaviors. Existing approaches, hence, choose a specific set of trace reconstruction rules as a basis of their causality notion. However, the suitability of a notion depends on the desired application. For example, while liability resolution requires conservative notions that give high confidence in their determinations of causes of failure, for system analysis and debugging less conservative notions are more appropriate, provided that they are cost-effective and focus on relevant components. One of the limitations of existing work in this area is that the various causality notions have been studied in isolation but no framework for comparing different notions of causality has been provided so far.

We present a language theoretic causality framework for concurrent reactive systems incorporating diverse counterfactual trace sets. A cause for a violation is a component set. Our analysis reasons about two classes of scenarios to determine if component set \mathcal{C} is a cause (for an observed system fault):

- *Fault Mitigation Capability* analysis asks whether the *correct behavior* of the component set \mathcal{C} is enough to mitigate the faults of all components (*including those of components not in \mathcal{C}*), by ensuring that the required system property holds.
- *Fault Manifestation* analysis asks whether the observed *faulty behavior* of the component set \mathcal{C} is enough to manifest a global fault (*i.e.*, a system-property violating global behavior), even if the components not in \mathcal{C} were to behave correctly.

These two classifications parallel the classifications of [8, 7] of causes into *necessary causes* and *sufficient causes*. However, our analysis is not limited to specific definitions of counterfactual sets. In contrast, we provide a reasoning framework based on generic counterfactual sets, and introduce several natural instantiations.

We will use the following example throughout the paper to illustrate the key notions of our framework for causality analysis.

Example 1. Consider a system with three components, C_1 , C_2 and C_3 , with a common shared resource. Access to the resource is regulated by C_3 , and there are in total M units of the resource available per unit of time. In particular, consider a solar battery for which the charge rate is M energy units per time unit. If the initial charge is $E > M$, then the components cannot utilize more than M units of energy for more than $E - M$ steps. Thus, to be safe, we require that in each step the combined consumption should be at most M . This system-level requirement is denoted as φ (in a concrete execution, however, components can consume more than the allowed M units for a small number of steps without dire consequences).

With a view towards robust satisfaction of φ , the local specifications of the components constrain their behaviors further than what is absolutely necessary for the satisfaction of this property. For example, let the specification φ_3 of C_3 require that the resource allocation respects a given safety margin, namely, that the combined allocation by C_3 to all the components should not be more than $M - 1$ units in any step. Furthermore suppose that φ_3 specifies that, if component C_1 or C_2 performs a violation, that is, consumes more units of energy than it has been allocated, then C_3 should attempt to compensate for that by reducing its own consumption. More concretely, if at time t , component C_3 indicates that C_1 should consume at most Δ units in the next time instant $t + 1$, and at $t + 1$ component C_1 consumes $\Delta + \alpha$ instead, then C_3 must decrease its own consumption at time step $t + 2$ (from the consumption at time $t + 1$) by α , if possible (or reduce it to 0 if not), in an attempt to prevent a violation of the global property φ . Note that there is a delay of one time unit for components to react to their inputs. Let the only requirement on C_1 (and also on C_2) be that if C_3 allocates it a resource units in the current step, then it will consume at most a units in the next step. Consider the following trace (with the global requirement

φ that the combined resource consumption be at most 31), where C_3 allocates at most 30 units in total to all components:

step	1	2	3	4	
allocated to C_1 for next time step	6	6	6	6	Observe that the combined consumption is 33 from step 3 onwards, violating the limit 31. C_2 exceeds its limit by 6 units at the first step, and by 8 units after that. C_3 is supposed to decrease its consumption from 7 units in step 2, to 1 unit in step 3 (and then to 0 units in step 4); but
consumption C_1 at current time step	0	6	6	6	
allocated to C_2 for next time step	12	12	12	12	
consumption C_2 at current time step	0	18	20	20	
allocated to C_3 for current time step	0	7	7	7	
consumption C_3 at current time step	0	7	7	7	

it violates its local specification and does not do this.

Had C_3 reduced its consumption as required, the global violation would not have occurred, even in the presence of C_2 's observed incorrect behavior. A causality analysis should report the component set $\{C_3\}$ as one of the possible causes for the observed violation.

Both singleton sets $\{C_2\}$ and $\{C_3\}$ of components have the capability to mitigate the observed error, i.e. the correct behavior of either component would have prevented the violation of the global requirement φ . Dually, both components C_2 and C_3 , i.e. the component set $\{C_2, C_3\}$, have to behave incorrectly as observed in order for the trace to manifest the observed error. \square

Contributions. We present a systematic, language theoretic study of causality for component-based concurrent reactive systems:

- We first describe a modular decomposition of counterfactual tracesets based on (i) hypotheses on possible incorrect behaviors (differing from the single observed trace); and (ii) interactions between different components due to the concurrent reactive nature of the components.
- Next, we show how composed counterfactual tracesets can be used to define various notions of causality in a uniform fashion (Equations (1) through (4)). Our approach uses basic language theoretic operations to reason about intricate *consistency* issues: issues which arise when repeated component interactions have to be reasoned about (e.g., two components are faulty, we repair one, and this leads to a different sequence of inputs to the unfixed faulty one).
- We demonstrate that the generality and modularity of our definition of causality allow us to seamlessly extend causality analysis to the case of heterogeneous fault models, where different components are examined under different fault scenarios.
- Our unified approach allows us to compare the resulting different causality notions, and the relationships between the causal sets, and thus to indicate the situations in which each of them is most appropriate.
- We present an automata-based method for determining various causal sets in the setting of heterogeneous component-fault models, and derive its algorithmic complexity.

2 Preliminaries

Languages. Let Σ be a non-empty finite alphabet. A *word* or a *trace* $w = \sigma_1, \sigma_2, \dots, \sigma_m$ over Σ is a finite sequence of letters from Σ . We denote by $w_{[i]}$ the i -th symbol σ_i in the word w , and by $w[i..j]$ the substring $\sigma_i, \dots, \sigma_j$. Σ^* is the set of all words over Σ , and ϵ is the empty word. A *language* is a set of words. The concatenation of two words u, w is denoted $u \cdot w$; and similarly for languages. For a word w , $\text{len}(w)$ is the length of w . For a language L and a positive integer k , let $L_{[k]}$ denote the words in L which have exactly k letters. A word u is a *prefix* of a word v , denoted $u \leq v$, iff there exists a word w such that $v = u \cdot w$. For a language L , the language $\text{Pref}(L)$ consists of the prefixes of words in L . We write $u < v$ when u is a strict prefix of v , that is $u \leq v$ and $u \neq v$. Given two words u, v , let $\text{lccmpref}(u, v)$ denote the longest common prefix of u and v . A language L is said to be *prefix closed* if whenever a word $v \in L$,

we have that every prefix of v is also in L , i.e., $\text{Prefs}(L) \subseteq L^1$.

Languages over Variables. For the purpose of modelling reactive systems in which components communicate via shared variables, we let an alphabet be a set of possible valuations of a set of variables over a given finite set. If an alphabet Σ is defined over a set of variables X we denote this by $\Sigma[X]$, omitting $[X]$ when X is clear from the context. Thus, a letter $\sigma \in \Sigma[X]$ is a function $\sigma : X \mapsto \bigcup_{x \in X} \mathcal{D}_\Sigma(x)$ where $\mathcal{D}_\Sigma(x)$ is the (finite) domain of the variable x . A word w over $\Sigma[X]$ is a sequence of valuations for X , i.e. every letter $w_{[i]}$ is a valuation of all variables in X . If $\Sigma[X]$ and $\Pi[Y]$ are alphabets with $Y \subseteq X$, and $w \in \Sigma^*$, then $w|_\Pi$ is the projection of w on Π defined in the usual way.

Alphabet and Language Composition. Given $\Sigma_1[X_1], \dots, \Sigma_n[X_n]$ for which we have that for every variable x such that $x \in X_i$ and $x \in X_j$ for $i \neq j$, we have $\mathcal{D}_{\Sigma_i}(x) = \mathcal{D}_{\Sigma_j}(x)$ (i.e. common variables have the same domain in each alphabet), we define the *composite alphabet* $\Sigma_1[X_1] \parallel \dots \parallel \Sigma_n[X_n]$ to be the alphabet $\Sigma[X]$ such that $X = \bigcup_{i=1}^n X_i$; such that the domain of a variable x is $\mathcal{D}_\Sigma(x) = \mathcal{D}_{\Sigma_i}(x)$ for $x \in X_i$. Given languages $L_i \subseteq \Sigma_i^*$ over Σ_i for $i = 1, \dots, n$, we define the *language composition* of L_1, \dots, L_n to be the language: $L_1 \parallel \dots \parallel L_n = \{w \in \Sigma^* \mid w|_{\Sigma_i} \in L_i \text{ for all } i\}$, over $\Sigma[X]$.

Example 2 (Languages and composition). We consider a language L_1 over an alphabet $\Sigma_1[X_1]$, with $X_1 = \{x_1, x_2\}$ and domain $\mathcal{D}_{\Sigma_1}(x_1) = \mathcal{D}_{\Sigma_1}(x_2) = \{0, 1\}$; and a language L_2 over an alphabet $\Sigma_2[X_2]$, with $X_2 = \{x_2, x_3\}$, also with Boolean domain. The language L_1 is:

$$\{(x_1: b_1^1, x_2: b_2^1), (x_1: b_1^2, x_2: b_2^2), \dots, (x_1: b_1^m, x_2: b_2^m) \mid b_1^j = b_2^j \text{ for all } 1 \leq j \leq m\}$$

i.e. words where the values of x_1 and x_2 are the same. L_2 consists of words with equal valued x_2 and x_3 :

$$\{(x_2: b_2^1, x_3: b_3^1), (x_2: b_2^2, x_3: b_3^2), \dots, (x_2: b_2^m, x_3: b_3^m) \mid b_2^j = b_3^j \text{ for all } 1 \leq j \leq m\}.$$

The language $L_1 \parallel L_2$ over $\Sigma[\{x_1, x_2, x_3\}]$ is defined as:

$$\left\{ \begin{array}{l} (x_1: b_1^1, x_2: b_2^1, x_3: b_3^1), (x_1: b_1^2, x_2: b_2^2, x_3: b_3^2), \dots \\ (x_1: b_1^m, x_2: b_2^m, x_3: b_3^m) \end{array} \mid \begin{array}{l} m \geq 0 \text{ and} \\ b_1^j = b_2^j = b_3^j \in \{0, 1\} \text{ for all } 1 \leq j \leq m \end{array} \right\}$$

i.e. words where x_1, x_2, x_3 have the same value at each step. \square

Component Model. A *component specification* is a tuple $C = (X, \text{inp}(X), \text{out}(X), \Sigma, \varphi)$, where

- $X = \text{inp}(X) \uplus \text{out}(X)$ is the set of variables of the component, consisting of the input variables $\text{inp}(X)$ and the output variables $\text{out}(X)$ (the sets of input and output variables being disjoint);
- Σ is the alphabet, consisting of all possible valuations of the variables X ;
- φ is a non-empty prefix-closed language over Σ , specifying the set of correct behaviours of C .

For a letter $\sigma \in \Sigma$ and a variable $x \in X$ we denote with $\sigma(x)$ the value of x according to σ . The input alphabet $\text{inp}(\Sigma)$ of C consists of the possible valuations of the input variables $\text{inp}(X)$, and, similarly, the output alphabet $\text{out}(\Sigma)$ consists of valuations of $\text{out}(X)$.

Example 3 (Component). Component C_1 from the example in the introduction can be modelled as $C_1 = (X_1, \text{inp}(X_1), \text{out}(X_1), \Sigma_1, \varphi_1)$, where² $X_1 = \{x_a^{3,1}, x_d^{1,3}\}$, and $\text{inp}(X_1) = \{x_a^{3,1}\}$, and $\text{out}(X_1) = \{x_d^{1,3}\}$; the alphabet Σ_1 consists of the possible valuations of $x_a^{3,1}$ and $x_d^{1,3}$ ranging over $[0, M]$, and φ_1 contains strings $w \in \Sigma_1^*$ such that either (i) w is the empty string; or (ii) $w_{[1]}(x_d^{1,3}) = 0$ and $w_{[j+1]}(x_d^{1,3}) \leq w_{[j]}(x_a^{3,1})$ for all $2 \leq j+1 \leq \text{len}(w)$. Intuitively, the value of $x_a^{3,1}$ specifies the units of resource allocated to C_1 by C_3 for the next step, the value of $x_d^{1,3}$ specifies the units depleted by C_1 in the current step (this number is given as input to C_3). The specification φ_1 ensures that at each step C_1 's consumption does not exceed the bound specified by the value of $x_a^{3,1}$ in the previous step. \square

¹Prefix-closed languages need not be regular

²In naming variables in the examples, we follow the convention that a variable x^{k,i_1,i_2,\dots,i_j} (i) is common to the components $C_k, C_{i_1}, \dots, C_{i_j}$; and (ii) is an output variable of component C_k , and an input variable to components C_{i_1}, \dots, C_{i_j} .

Component Compositions, Systems, & Global Specifications. Given a set of components $\mathcal{C} = \{C_1, \dots, C_n\}$ where each $C_i = (X_i, \text{inp}(X_i), \text{out}(X_i), \Sigma_i, \varphi_i)$, the component composition $C_1 \parallel \dots \parallel C_n$ is defined in case the following two conditions both hold.

1. The sets of output variables are pairwise disjoint, i.e., if $\text{out}(X_i) \cap \text{out}(X_j) = \emptyset$ for $i \neq j$; and
2. the composite alphabet $\Sigma_1[X_1] \parallel \dots \parallel \Sigma_n[X_n]$ exists.

The composition $C_1 \parallel \dots \parallel C_n$ is the component $(X_{\mathcal{C}}, \text{inp}(X_{\mathcal{C}}), \text{out}(X_{\mathcal{C}}), \Sigma_{\mathcal{C}}, \varphi_{\mathcal{C}})$ defined as follows

- $X_{\mathcal{C}} = \bigcup_{i=1}^n X_i$ is the set of all variables;
- $\text{out}(X_{\mathcal{C}}) = \bigcup_{i=1}^n \text{out}(X_i)$, i.e., the set $\text{out}(X_{\mathcal{C}})$ consist of all output variables of all components.
- $\text{inp}(X_{\mathcal{C}}) = \left(\bigcup_{i=1}^n \text{inp}(X_i) \right) \setminus \text{out}(X_{\mathcal{C}})$ i.e., the set of input variables contains those input variables which are are not output variables of any component in \mathcal{C} .
- $\Sigma_{\mathcal{C}}$ is the composite alphabet $\Sigma_1[X_1] \parallel \dots \parallel \Sigma_n[X_n]$.
- $\varphi_{\mathcal{C}}$ is the composite language $\varphi_1 \parallel \dots \parallel \varphi_n$.

A collection of composable components is called a *system*. Given a system $\mathcal{S} = \{C_1, \dots, C_n\}$, a (global) system *specification* φ is a language over the composite alphabet $\Sigma_1 \parallel \dots \parallel \Sigma_n$. In this work, we require that the system specification φ be prefix closed, and in addition, that φ contains $\varphi_1 \parallel \dots \parallel \varphi_n$. Thus, the global requirement is more relaxed than the promised behaviors of the individual components. In other words, the system $\{C_1, \dots, C_n\}$ promises to *implement* or *refine* the global requirement φ . Abusing notation, we let \mathcal{S} also denote the component composition $C_1 \parallel \dots \parallel C_n$.

Note. The composition of components as defined above implies that components execute synchronously in lock-step³. All definitions and results presented in this paper can be easily extended to the asynchronous setting, which we do not do here for the sake of simplicity of the presentation.

System Traces. A *global trace* of \mathcal{S} is a word $\mathbf{tr} \in \Sigma^*$. The trace \mathbf{tr} is *correct* if $\mathbf{tr} \in \varphi$; otherwise it is an *error trace*. A *local trace* for component C_i is a word $w \in \Sigma_i^*$.




Figure 1: Resource sharing system described in Example 4.

Example 4 (Systems and Traces). We define component

$C_2 = (X_2, \text{inp}(X_2), \text{out}(X_2), \Sigma_2, \varphi_2)$ analogously to C_1 in our running example. Let component $C_3 = (X_3, \text{inp}(X_3), \text{out}(X_3), \Sigma_3, \varphi_3)$, where $\text{inp}(X_3) = \{x_d^{1,3}, x_d^{2,3}\}$, and $\text{out}(X_3) = \{x_a^{3,1}, x_a^{3,2}, x_d^3\}$. The alphabet Σ_3 consists of the possible valuations of the variables in $X_3 = \text{inp}(X_3) \cup \text{out}(X_3)$, the range of each variable being $[0, M]$. The variables $x_d^{1,3}, x_d^{2,3}$ denote the current-step depletions of the resource by C_1 and C_2 respectively, the values of which are read in by component C_3 . The value of x_d^3 is the depletion by C_3 . The variables $x_a^{3,1}, x_a^{3,2}$ are the allocations of the resource to C_1 and C_2 for the next step. The system is depicted in Figure 1.

The local specification φ_3 of C_3 is defined as containing all words $w \in \Sigma_3^*$ which satisfy each of the following four requirements (letting $\sigma_j = w_{[j]}$, and $\sigma_{j+1} = w_{[j+1]}$, and $\sigma_{j+2} = w_{[j+2]}$),

1. φ_3'''' specifies that for every $j \leq \text{len}(w)$ we have $\sigma_j(x_a^{3,1}) + \sigma_j(x_a^{3,2}) + \sigma_{j+1}(x_d^3) \leq M - 1$; i.e., the planned combined depletion at step $j + 1$ should be at most $M - 1$ (leaving a safety margin of 1).
2. φ_3'''' specifies that $w_{[1]}(x_d^3) = 0$, i.e. C_3 should not deplete the resource at all in the first step.
3. φ_3'' specifies how component C_3 should change its behavior at step $j + 2$ based on C_2 's behavior at step $j + 1$. It requires one of the following conditions to hold.
 - $\sigma_{j+1}(x_d^{2,3}) \leq \sigma_j(x_a^{3,2})$, i.e., the depletion by C_2 in step $j + 1$ is at most what it was allocated to it by C_3 in the previous step.
 - If $\sigma_{j+1}(x_d^{2,3}) > \sigma_j(x_a^{3,2})$, i.e., if the previous case does not hold, then:
$$\sigma_{j+2}(x_d^3) \leq \max \left(0, \sigma_{j+1}(x_d^3) - \left(\sigma_{j+1}(x_d^{2,3}) - \sigma_j(x_a^{3,2}) \right) \right).$$

³The components in our examples are such that in every step of the system execution each component reads the values of its input variables that were written in the previous step and updates its output variables. This delay is not imposed by the formal model.

That is, if C_2 exceeds its allocation by an amount α at step $j + 1$, then C_3 reduces its own consumption at step $j + 2$ from that at step $j + 1$ by α (if possible).

4. φ'_3 is the condition analogous to φ'_3 for component C_1 , and specifies how component C_3 should change its behavior based on C_1 's behavior.

The system specification of $C_1 \parallel C_2 \parallel C_3$ is defined to be the language φ containing words w such that for every j the combined depletion is at most M . Formally, φ equals:

$$\{w \in \Sigma^* \mid \text{for all } j, \text{ we have } w_{[j]}(x_d^3) + w_{[j]}(x_d^{1,3}) + w_{[j]}(x_d^{2,3}) \leq M\}$$

We present two sample traces, letting $M = 31$ (the global specification φ is that at each step, the combined depletion for that step must not exceed 31). The first trace satisfies φ and is as follows.

Note that even though C_2 violates its local spec φ_2 in steps 2 and 3 (as it depletes by 16 units when it was allowed only 10 as specified by $x_a^{3,2}$ in steps 1 and 2), the global specification is still satisfied due to C_3 reducing its own depletion amount at step 3 from 10 (in the previous step) to 4.

j, φ	1, φ	2, φ	3, φ
local specs	$\varphi_1, \varphi_2, \varphi_3$	$\varphi_1, \neg\varphi_2, \varphi_3$	$\varphi_1, \neg\varphi_2, \varphi_3$
$x_a^{3,1}$ and $x_a^{3,2}$	10	10	10
$x_d^{1,3}$	0	4	10
$x_d^{2,3}$	0	16	16
x_d^3	0	10	4

The second trace given to the right violates φ . Component C_3 violates its local specification φ_3 at step 3 because it should have reduced its consumption (from that in step 2) by α where α is the amount by which the resource depletion by C_2 exceeded its allocated 10 units (in this case $\alpha = 6$ units).

j, φ	1, φ	2, φ	3, φ	4, $\neg\varphi$
local specs	$\varphi_1, \varphi_2, \varphi_3$	$\varphi_1, \neg\varphi_2, \varphi_3$	$\varphi_1, \neg\varphi_2, \neg\varphi_3$	$\varphi_1, \neg\varphi_2, \neg\varphi_3$
$x_a^{3,1}$	10	10	10	10
$x_a^{3,2}$	10	10	10	10
$x_d^{1,3}$	0	6	6	8
$x_d^{2,3}$	0	16	16	16
x_d^3	0	8	8	8

3 A Framework for Causality

In this section, we fix a system $\{C_1, \dots, C_n\}$, where $C_i = (\Sigma_i, \text{inp}(\Sigma_i), \text{out}(\Sigma_i), \varphi_i)$; a specification φ ; and an observed trace $\mathbf{tr} \notin \varphi^4$. We also fix a non-empty collection of components $\mathfrak{C} \subseteq \{C_1, \dots, C_n\}$ for causality analysis. Let $\overline{\mathfrak{C}}$ be the components not in \mathfrak{C} . Assume, w.l.o.g., $\mathfrak{C} = \{C_1, C_2, \dots, C_{n_{\mathfrak{C}}}\}$ (thus, $\overline{\mathfrak{C}} = \{C_{1+n_{\mathfrak{C}}}, \dots, C_n\}$). Let $\Sigma_{\mathfrak{C}}$ denote the composition of the alphabets of the components of \mathfrak{C} .

3.1 Counterfactual Traces & Faulty Behaviors

Counterfactual Traces. Informally, the set of counterfactual traces for a given observed trace \mathbf{tr} , consists of traces obtained from \mathbf{tr} by correcting the behavior of some faulty components. These traces are used to reason about hypothetical scenarios where a subset of the components behave (correctly) in a way that differs from the incorrect behavior in the observed trace \mathbf{tr} . Depending on the hypothetical scenario, the set of counterfactual traces is obtained as (a subset of) the composition of trace sets of individual component behaviors appropriately altered with respect to the trace \mathbf{tr} .

In *reactive systems*, the behaviors of individual components are intertwined; This results in consistency dependencies between the component behaviors that must be taken into account. As the effect of the change in behaviors of other components that affect a particular component C_i is not easily determined, there does not exist a unique definition of counterfactual tracesets for C_i that is applicable for all purposes. We present different constructions of counterfactual tracesets, and indicate the situations in which each is useful. In each of these constructions, the set of counterfactual traces for a component C_i

⁴Global system traces (obtained by composing the local traces of individual components) are denoted in bold font.

whose observed behavior in \mathbf{tr} is incorrect will include *some of the correct behaviors* of C_i (according to φ_i), as well as *some incorrect behaviors*. The latter are determined according one of the *fault models* $F1$ and $F2$ (presented below).

Counterfactual Sets of Incorrect Behaviors: In this paper we consider several possible scenarios regarding the counterfactual behaviors of the incorrectly behaving components, described in the following list. It is important to note that these are just a few representative scenarios among all that can be captured within our framework. For all components C_i ,

- F1. the only incorrect* local traces for component C_i that may be included in counterfactual sets are $\mathbf{tr}|_{\Sigma_i}$ and its prefixes. Essentially, this fault model assumes that if the inputs to C_i change, then the faulty behavior disappears, and is replaced by correct behaviors (according to φ_i) over the new input. Thus we assume that the faulty behavior of C_i was *only* for the particular input in $\mathbf{tr}|_{\Sigma_i}$.
- F2. the only incorrect* local traces for component C_i that may be included in counterfactual sets are ones that agree with both: (i) w_{mcp} , where w_{mcp} is the maximal correct prefix (with respect to φ_i) of $\mathbf{tr}|_{\Sigma_i}$; and (ii) $\mathbf{tr}|_{\text{out}(\Sigma_i)}$. Thus, any counterfactual trace must be as the original one of C_i till the first error there, and after that it must follow the same sequence of $\text{out}(\Sigma_i)$ output symbols as $\mathbf{tr}|_{\text{out}(\Sigma_i)}$. This model implies that after the first error in C_i , if the input were to change, C_i would either (a) behave correctly on the new input, or (b) ignore the new input altogether and output the same sequence of output symbols as in the original trace $\mathbf{tr}|_{\text{out}(\Sigma_i)}$.

While each fault model is imperfect, there is not much else that can be done given that the input for the analysis consists only of the properties $\varphi_1, \dots, \varphi_n$ and a single execution trace \mathbf{tr} . Thus, there is no mechanism to predict what the output of a component will be when its input changes, without paying the cost of running additional simulations. If such additional data is available (as in [23]), it can be easily incorporated in our models. In our work, we focus on the fault models $F1$ and $F2$. \square

Now we present several ways of constructing counterfactual tracesets, which use the notion of the *maximal correct prefix* of the observed trace. The *maximal correct prefix* of the trace \mathbf{tr} , denoted $\text{maxcp}(\mathbf{tr})$ is defined to be the maximal prefix $\mathbf{tr}_{\text{mcp}} \leq \mathbf{tr}$ that satisfies all local specifications, i.e. (a) \mathbf{tr}_{mcp} projected onto Σ_i is a subset of φ_i for all i ; and (b) for every prefix \mathbf{tr}_p of \mathbf{tr} such that \mathbf{tr}_{mcp} is a strict prefix of \mathbf{tr}_p , there is a j such that $\mathbf{tr}_p|_{\Sigma_j} \notin \varphi_j$.

Local Counterfactual Tracesets. We define the following counterfactual tracesets for component C_i .

★ **Repair_{tr}(C_i)** defined as: $\{w \in \varphi_i \mid \text{maxcp}(\mathbf{tr})|_{\Sigma_i} \leq w\}$.

That is, we keep the prefix $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$ for component C_i , and then take all possible correct C_i behavior extensions following $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$; i.e. we *repair* the errors following $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$, as well as the effects in C_i of errors in other components after $\text{maxcp}(\mathbf{tr})$. Observe that $\text{Repair}_{\mathbf{tr}}(C_i)$ is a subset of φ_i . Intuitively, this set captures the set of possible outcomes of C_i after $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$, if no error had occurred in *any* component. We illustrate this traceset in Figure 2.

The first trace in the Figure is the local trace for component C_i , obtained as the projection of the global trace on Σ_i . The point v_i denotes the place of the first violation of the local property φ_i by C_i . The point v denotes the place of the first violation of *some* local property φ_k by C_k where k may be different from i . Thus, the portion of $\mathbf{tr}|_{\Sigma_i}$ until v is equal to $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$. The set $\text{Repair}_{\mathbf{tr}}(C_i)$ is obtained by taking the cone of all correct executions of C_i from the prefix $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$.

Observe, as depicted in the Figure, that there might be a strict prefix $\mathbf{tr}_p < \mathbf{tr}$ such that $\text{maxcp}(\mathbf{tr})|_{\Sigma_i} < \mathbf{tr}_p|_{\Sigma_i}$, and $\mathbf{tr}_p|_{\Sigma_i} \in \varphi_i$, i.e., component C_i might continue to behave correctly in $\mathbf{tr}|_{\Sigma_i}$ after $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$; however the behavior after $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$ is considered to be *tainted*. This is because after $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$ there is some component which behaves incorrectly, and that incorrect behavior might affect other components. Thus, we consider the cone of all possible behaviors after $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$. Before $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$, no component is in error, and all are behaving according to their specifications; thus, we need not con-




Figure 2: Traceset $\text{Repair}_{\mathbf{tr}}(C_i)$.

sider alternate traces before $\text{maxcp}(\mathbf{tr})|_{\Sigma_i}$.

★ **Feasible $_{\mathbf{tr}}^{F1}(C_i)$** is defined as: $\text{Prefs}(\mathbf{tr}|_{\Sigma_i}) \cup \text{Repair}_{\mathbf{tr}}(C_i)$.

This traceset is obtained by adding all the prefixes of the observed (possibly incorrect) trace $\mathbf{tr}|_{\Sigma_i}$ to the set $\text{Repair}_{\mathbf{tr}}(C_i)$. Thus, this traceset consists of all local traces for C_i , that are considered feasible according to either the observed trace; or to the promised behavior of component C_i after the prefix $\text{maxcp}(\mathbf{tr})$. This set models the faulty behavior of \mathcal{C} under fault model $F1$.

In $\text{Feasible}(C_i)$, we take the prefix set of the incorrect behavior, instead of only the whole incorrect trace, because we want the causality analysis to be robust: the analysis should consider every intermediate trace prefix which is in error. We also include correct behaviors in $\text{Feasible}(C_i)$, because (i) although we want $\text{Feasible}(C_i)$ to model incorrect behaviors, we do not want other components to *count on* C_i behaving incorrectly; and (ii) correcting the behavior of some components might lead to inconsistencies with the original local incorrect traces.

★ **Feasible $_{\mathbf{tr}}^{F2}(C_i)$** .

This set is used to model the faulty behavior of C_i under fault model $F2$. We first obtain the incorrect traces for C_i under $F2$. Let w_{mcp} be the maximal correct prefix (with respect to φ_i) of $\mathbf{tr}|_{\Sigma_i}$. Let $L_{w_{\text{mcp}}}^{F2}(C_i) \subseteq \Sigma_i^*$ be the language such that $u \in L_{w_{\text{mcp}}}^{F2}(C_i)$ iff $u = w_{\text{mcp}} \cdot v$ for some $v \in \Sigma_i^*$ such that $\text{len}(u) = \text{len}(\mathbf{tr}|_{\Sigma_i})$ and $u|_{\text{out}(\Sigma_i)} = \mathbf{tr}|_{\text{out}(\Sigma_i)}$. Thus, to obtain $L_{w_{\text{mcp}}}^{F2}(C_i)$, we cement the maximal correct local prefix w_{mcp} , and for the positions after that we keep the *same output* as in $\mathbf{tr}|_{\Sigma_i}$ and we allow for all possible inputs. The set $\text{Feasible}_{\mathbf{tr}}^{F2}(C_i)$ is defined to be:

$$\text{Feasible}_{\mathbf{tr}}^{F2}(C_i) = \text{Prefs}(L_{w_{\text{mcp}}}^{F2}(C_i)) \cup \text{Repair}_{\mathbf{tr}}(C_i).$$

Observe that since $L_{w_{\text{mcp}}}^{F2}(C_i)$ contains $\mathbf{tr}|_{\Sigma_i}$, we have $\text{Feasible}_{\mathbf{tr}}^{F1}(C_i)$ to be a subset of $\text{Feasible}_{\mathbf{tr}}^{F2}(C_i)$.⁵

Example 5 (Counterfactual Sets). Consider the error trace \mathbf{tr} from Example 4 and component C_3 .

- $\text{Repair}_{\mathbf{tr}}(C_3)$ consists of all traces in φ_3 that agree with $\mathbf{tr}|_{\Sigma_3}$ up to and including position 2 (recall that 3 was the first position at which φ_3 was violated).
- $\text{Feasible}_{\mathbf{tr}}^{F1}(C_3)$ extends $\text{Repair}_{\mathbf{tr}}(C_3)$ with all prefixes of $\mathbf{tr}|_{\Sigma_3}$.
- $\text{Feasible}_{\mathbf{tr}}^{F2}(C_3)$ extends $\text{Feasible}_{\mathbf{tr}}^{F1}(C_3)$ by including all traces in $w \in \Sigma_3^*$ such that: (i) $\text{len}(w) = 4$; and (ii) $\mathbf{tr}|_{\Sigma_3}[1..2]$ is a substring of w ; and (iii) w agrees with $\mathbf{tr}|_{\Sigma_3}$ on the variables $\text{out}(X_3)$. □

Many of the traces in these local tracesets are infeasible due to interaction with other components. These infeasibilities will be taken care of in the construction of global counterfactual tracesets explained later in this subsection.

In addition to the two counterfactual tracesets above, we have the most expansive tracesets:

- (a) φ_i , which is a superset of $\text{Repair}_{\mathbf{tr}}(C_i)$,
- (b) $\text{Prefs}(\mathbf{tr}|_{\Sigma_i}) \cup \varphi_i$, a superset of $\text{Feasible}_{\mathbf{tr}}^{F1}(C_i)$,
- (c) $\text{Prefs}(L_{w_{\text{mcp}}}^{F2}(C_i)) \cup \varphi_i$, a superset of $\text{Feasible}_{\mathbf{tr}}^{F2}(C_i)$.

Notation. For a set $\mathcal{D} = \{D_1, \dots, D_m\}$ of components we denote $\text{Repair}_{\mathbf{tr}}(\mathcal{D}) = \text{Repair}_{\mathbf{tr}}(D_1) \parallel \dots \parallel \text{Repair}_{\mathbf{tr}}(D_m)$, and similarly for the functions $\text{Feasible}_{\mathbf{tr}}^{F1}$ and $\text{Feasible}_{\mathbf{tr}}^{F2}$.

Global Counterfactual Tracesets. The global counterfactual tracesets of a system with respect to the component collection \mathcal{C} are obtained by composing appropriately chosen local counterfactual tracesets, for components both in \mathcal{C} and in $\bar{\mathcal{C}}$. That is, for each component C_i , we pick a counterfactual traceset T_i , e.g., $T_i = \text{Repair}_{\mathbf{tr}}(C_i)$, or $T_i = \text{Feasible}_{\mathbf{tr}}^{F1}(C_i)$, or $T_i = \text{Feasible}_{\mathbf{tr}}^{F2}(C_i)$. The global counterfactual traceset is then $T_1 \parallel \dots \parallel T_n$. Local traces from T_i which become infeasible due to component interactions get automatically eliminated by the language composition definition. In the next section we show what are the appropriate local counterfactual tracesets that need to be chosen; and how global counterfactual sets can be used for various kinds of causality inference.

⁵Note that in case the observed behavior $\mathbf{tr}|_{\Sigma_i}$ satisfies the local specification φ_i , we have $\text{Feasible}_{\mathbf{tr}}^{F2}(C_i)$ and $\text{Feasible}_{\mathbf{tr}}^{F1}(C_i)$ both to be equal to $\text{Repair}_{\mathbf{tr}}(C_i)$.

3.2 Causality Analysis with Counterfactuals

Causality analysis uses counterfactual sets for reasoning about the following two scenarios:

1. **Fault Mitigation Capability:** Would the *correct behavior* of the component set \mathcal{C} be enough to mitigate the faults of all components (*including those of components that are not in \mathcal{C}*), by ensuring that the global property φ holds?
2. **Fault Manifestation:** Is the observed *faulty behavior* of the component set \mathcal{C} enough to manifest a global fault (*i.e.*, does it lead to global behaviors that violate φ), even if the components in $\overline{\mathcal{C}}$ were to behave correctly?

If the answer to the first question above is affirmative, we classify the component set \mathcal{C} as *fault mitigation-capable*. If the answer to the second question is affirmative, we classify \mathcal{C} as *fault manifesting*⁶. (In [8, 7], a fault mitigation-capable set is known as a *necessary cause*; and a set which manifests faults is known as a *sufficient cause*.) Here, we use the more reasoning-mechanism explicit names, and try avoid referring to these sets as causes, to keep the trace analysis separate from the philosophical aspects of causality.⁷ In Subsection 3.3 we analyze fault mitigation-capable component sets. Fault manifestation analysis is presented in Subsection 3.4.

Remark. Before we formally define the sets, note that correcting an individual component does not always make things better with respect to the global requirement φ , i.e., two wrongs can make a right.

3.3 Causality Analysis: Fault Mitigation

Fault Mitigation-Capable Sets. Intuitively, a component set \mathcal{C} is fault mitigation-capable if it can, were it to behave correctly, mask the faults of $\overline{\mathcal{C}}$ in the observed trace \mathbf{tr} with respect to φ by ensuring that *every* trace in the counterfactual traceset belongs to φ . Here we present the definitions of two possible such sets that arise from two natural choices of counterfactual tracesets.

★ **MitigCbl-1.** Component set \mathcal{C} is fault mitigation-capable if

$$\text{Repair}_{\mathbf{tr}}(\mathcal{C}) \parallel \text{Feasible}_{\mathbf{tr}}^{F1}(\overline{\mathcal{C}}) \subseteq \varphi \quad (1)$$

Thus, we correct the behavior of the components in \mathcal{C} ; and take the incorrect *together* with the correct behaviors of components in $\overline{\mathcal{C}}$, and ask if all the resultant traces are in φ . An obvious question is why the correct behaviors of $\overline{\mathcal{C}}$ need to be taken – since $\text{Repair}_{\mathbf{tr}}(\mathcal{C})$ contains only correct behaviors of \mathcal{C} , composing these correct behaviors with correct behaviors from $\overline{\mathcal{C}}$ would automatically result in the satisfaction of φ . The reason for this is the following subtlety. Let $\mathcal{C} = \{C_1\}$ and $\overline{\mathcal{C}} = \{C_2, C_3\}$. Suppose all components are faulty in the observed trace. If we correct C_1 , then the situation can arise where $\text{Repair}_{\mathbf{tr}}(C_1) \parallel \text{Prefs}(\mathbf{tr}|_{\Sigma_2}) \parallel \text{Prefs}(\mathbf{tr}|_{\Sigma_3})$ is the empty set due to inconsistencies between $\text{Repair}_{\mathbf{tr}}(C_1)$ and $\text{Prefs}(\mathbf{tr}|_{\Sigma_3})$ (and thus $\text{Repair}_{\mathbf{tr}}(C_1) \parallel \text{Prefs}(\mathbf{tr}|_{\Sigma_2}) \parallel \text{Prefs}(\mathbf{tr}|_{\Sigma_3}) \subseteq \varphi$ vacuously), but $\text{Repair}_{\mathbf{tr}}(C_1) \parallel \text{Prefs}(\mathbf{tr}|_{\Sigma_2}) \parallel \text{Feasible}_{\mathbf{tr}}^{F1}(C_3)$ is not empty; and moreover is not a subset of φ . That is, including correct behaviors of some components in $\overline{\mathcal{C}}$ can help us in finding out that correcting the behaviors of the components in \mathcal{C} does not suffice to ensure satisfaction of the global property.

Approximation introduced by the analysis. As mentioned previously, the counterfactual analysis procedure can only rely on a single observed trace \mathbf{tr} and the expected behavior

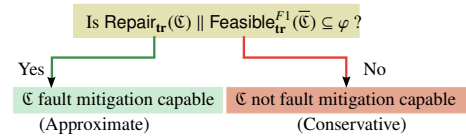


Figure 3: Fault mitigation capability analysis under F1.

⁶ \mathcal{C} can contain both faulty, and non-faulty components. It follows from our definitions in the following sections that if \mathcal{C} does not contain *any* faulty components, then \mathcal{C} is neither fault mitigation-capable, nor fault manifesting.

⁷Readers who are more comfortable with causality terminology can regard “fault mitigation-capable set” as an alias for “necessary cause”; and “fault manifesting set” as an alias for “sufficient cause”.

component specifications. It assumes that all global traces resulting from composing local projections of \mathbf{tr} and correct local traces are *possible executions* of the system. Under this assumption, a *conservative* result is one that relies on the existence of an execution (with certain properties) *in this set* (the set being $\text{Feasible}_{\mathbf{tr}}^{F1}()$).

In particular, the answer “No” in Figure 3, that is, when Equation 1 is *not* satisfied is conservative. A negative answer is given when one of the following two cases arises:

- After correcting the components in \mathbb{C} the violation of φ will remain under the original *observed* faulty behaviors of the components in $\bar{\mathbb{C}}$. That is, we have $\text{Repair}_{\mathbf{tr}}(\mathbb{C}) \parallel \text{Prefs}(\mathbf{tr}|_{\Sigma_{\bar{\mathbb{C}}}}) \not\subseteq \varphi$.
- After correcting the components in \mathbb{C} the violation of φ will remain in the case when *some components in $\bar{\mathbb{C}}$ are corrected*. That is, when we have that $\text{Repair}_{\mathbf{tr}}(\mathbb{C}) \parallel \text{Prefs}(\mathbf{tr}|_{\Sigma_{\bar{\mathbb{C}}}}) \subseteq \varphi$ but it holds that $\text{Repair}_{\mathbf{tr}}(\mathbb{C}) \parallel \text{Feasible}_{\mathbf{tr}}^{F1}(\bar{\mathbb{C}}) \not\subseteq \varphi$.

In both cases there exists a global counterfactual trace in the set of possible executions discussed above that violates φ , and thus Equation 1 is conservative when it gives a “No” answer.

Since Equation 1 is based on the fault model $F1$ (see Subsection 3.1), it is based on the assumption that in case the repair of \mathbb{C} changes the input to the faulty components in $\bar{\mathbb{C}}$, these components will react correctly (that is, satisfying their local specifications) to the new input. As this assumption is not always guaranteed (as mentioned before, there is no mechanism to predict what happens when we change inputs to faulty components), the “Yes” answer in Figure 3 is *approximate*. That is, in the case when the actual components do not satisfy the fault model $F1$ a positive answer need not imply that correcting the components in \mathbb{C} will result only in executions that satisfy φ (this may or may not be the case, since changing the input to $\bar{\mathbb{C}}$ may lead to new faulty behaviors).

Our next definition of fault mitigation-capable sets is based on the fault model $F2$ from Subsection 3.1 and includes into consideration additional counterfactual behaviors thus allowing for a finer analysis in the cases when Equation 1 is satisfied.

★ **MitigCbl-2.** Component set \mathbb{C} is fault mitigation-capable if

$$\text{Repair}_{\mathbf{tr}}(\mathbb{C}) \parallel \text{Feasible}_{\mathbf{tr}}^{F2}(\bar{\mathbb{C}}) \subseteq \varphi \quad (2)$$

This corresponds to the fault model $F2$ from Subsection 3.1.

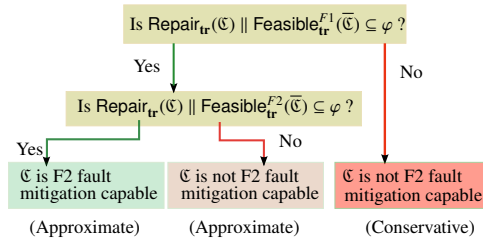


Figure 4: Fault mitigation capability analysis under $F2$.

components that behave correctly in the observed trace \mathbf{tr} we consider only sets of correct behaviors. A discussion of the approximations mentioned in the figure can be found in the appendix.

Example 6 (Fault Mitigation Capability). Let us consider again the system and the error trace from Example 4. Using the analysis above, we conclude that under each of the fault models $F1$ and $F2$:

- $\{C_1\}$ is not fault mitigation capable. This is obvious, since the component C_1 behaves correctly in the observed trace \mathbf{tr} .
- $\{C_2\}$ is fault mitigation capable. In the observed trace \mathbf{tr} , C_2 violates its safety requirement at positions greater or equal to 2, and the violation at position 4 results in a violation of the

global specification φ . Composing the set $\text{Repair}_{\text{tr}}(C_2)$ with $\text{Feasible}_{\text{tr}}^{F1}(C_1) = \text{Repair}_{\text{tr}}(C_1)$ and $\text{Feasible}_{\text{tr}}^{F1}(C_3)$ yields traces in φ , since by correcting C_2 we also eliminate the observed violation of φ_3 by C_3 (recall that in tr the first violation of φ_2 occurs at position 2 and the first violation of φ_3 at 3).

In this example, even under the fault model $F2$, where $\text{Feasible}_{\text{tr}}^{F2}(C_3)$ also includes local traces where the output of C_3 remains as in tr , the violation of φ_3 and φ gets ruled out.

- $\{C_3\}$ is fault mitigation capable. In the observed error trace tr C_2 's consumption at step 2 exceeds the allocated amount in a way that can be tolerated by C_3 in step 3, where the violation of φ_3 occurs. Thus composing any trace in $\text{Repair}_{\text{tr}}(C_3)$ with traces from $\text{Feasible}_{\text{tr}}^{F1}(C_1) = \text{Repair}_{\text{tr}}(C_1)$ and $\text{Feasible}_{\text{tr}}^{F1}(C_2)$ results in a trace on which φ is satisfied (even if φ_2 might still be violated). As the elimination of the global violation does not depend on C_2 reacting to the changed input it receives from C_3 , this holds also under the fault model $F2$. \square

In Example 6 we saw that the set $\{C_2\}$ is fault mitigation capable under both fault models $F1$ and $F2$ because the elimination of the global violation did not depend on C_3 reacting to the changed input from C_2 . In the Appendix, we give two additional examples which present situations in which this is not the case, and demonstrate scenarios in which fault mitigation capability classification depends on the fault mode used.

Notes. (A.) For \mathcal{C} to be fault mitigation-capable, we require that *all* traces in the counterfactual sets of Equations 1 and 2 be in φ . Thus, we have a universal quantifier over the traces in the counterfactual set. A bigger counterfactual set makes it harder to classify \mathcal{C} as fault mitigation-capable. As a corollary, if \mathcal{C} is fault mitigation-capable under Equation 2, it is also so under Equation 1 (recall that $\text{Feasible}_{\text{tr}}^{F1}(C_i)$ is a subset of $\text{Feasible}_{\text{tr}}^{F2}(C_i)$). If \mathcal{C} is *not* fault mitigation-capable under Eq. 1, it is also not so under Equation 2.

(B.) The conditions in Equations 1 and 2 allow for the possibility that some components in $\bar{\mathcal{C}}$ might behave correctly, even if their observed behavior in tr was incorrect. This makes it harder to classify \mathcal{C} as being fault mitigation-capable (we recall that since two wrongs can make a right, correcting a component does not always help). \square

Application. In *fortification*, we want to know if fixing some of the components under our control would suffice to “absorb” the observed errors of the other components so that the global requirement is satisfied; and this notion is what fault mitigation-sets capture.

3.4 Causality Analysis: Fault Manifestation

In this section we analyze fault manifesting sets; which are the component sets \mathcal{C} such that their observed faulty behavior alone is sufficient to manifest a violation of the global specification φ . An immediate question for fault manifestation is whether the faulty behavior of \mathcal{C} is such that *some* resultant behavior is faulty w.r.t. φ . The answer to this question is useful in debugging contexts.

Fault Manifesting Sets. Formally, a component set \mathcal{C} is fault manifesting if its observed faulty behavior alone is enough to manifest in a global error (with respect to φ) in *some* resultant trace, even if the components in $\bar{\mathcal{C}}$ were to behave correctly. One natural fault manifesting set is as follows.

- ★ **Manifest-1.** Component set \mathcal{C} is fault manifesting if

$$\text{Feasible}_{\text{tr}}^{F1}(\mathcal{C}) \parallel \text{Repair}_{\text{tr}}(\bar{\mathcal{C}}) \not\subseteq \varphi \quad (3)$$

The resulting classification analysis is depicted in Figure 5.

In case of an answer “Yes” in Figure 5, we have that there exists a scenario in which the observed behavior of *some* components in \mathcal{C} is sufficient to lead to a violation of the global specification φ , assuming that the remaining components in \mathcal{C} (if any) and the components in $\bar{\mathcal{C}}$ behave correctly. A discussion of the approximation can be found in the appendix.

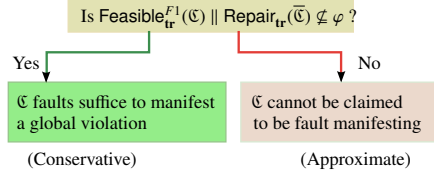


Figure 5: Fault manifestation analysis under F1 (Manifest-1).

Application. In *debugging*, we wish to find the group of components, whose erroneous behaviors may cause a violation of the global specification. Of these, groups \mathbb{C} whose erroneous behaviors are *sufficient* to manifest a global error are the most urgent ones; unless these are fixed, errors will be manifested. These sets are the Manifest-1 sets.

A Stronger Notion of Fault Manifestation. Now we present a stronger definition of fault manifesting sets, in which the possible counterfactual behaviors of components \mathbb{C} are restricted to prefixes of the observed behavior, and the requirement for the existence of traces violating φ is stronger.

Strong Fault Manifesting Sets. Intuitively, a component set \mathbb{C} is strong-fault manifesting if its observed faulty behavior alone is enough to manifest in a global error (with respect to φ) in some resultant trace, whether the components in $\overline{\mathbb{C}}$ were to behave correctly or incorrectly. Recall that $\overline{\mathbb{C}} = \{C_{1+n_{\mathbb{C}}}, \dots, C_n\}$, thus $\overline{\mathbb{C}}$ has $n_{\overline{\mathbb{C}}} = n - n_{\mathbb{C}}$ elements. For each component, consider a function $G^{C_i} : \{0, 1\} \mapsto \{\text{Prefs}(\text{tr}|_{\Sigma_{C_i}}), \text{Repair}_{tr}(C_i)\}$ defined by:

$$G^{C_i}(0) = \text{Prefs}(\text{tr}|_{\Sigma_{C_i}}); \quad G^{C_i}(1) = \text{Repair}_{tr}(C_i).$$

Now consider the natural extension to $\overline{\mathbb{C}}$, where $G^{\overline{\mathbb{C}}} : \{0, 1\}^{n_{\overline{\mathbb{C}}}} \mapsto \text{Feasible}_{tr}^{F1}(\overline{\mathbb{C}})$ defined by:

$$G^{\overline{\mathbb{C}}}(b_1, b_2, \dots, b_{n-n_{\mathbb{C}}}) = G^{C_{1+n_{\mathbb{C}}}}(b_1) \parallel G^{C_{2+n_{\mathbb{C}}}}(b_2) \parallel \dots \parallel G^{C_n}(b_{n-n_{\mathbb{C}}})$$

That is, the boolean vector $(b_1, b_2, \dots, b_{n-n_{\mathbb{C}}})$ tells us whether to choose $\text{Prefs}(\text{tr}|_{\Sigma_{C_i}})$, or $\text{Repair}_{tr}(C_i)$ for each component of $\overline{\mathbb{C}}$ in the composition.

★ **Manifest-Strong.** Set \mathbb{C} is strongly-fault manifesting if

$$\forall (b_1, \dots, b_{n-n_{\mathbb{C}}}) \in \{0, 1\}^{n-n_{\mathbb{C}}} \quad \text{Prefs}(\text{tr}|_{\Sigma_{\mathbb{C}}}) \parallel G^{\overline{\mathbb{C}}}(b_1, b_2, \dots, b_{n-n_{\mathbb{C}}}) \not\subseteq \varphi \quad (4)$$

we have

That is, for each component $C_i \in \overline{\mathbb{C}}$, no matter whether we consider only the observed behavior $\text{Prefs}(\text{tr}|_{\Sigma_{C_i}})$, or the corrected behaviors $\text{Repair}_{tr}(C_i)$, there will be some resultant trace tr' in composition with the observed behavior $\text{Prefs}(\text{tr}|_{\Sigma_{\mathbb{C}}})$ of \mathbb{C} such that this trace tr' will violate φ . This means that the observed faulty behavior of \mathbb{C} is sufficient to manifest in a global error in some trace, no matter which components of $\overline{\mathbb{C}}$ are repaired or kept as they are.

Equation 4 enables us to make a bullet-proof argument that \mathbb{C} is for sure to blame for the φ violation. The defining criterion ensures that *no matter which subset of $\overline{\mathbb{C}}$ components were to be corrected*, some resulting trace in composition with the *observed* \mathbb{C} behavior would have resulted in a φ violation.

4 Fault Models & Language-Theoretic Algorithms

4.1 Analysis of Heterogeneous Fault Models

An important distinguishing feature of our language theoretic framework for counterfactual analysis is its modularity. This modularity yields a powerful reasoning technique that cleanly generalizes existing causality notions, *e.g.*, of [8, 7] to more expressive cases. Previous work required the same fault model for all components. Our approach allows us to drop this requirement by just changing the individual

counterfactual sets: since the set of counterfactual traces is constructed locally for each component, we do not have to assume that they follow the same fault model.

Formally, the generalization is done as follows. A *fault-model profile* for a system $\mathcal{S} = \{C_1, \dots, C_n\}$ is a tuple $\hat{f} = (f_i)_{i=1}^n$ of functions where each $f_i : \Sigma^* \rightarrow 2^{\Sigma_i^*}$ maps system traces to a set of local traces for C_i such that for every trace \mathbf{tr} , we have $\mathbf{tr}|_{\Sigma_i} \in f_i(\mathbf{tr})$, and $f_i(\mathbf{tr})$ is prefix closed. Intuitively, f_i describes the fault model for component C_i , and given an observed (error) trace $\mathbf{tr} \in \Sigma^*$, the set $f_i(\mathbf{tr})$ is the set of possible local counterfactual traces for C_i (which includes $\mathbf{tr}|_{\Sigma_i}$ since it was observed). In this generalized setting, the sets $\text{Feasible}_{\mathbf{tr}}^{F1}$ and $\text{Feasible}_{\mathbf{tr}}^{F2}$ define two specific functions: $f'_i(\mathbf{tr}) = \text{Feasible}_{\mathbf{tr}}^{F1}(C_i)$ and $f''_i(\mathbf{tr}) = \text{Feasible}_{\mathbf{tr}}^{F2}(C_i)$. Let

$$\text{CFac}_{\mathbf{tr}}^{\hat{f}}(\bar{\mathcal{C}}) = f_k(\mathbf{tr}) \parallel f_{k+1}(\mathbf{tr}) \parallel \dots \parallel f_n(\mathbf{tr}),$$

where $\bar{\mathcal{C}} = \{C_k, C_{k+1}, \dots, C_n\}$ (the set $\text{CFac}_{\mathbf{tr}}^{\hat{f}}(\bar{\mathcal{C}})$ is defined similarly). Using the counterfactual set $\text{CFac}_{\mathbf{tr}}^{\hat{f}}$, we give a general definition of mitigation capability based causality, generalizing Equations 1 and 2, as follows: the component set \mathcal{C} is fault mitigation-capable under the fault-model profile \hat{f} if

$$\text{Repair}_{\mathbf{tr}}(\mathcal{C}) \parallel \text{CFac}_{\mathbf{tr}}^{\hat{f}}(\bar{\mathcal{C}}) \subseteq \varphi \quad (5)$$

Similarly, the general definition of fault manifestation based causality as follows: component set \mathcal{C} is fault manifesting under \hat{f} if

$$\text{CFac}_{\mathbf{tr}}^{\hat{f}}(\mathcal{C}) \parallel \text{Repair}_{\mathbf{tr}}(\bar{\mathcal{C}}) \not\subseteq \varphi \quad (6)$$

Employing heterogeneous fault models leads to improved precision of the causality analysis, easily incorporating designer knowledge about the behaviour of components and available simulation data.

Bridging the Gap to Structural-Model based Causality. The seminal work of Halpern and Perl [12] investigated a notion of causality in non-reactive settings which was based on *structural equations* between variables which specified which variables affect which others. The fault model profile \hat{f} , and its utilization in causality definitions 5 and 6 bridges the gap between structural-model based causality and causality notions in a reactive setting as follows. A component C_i has an associated variable dependency given by structural equations which specify which variables affect which others (possibly in the future), e.g. a change in input variable x at a time-step will lead to a change in output variable y at some point in the future; but another output variable z will remain unaffected by the change in x — this means that y depends on x , and z does not depend on x . This variable dependency can be utilized in the fault model profile f_i for C_i : the set $f_i(\mathbf{tr})$ will only contain strings which satisfy the structural variable dependencies mentioned above. Of course, a change in variable x may lead to a change in y in the future, and if y is an input to some other component C_j , this may lead to a change in some other variable u , and this change may flow back to C_i in effect changing z . Thus, we have two manners in which changes in variable values propagate: (i) locally inside a component (perhaps through states), and (ii) in an inter-component fashion in the reactive setting. A fault-model profile based on structural equations captures the first kind of variable change effects. Language composition *automatically* accounts for the second kind of variable change flow for counterfactual reasoning in a modular fashion. Thus, our causality framework using fault-model profiles lays down the theoretical foundations for connecting the work in structural-model based non-reactive causality, to causality in a reactive setting.

4.2 Algorithm Complexity using Language-Theoretic Analysis

We now analyze the time complexity of determining causality based on the language theoretic framework of Section 3. We discuss the language theoretic operations employed, and give bounds for the case

where the components are given as finite state automata. For more expressive models, the time bounds correspond to time bounds of analogous language operations.

For an alphabet Σ and a word w , let

- $\text{Cone}_\Sigma(w) = \{w\} \cdot \Sigma^*$, i.e., the word w followed by all possible strings in Σ^* ; and
- for $\Sigma(X) = \Sigma'(X') \parallel \Xi(X'')$, for some alphabets $\Sigma'(X')$ and $\Xi(X'')$, and for w_p a prefix of w , let

$$\text{AlterRest}_\Xi(w, w_p, \Sigma) = \{w_p \cdot u \mid u \in \Sigma^{|w|-|w_p|} \text{ and } (w_p \cdot u)|_\Xi = w|_\Xi\}.$$

The set $\text{AlterRest}_\Xi(w, w_p, \Sigma)$ contains words of length $|w|$ obtained from w by keeping the first $|w_p|$ letters unchanged, and then changing all letters *not* in Ξ to all possible values (this corresponds to changing valuations of variables in $X' \setminus X''$ after w_p).

The counterfactual sets from Section 3.1 can be defined using these languages and basic operations on languages as given below. Here, w_{mxp} is the maximal correct prefix (with respect to φ_i) of $\text{tr}|_{\Sigma_i}$.

$$\begin{aligned} \text{Repair}_{\text{tr}}(C_i) &= \varphi_i \cap \text{Cone}_{\Sigma_i}(\text{maxcp}(\text{tr})|_{\Sigma_i}), \\ \text{Feasible}_{\text{tr}}^{F1}(C_i) &= \text{Prefs}(\text{tr}|_{\Sigma_i}) \cup \text{Repair}_{\text{tr}}(C_i), \\ \text{Feasible}_{\text{tr}}^{F2}(C_i) &= \text{Prefs}(\text{AlterRest}_{\text{out}(\Sigma_i)}(\text{tr}|_{\Sigma_i}, w_{\text{mxp}}, \Sigma_i)) \\ &\quad \cup \text{Repair}_{\text{tr}}(C_i), \end{aligned}$$

Specific algorithms for the case of finite automata to obtain the basic sets on the right are as follows.

Recall that a non-deterministic finite automaton (NFA) over an alphabet Σ is a tuple $\mathcal{A} = (Q, q_0, \Sigma, \rho, Q_f)$, where Q is a finite set of states, $q_0 \in Q$ is an initial state, Σ is the input alphabet, $\rho \subseteq Q \times \Sigma \times Q$ is a transition relation, and $Q_f \subseteq Q$ is a set of accepting states. A *deterministic* automaton (DFA) is one where for any $q \in Q$ and $\sigma \in \Sigma$, there is at most one q' such that $(q, \sigma, q') \in \rho$. We denote $\mathcal{L}(\mathcal{A})$ as the language of words in Σ^* accepted by \mathcal{A} . Define $|\mathcal{A}| = |Q| + |\rho|$ (thus $|\mathcal{A}| \leq |Q| \times |\Sigma|$). Let the local and global specifications $\varphi_1, \dots, \varphi_n$ and φ be given as DFAs or NFAs $\mathcal{A}_1, \dots, \mathcal{A}_n$ and \mathcal{A} respectively. Note that since the specifications are prefix closed, we can assume that all reachable states are final states [14].

The various entities in the previous equations are obtained as follows.

- The string w_{mxp} can be obtained from tr and φ_i in time $O(|\text{tr}| \cdot |Q_i|^2)$ by running the automaton \mathcal{A}_i on $\text{tr}|_{\Sigma_i}$ (if \mathcal{A}_i is a DFA, this can be done in $O(|\text{tr}|)$ time). Similarly, string $\text{maxcp}(\text{tr})$ can be obtained in $O(|\text{tr}| \cdot \sum_{i=1}^n |Q_i|^2)$ time ($O(n \cdot |\text{tr}|)$ in the DFA case).
- A DFA \mathcal{D}_i with $|\text{tr}|_{\Sigma_i}|$ states (and size $|\text{tr}|_{\Sigma_i}| + |\Sigma_i|$) can be constructed such that $\mathcal{L}(\mathcal{D}_i) = \text{Cone}_{\Sigma_i}(\text{tr}|_{\Sigma_i})$.
- We can construct a DFA \mathcal{D}'_i with $|\text{tr}|_{\Sigma_i}|$ states and size such that $\mathcal{L}(\mathcal{D}'_i) = \text{Prefs}(\text{tr}|_{\Sigma_i})$ using the standard prefix construction.
- We can construct a DFA for $\text{AlterRest}_{\text{out}(\Sigma_i)}(\text{tr}|_{\Sigma_i}, w_{\text{mxp}}, \Sigma_i)$ with $|\text{tr}|_{\Sigma_i}|$ states (and size $|\text{tr}|_{\Sigma_i}| \cdot |\Sigma_i|$).

It can be modified to accept $\text{Prefs}(\text{AlterRest}_{\text{out}(\Sigma_i)}(\text{tr}|_{\Sigma_i}, w_{\text{mxp}}, \Sigma_i))$ by making all states final.

Union and intersection are standard operations on NFAs/DFAs. Thus, the sets $\text{Repair}_{\text{tr}}(C_i)$, and $\text{Feasible}_{\text{tr}}^{F1}(C_i)$ and $\text{Feasible}_{\text{tr}}^{F2}(C_i)$ can all be obtained in polynomial time and represented as NFAs of size polynomial in the sizes of $\mathcal{A}_1, \dots, \mathcal{A}_n$.

Consider a fault-model profile $\hat{f} = (f_i)_{i=1}^n$ such that $f_i(\text{tr}) = \text{Feasible}_{\text{tr}}^{F1}(C_i)$ or $f_i(\text{tr}) = \text{Feasible}_{\text{tr}}^{F2}(C_i)$. Recall Equations 5 and 6. The equations involve taking the parallel composition of the $f_i(\text{tr})$ sets. As we just showed, each $f_i(\text{tr})$ set can be represented as the language of a NFA (or DFA) of polynomial size. The parallel composition of the $f_i(\text{tr})$ sets can be obtained by taking the parallel composition of the corresponding automata using the product construction (in polynomial time). Finally, the equations involve making language inclusion checks, which involve checking $\mathcal{L}(\mathcal{B}_1) \parallel \dots \parallel \mathcal{L}(\mathcal{B}_n) \subseteq \mathcal{L}(\mathcal{A})$, where \mathcal{B}_i are (polynomial sized) automata derived as above for either the $\text{Repair}_{\text{tr}}$ or $\text{Feasible}_{\text{tr}}^{F1}$ or $\text{Feasible}_{\text{tr}}^{F2}$ sets. This check can be performed by checking $\mathcal{L}(\mathcal{B}_1) \parallel \dots \parallel \mathcal{L}(\mathcal{B}_n) \subseteq \mathcal{L}(\bar{\mathcal{A}})$ where $\bar{\mathcal{A}}$ is the deterministic automaton for \mathcal{A} . Putting everything together, we get the following.

Theorem 1. *Let $\hat{f} = (f_i)_{i=1}^n$ be a fault-model profile such that $f_i(\mathbf{tr}) = \text{Feasible}_{\mathbf{tr}}^{F1}(C_i)$ or $f_i(\mathbf{tr}) = \text{Feasible}_{\mathbf{tr}}^{F2}(C_i)$. Let the local and global specifications $\varphi_1, \dots, \varphi_n$ and φ (such that $\varphi_1 \parallel \dots \parallel \varphi_n \subseteq \varphi$) be given as NFAs (or DFAs) $\mathcal{A}_1, \dots, \mathcal{A}_n$ and \mathcal{A} respectively. Given an observed trace $\mathbf{tr} \notin \varphi$, for the fault profile \hat{f} , a component set \mathcal{C} can be determined to be: fault mitigation capable (Equation 5), or fault manifesting (Equation 6) in time (i) polynomial in the sizes of $\mathcal{A}_1, \dots, \mathcal{A}_n$, and \mathbf{tr} ; and (ii) exponential in $|\mathcal{A}|$ in case \mathcal{A} is an NFA, or polynomial in $|\mathcal{A}|$ in case \mathcal{A} is a DFA. \square*

Determining whether \mathcal{C} is strongly-fault manifesting according to Equation (4) is harder as there is an additional for all quantifier, and thus is exponential time even in the case \mathcal{A} is a DFA.

A careful analysis shows that for fault-model profiles with $f_i(\mathbf{tr}) = \text{Feasible}_{\mathbf{tr}}^{F1}(C_i)$ or $f_i(\mathbf{tr}) = \text{Feasible}_{\mathbf{tr}}^{F2}(C_i)$, the fault mitigation capability (Equation 5) check can be performed only with word-sets that are of length at most $|\mathbf{tr}|$. As a result, the fault mitigation capability problem is in co-NP. A similar argument shows fault manifestation determination to be in NP. This also allows us to check language inclusion, without determinizing \mathcal{A} , in time polynomial in $|\mathcal{A}|$ and exponential in $|\mathbf{tr}|$.

For general regular language fault-model profiles when $\varphi_1, \dots, \varphi_n, \varphi$ are given as NFAs language inclusion for two nondeterministic automata can be encoded in each of the causality analysis questions. Thus, the causality analysis problem is PSPACE-complete (membership in PSPACE follows from PSPACE membership of NFA language inclusion).

5 Discussion

We discuss some of the related work in the Appendix. Our work overcomes the shortcomings of existing work in the reactive setting as follows. As evidenced by works [8, 7, 22, 6, 9], the main challenge in causality analysis for concurrent systems is in the construction of counterfactual sets. Definitions that do not account for the effect of repairing some components on the behavior of others [8] result in vacuous causes. This implies the need for definitions that take component interactions into account. However, the ones existing in the literature are overly complicated, which hinders understanding and ensuring their correctness. For example, in [7] the cone of influence of a set of components is defined by means of the fixpoint of a function g , where g itself is defined by a formula spanning several lines containing six quantified variables, with five of them being trace temporal-position variables (cf. Definition 4 in [7]). Similarly, in [6, 9] the definition of unaffected prefixes is described as a fixpoint computation involving a sequence of four connected definitions. In addition, the minimal unaffected prefixes in [6] are not always composable. More precisely, the expression for tr_i^* in Definition 9 of [6] allows for having different extensions w for different components j , thus resulting in local unaffected prefixes that are *not* consistently extendable taken together, resulting in an empty set of global counterfactual traces. This problem stems from the mixed-up treatment of local and global traces.

Our work demonstrates that we can leverage the theory of language composition to obtain modular and transparent definitions for counterfactual sets. This decomposition allows us to focus on local alternative scenarios when comparing different causality notions; our framework based on language composition takes care of global-level reactive reasoning in a uniform way across the different causal sets. This machinery allows us to (1) define and easily reason about *new* causal sets (eg strongly-fault manifesting sets in Equation (4) that had not been considered before) based on the application need, (2) seamlessly incorporate *heterogeneous fault models* in causal sets (existing work assumes a common fault model across all components), (3) compare different causality notions; and (4) automatically obtain *algorithms* for computing different causal sets based on standard language theoretic operations.

References

- [1] Adrian Beer, Stephan Heindinger, Uwe Kühne, Florian Leitner-Fischer, and Stefan Leue. Symbolic causality checking using bounded model checking. In *SPIN*, volume 9232 of *Lecture Notes in Computer Science*, pages 203–221. Springer, 2015.
- [2] I. Beer, S. Ben-David, H. Chockler, A. Orni, and R.J. Treffler. Explaining counterexamples using causality. In *CAV 09*, LNCS 5643, pages 94–108. Springer, 2009.
- [3] F.D. Busnelli. Causation. In *Principles of European Tort Law*, pages 43–63. Springer, 2005.
- [4] H. Chockler, O. Grumberg, and A. Yadgar. Efficient automatic STE refinement using responsibility. In *TACAS 08*, LNCS 4963, pages 233–248. Springer, 2008.
- [5] H. Chockler, J.Y. Halpern, and O. Kupferman. What causes a system to satisfy a specification? *ACM Trans. Comput. Logic*, 9(3):20:1–20:26, 2008.
- [6] G. Gössler and L. Aştefănoaei. Blaming in component-based real-time systems. In *EMSOFT 14*, pages 7:1–7:10. ACM, 2014.
- [7] G. Gößler and D.L. Métayer. A general trace-based framework of logical causality. In *FACS 13*, LNCS 8348, pages 157–173. Springer, 2013.
- [8] G. Gößler, D.L. Métayer, and J.-B. Raclet. Causality analysis in contract violation. In *RV 10*, LNCS 6418, pages 270–284. Springer, 2010.
- [9] Gregor Gößler and Daniel Le Métayer. A general framework for blaming in component-based systems. *Sci. Comput. Program.*, 113:223–235, 2015.
- [10] Gregor Gößler and Jean-Bernard Stefani. Fault ascription in concurrent systems. In Pierre Ganty and Michele Loreti, editors, *Trustworthy Global Computing - 10th International Symposium, TGC 2015, Madrid, Spain, August 31 - September 1, 2015 Revised Selected Papers*, volume 9533 of *Lecture Notes in Computer Science*, pages 79–94. Springer, 2015.
- [11] S. Halle. Causality in message-based contract violations: A temporal logic “whodunit”. pages 171–180, 2011.
- [12] J.Y. Halpern and J. Pearl. Causes and explanations: A structural-model approach. part I: Causes. *The British journal for the philosophy of science*, 56(4):843–887, 2005.
- [13] D. Hume. *An Enquiry concerning Human Understanding*. 1748.
- [14] J.-Y. Kao, N. Rampersad, and J. Shallit. On NFAs where all states are final, initial, or both. *Theor. Comput. Sci.*, 410(47-49):5010–5021, 2009.
- [15] Matthias Kuntz, Florian Leitner-Fischer, and Stefan Leue. From probabilistic counterexamples via causality to fault trees. In *SAFECOMP*, volume 6894 of *LNCS*, 2011.
- [16] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978.
- [17] Florian Leitner-Fischer and Stefan Leue. Causality checking for complex system models. In *VMCAI*, volume 7737 of *Lecture Notes in Computer Science*, pages 248–267. Springer, 2013.
- [18] Florian Leitner-Fischer and Stefan Leue. Probabilistic fault tree synthesis using causality computation. *IJC-CBS*, 4(2):119–143, 2013.
- [19] D. Lewis. Void and object. In *Causation and Counterfactuals*, pages 277–290. MIT Press, 2004.
- [20] Chao Wang, Zijiang Yang, Franjo Ivancic, and Aarti Gupta. Whodunit? causal analysis for counterexamples. In *ATVA 2006*, volume 4218 of *LNCS*, pages 82–95, 2006.
- [21] S. Wang, A. Ayoub, R. Ivanov, O. Sokolsky, and I. Lee. Contract-based blame assignment by trace analysis. In *HiCoNS 13*, pages 117–126. ACM, 2013.
- [22] S. Wang, A. Ayoub, B. Kim, G. Gößler, O. Sokolsky, and I. Lee. A causality analysis framework for component-based real-time systems. In *RV 13*, LNCS 8174, pages 285–303. Springer, 2013.
- [23] Shaohui Wang, Yoann Geoffroy, Gregor Gößler, Oleg Sokolsky, and Insup Lee. A hybrid approach to causality analysis. In Ezio Bartocci and Rupak Majumdar, editors, *Runtime Verification - 6th International Conference, RV 2015 Vienna, Austria, September 22-25, 2015. Proceedings*, volume 9333 of *Lecture Notes in Computer Science*, pages 250–265. Springer, 2015.

Appendix

A Examples: Fault Mitigation Analysis

Example 7 (Fault Mitigation Capability under F2). Consider the system of Example 4. Suppose we restrict the correct behaviors of C_3 further by adding the following condition φ_3^\dagger to φ_3 : for every $1 \leq j < \text{len}(w)$, we have $w_{[j]}(x_a^{3,1}) + w_{[j]}(x_a^{3,2}) + w_{[j+1]}(x_d^3) \geq 29$. That is, component C_3 tries to optimize the resource allocation so that the combined availability (usage by C_3) is at least 29 in each step. All

j, φ	1, φ	2, φ	3, $\neg\varphi$	4, $\neg\varphi$	components are faulty in the trace on the left. In the trace, both C_1 , and C_2 exceed their allocations by 4 units each at each step (after step 3). Component C_3 is faulty from step 3 on, as it should have decreased its consumption by 5 (= 16 – 11) units from the depletion at step
local specs	$\varphi_1, \varphi_2, \varphi_3$	$\varphi_1, \neg\varphi_2, \varphi_3$	$\neg\varphi_1, \neg\varphi_2, \neg\varphi_3$	$\neg\varphi_1, \neg\varphi_2, \neg\varphi_3$	
$x_a^{3,1}$	11	12	12	12	
$x_a^{3,2}$	11	12	12	12	
$x_d^{1,3}$	0	8	16	16	
$x_d^{2,3}$	0	16	16	16	
x_d^3	0	7	5	5	

2, but instead it only decreases by 2 units (from 7 to 5), and does not decrease at all in step 4.

Analysis under fault model F1 (Equation 1) classifies $\{C_3\}$ as fault mitigation-capable, i.e. able to absorb the faults of both C_1 and C_2 . Intuitively, this seems false: if both C_1 and C_2 are consuming 16 units as observed, there is nothing C_3 can do. The reason why we get this false answer is due to the shortcoming of fault model F1. Consider any fix of C_3 . Let this fixed word be $\sigma'_1, \sigma'_2, \sigma'_3, \sigma'_4$. A fix requires that component C_3 reduce its resource depletion to 2 at step 3, as C_2 had exceeded its allocation by 5 units in the previous step (16 – 11). Because of the new optimized resource allocation requirement introduced at the beginning of the example, this reduction of $\sigma'_3(x_d^3)$ to 2 implies that $28 \geq \sigma'_2(x_a^{3,1}) + \sigma'_2(x_a^{3,2}) \geq 27$, thus, the values of at least one of $x_a^{3,1}, x_a^{3,2}$ must change from the observed 12 units in the trace at step 2 to something higher. However, F1 assumes that whenever inputs change to a faulty component, the outputs must change to correct ones, thus F1 implies that the combined consumption of C_1, C_2 will reduce to 28 or lower (from the observed 32). Thus, F1 forces us to assume that if C_3 gives a higher allocation to C_1, C_2 , it will result in a lower consumption by C_1, C_2 as their inputs have changed.

An analysis under F2 on the other hand assumes that C_1, C_2 will keep consuming 16 units from step 3 onwards, and thus will not classify $\{C_3\}$ as fault mitigation-capable. \square

Example 8 (Fault Mitigation Capability under F1). Consider the system of Example 4 (with the additional requirement φ_3^\dagger added to φ_3). In addition, let component C_2 have another output variable $x_r^{2,3}$ denoting the requested amount for the next to next time step; and let this variable be read by C_3 . Let us add to φ_2 the requirement φ_2^\dagger that the value requested $x_r^{2,3}$ at step j is at least as much as the depleted amount $x_d^{2,3}$ by C_2 in step $j + 2$. Finally, let us add another requirement φ_3^\ddagger to φ_3 saying that the value of the allocation to C_2 , i.e. $x_a^{3,2}$ is equal to its requested resource $x_r^{2,3}$ in the previous time step. Thus, C_3 tries to maintain the invariant $x_a^{3,2} = x_r^{2,3}$. Consider the trace $\sigma_1, \sigma_2, \sigma_3, \sigma_4$ on the left in which C_2, C_3 are faulty, C_1 is not. φ_3 is

j, φ	1, φ	2, φ	3, $\neg\varphi$	4, $\neg\varphi$	violated at step two because $\sigma_1(x_a^{3,1}) + \sigma_1(x_a^{3,2}) + \sigma_2(x_d^3) < 29$ violating the added optimization criteria φ_3^\dagger . This is a benign violation. The global specification stays violated at step 4 even though the combined utilization is less than 31 as the bound was violated in the previous step (we require φ to be
local specs	$\varphi_1, \varphi_2, \varphi_3$	$\varphi_1, \neg\varphi_2, \neg\varphi_3$	$\neg\varphi_1, \neg\varphi_2, \neg\varphi_3$	$\neg\varphi_1, \neg\varphi_2, \neg\varphi_3$	
$x_a^{3,1}$	13	25	7	7	
$x_a^{3,2}$	13	5	23	23	
$x_d^{1,3}$	0	10	9	7	
$x_d^{2,3}$	0	16	23	23	
$x_r^{2,3}$	5	6	8	6	
x_d^3	0	2	0	0	

prefix-closed).

Intuitively, looking at the example, the problem with C_3 is that it blindly trusted C_2 's estimates and did not increase allocation to C_2 by the end of step 2 (and concomitantly decrease allocation to C_1). While the added variable $x_r^{2,3}$ is available to C_3 , it is of no use as C_2 is giving incorrect estimates of its future resource usage. Observe that if C_1 is working perfectly (and depleting the resource way less than C_2). It appears that if C_2 had given correct values in its estimates $x_r^{2,3}$, then C_3 could have allocated correctly (telling C_1 to decrease its usage), and avoided a global violation, thus we expect $\{C_2\}$ to be fault mitigation-capable.

We claim $\{C_2\}$ is fault mitigation-capable under $F1$, but not under $F2$. The reason is that under $F2$, even if the inputs to C_3 change (in particular the estimates $x_r^{2,3}$ by component C_2), the behavior of C_3 will be assumed to be the same as observed, with the same old output values of $x_a^{3,1}, x_a^{3,2}$. However, under $F1$, with the changed inputs, the behavior of C_3 is assumed to be different, and correct; and will correctly set the changed $x_a^{3,1}, x_a^{3,2}$ values (in the process telling C_1 to reduce its usage). Thus in this example, $F1$ is the fault model which gives the intuitively correct answer to fault mitigation capability, compared to Example 7 where $F2$ gave the intuitively correct answer. This example, and Example 7 show that which fault model to choose depends on the application dynamics \square

B Example: Fault Manifestation Analysis

Example 9 (Fault Manifestation). We consider the system and the error trace from Example 4 and determine that $\{C_2, C_3\}$ is fault manifesting. The set $\text{Prefs}(\mathbf{tr}|_{\Sigma_2 \parallel \Sigma_3})$ contains the projection of \mathbf{tr} on the alphabet $\Sigma_2 \parallel \Sigma_3$. Since C_1 behaves correctly in \mathbf{tr} , the observed error trace is actually an element of the set of counterfactual traces which implies that the set $\{C_2, C_3\}$ is fault manifesting. \square

C Casual Sets: Approximations Resulting from Fault Model Assumptions

C.1 MitigCbl-1

Quantifying the confidence in the approximation. In the case of a “Yes” answer to the decision question in Figure 3, we can quantify our confidence the given answer as follows.

- If the set $\text{Repair}_{\mathbf{tr}}(\mathbb{C}) \parallel \{\mathbf{tr}|_{\Sigma_{\bar{\mathbb{C}}}}\}$ is non-empty, then it means that the behavior of \mathbb{C} can be corrected in such a way that the original faulty behavior of $\bar{\mathbb{C}}$ is compatible with the new behavior of \mathbb{C} . According to Equation 1 we have that every trace in this set satisfies the global specification φ . Thus, in this case, the answer “Yes” is actually *exact*, meaning that repairing \mathbb{C} suffices to absorb the remaining faults of $\bar{\mathbb{C}}$ that were observed.
- If, on the other hand,
 - $\text{Repair}_{\mathbf{tr}}(\mathbb{C}) \parallel \{\mathbf{tr}|_{\Sigma_{\bar{\mathbb{C}}}}\}$ is empty and
 - $\text{Repair}_{\mathbf{tr}}(\mathbb{C}) \parallel \text{Prefs}(\mathbf{tr}|_{\Sigma_{\bar{\mathbb{C}}}})$ is not empty,

then we use the maximal prefix $\mathbf{w}_{\text{mxp}} \in \text{Prefs}(\mathbf{tr}|_{\Sigma_{\bar{\mathbb{C}}}})$ such that $\text{Repair}_{\mathbf{tr}}(\mathbb{C}) \parallel \{\mathbf{w}_{\text{mxp}}\}$ is non-empty to determine our confidence in the result. The prefix \mathbf{w}_{mxp} is the longest prefix of $\mathbf{tr}|_{\Sigma_{\bar{\mathbb{C}}}}$ which is consistent with the corrected behaviors $\text{Repair}_{\mathbf{tr}}(\mathbb{C})$ of \mathbb{C} . The confidence in the “Yes” answer is bigger the longer \mathbf{w}_{mxp} is (after occurrences of errors in $\bar{\mathbb{C}}$); if e.g. $\mathbf{tr}|_{\Sigma_{\bar{\mathbb{C}}}}$ is only a one-step extension of \mathbf{w}_{mxp} , then Equation 1 tells us that correcting the components \mathbb{C} will be enough to absorb the faults of $\bar{\mathbb{C}}$ of $\mathbf{tr}|_{\Sigma_{\bar{\mathbb{C}}}}$, except possibly at the last step for which we do not know.

C.2 MitigCbl-2

Approximation introduced by the analysis. Although here we consider a larger set of counterfactual traces than in the previous definition, a positive answer to the question whether Equation 2 is satisfied is again approximate. The reason is that the components in $\bar{\mathcal{C}}$ are not guaranteed to satisfy the fault model $F2$, and, again, unconsidered behaviors (on which φ might not hold) might arise in the actual system after repairing \mathcal{C} .

Unlike in the previous definition, now an *approximation occurs even in the case of a negative answer*, as depicted in Figure 4. The reason is that we might classify \mathcal{C} as not being fault mitigation-capable under Equation 2 based on a trace which is not a possible trace of the actual system, while it can happen that in the actual system repairing \mathcal{C} leads to a completely new output behavior of $\bar{\mathcal{C}}$ that results in traces on which φ holds. Thus, the “No” answer to the check of Equation 2 is approximate.

Quantifying the confidence in the approximation. The confidence of a “Yes” answer in this case is determined in a way similar to before. The more interesting case is that of a “No” answer.

Recall that $\text{Feasible}_{\text{tr}}^{F1}(\bar{\mathcal{C}}) \subseteq \text{Feasible}_{\text{tr}}^{F2}(\bar{\mathcal{C}})$, and thus if \mathcal{C} is not fault mitigation capable under Equation 1, then it is not mitigation capable under Equation 2 either. Therefore, as shown in Figure 4, it only makes sense to consider Equation 2 when the answer to the check of Equation 1 is “Yes”. This allows us to use the confidence in the “Yes” answer for Equation 1 determined as before to determine the confidence in the “No” answer for Equation 2.

- If the “Yes” answer for Equation 1 was exact, then we should have low confidence in the “No” answer for Equation 2.
- Otherwise, the higher the confidence in the “Yes” answer for Equation 1 is, the lower our confidence in the “No” answer for Equation 2 should be.

C.3 Manifest-1

Approximation introduced by the analysis. The “Yes” decision in Figure 5 is conservative, since the witness trace that violates φ is composed of observed and correct behaviors of the components. A “No” answer, however, is bound to be approximate, in the actual system correcting the components in $\bar{\mathcal{C}}$ may result in new executions that are not considered as part of our set of counterfactuals.

Quantifying confidence in the approximation. In case of a “No” decision in Figure 5, we quantify our confidence in the “No” answer analogously to the analysis of the “Yes” decision for MitigCbl-1.

- If the set $\{\text{tr}|_{\Sigma_{\bar{\mathcal{C}}}}\} \parallel \text{Repair}_{\text{tr}}(\bar{\mathcal{C}})$ is non-empty, then it means that some correct behavior of $\bar{\mathcal{C}}$ is compatible with the observed behavior of \mathcal{C} , and all the traces in the composition satisfy φ . Thus, if $\bar{\mathcal{C}}$ were to behave correctly, the violation of φ would disappear, irrespective of the observed faulty behavior of \mathcal{C} . This indicates that the confidence in the “No” decision in this case is high.
- If, on the other hand,
 - $\{\text{tr}|_{\Sigma_{\bar{\mathcal{C}}}}\} \parallel \text{Repair}_{\text{tr}}(\bar{\mathcal{C}})$ is empty and
 - $\text{Prefs}(\text{tr}|_{\Sigma_{\bar{\mathcal{C}}}}) \parallel \text{Repair}_{\text{tr}}(\bar{\mathcal{C}})$ is not empty,
 then we consider the maximal prefix $\mathbf{w}_{\text{mxp}} \in \text{Prefs}(\text{tr}|_{\Sigma_{\bar{\mathcal{C}}}})$ such that $\{\mathbf{w}_{\text{mxp}}\} \parallel \text{Repair}_{\text{tr}}(\bar{\mathcal{C}})$ is non-empty, exactly like in the “Yes” decision for MitigCbl-1, and quantify our confidence according to the length of \mathbf{w}_{mxp} (after occurrences of errors in \mathcal{C}).

D Relationships Between Causal Sets

In this subsection we establish relations between some of the causal sets defined in Subsections 3.3 and 3.4. First, we compare the different sets with respect to the strength of the corresponding causality

notions. Recall again that in the literature on causality, fault mitigation-capable sets are called necessary causes, and that fault manifesting sets are called sufficient causes.

Proposition 1 (Fault Mitigation-Capable Sets). *If set \mathcal{C} is fault mitigation-capable under Equation (2), then it is also fault mitigation-capable under Equation (1).*

The following proposition formalizes the relationship between fault mitigation-capable sets and fault manifesting sets.

Proposition 2 (Interrelationships). *1. If \mathcal{C} is a fault mitigation-capable set according to Equation (1), then $\bar{\mathcal{C}}$ is not a fault manifesting set according to Equation (3).
2. If \mathcal{C} is a fault manifesting set according to Equation (3), then $\bar{\mathcal{C}}$ is not fault mitigation-capable under Equation (1).*

Note that the set of all components $\{C_1, \dots, C_n\}$ is trivially both a fault mitigation-capable set, and also a fault manifesting set. However, in applications we are interested in such sets that are *minimal* with respect to the subset relation, and identifying such sets is a non-trivial task. The definitions we studied here enjoy the monotonicity properties stated in the proposition below.

Proposition 3 (Monotonicity). *1. If a set \mathcal{C} satisfies Equation (1), then any superset $\mathcal{D} \supseteq \mathcal{C}$ also satisfies Equation (1).
2. If \mathcal{C} satisfies Equation (3) any superset $\mathcal{D} \supseteq \mathcal{C}$ also satisfies it.*

E Computational Complexity of Causality Problems

Let \mathcal{A}, \mathcal{B} be NFAs over an alphabet Ξ , such that $\mathcal{L}(\mathcal{A})$ and $\mathcal{L}(\mathcal{B})$ are prefix closed. Note that for NFAs where all states are final, language inclusion has the same complexity as for general NFA, i.e., it is PSPACE-complete [14]. We will now show how to reduce the question $\mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{A})$ to the fault mitigation and fault manifestation questions, thus proving their PSPACE-hardness.

We define two components C_1 and C_2 as follows. Assume w.l.o.g. that we have letters $a_1, b_1, a_2, b_2 \notin \Xi$. Let x_1 and x_2 be two variables with domains $O_1 = \{a_1, b_1\}$ and $O_2 = \{a_2, b_2\}$ respectively, and x_Ξ be a variable with domain Ξ .

We define the component $C_1 = (X_1, \text{inp}(X_1), \text{out}(X_1), \Sigma_1, \varphi_1)$, where $X_1 = \{x_\Xi, x_1, x_2\}$, $\text{inp}(X_1) = \{x_2\}$, $\text{out}(X_1) = \{x_1, x_\Xi\}$ and

$$\begin{aligned} \varphi_1 = & (\mathcal{L}(\mathcal{A}) \parallel (\{a_1\} \cdot O_1^*) \parallel (\{a_2\} \cdot O_2^*)) \cup \\ & (\mathcal{L}(\mathcal{B}) \parallel (\{a_1\} \cdot O_1^*) \parallel (\{b_2\} \cdot O_2^*)) \cup \{\epsilon\}. \end{aligned}$$

Intuitively, if the first value of x_2 is a_2 , then C_1 has to output strings from $\mathcal{L}(\mathcal{A})$, and if this value is b_2 , C_1 has to output strings from $\mathcal{L}(\mathcal{B})$.

For the other component, let $C_2 = (X_2, \text{inp}(X_2), \text{out}(X_2), \Sigma_2, \varphi_2)$, where $X_2 = \{x_2\}$, $\text{inp}(X_2) = \emptyset$, $\text{out}(X_2) = \{x_2\}$ and $\varphi_2 = \{a_2\} \cdot O_2^* \cup \{\epsilon\}$.

Finally, we define the global specification as

$$\varphi = \mathcal{L}(\mathcal{A}) \parallel (\{a_1\} \cdot O_1^*) \parallel O_2^* \cup \{\epsilon\}.$$

Thus, we clearly have that $\varphi_1 \parallel \varphi_2 \subseteq \varphi$, regardless of $\mathcal{L}(\mathcal{B})$.

Fix the fault profile \hat{f} : $f_1(\mathbf{tr}) = \Xi^* \parallel O_1^*$ and $f_2(\mathbf{tr}) = O_2^*$.

Consider the trace \mathbf{tr} of length 1, where for the first letter we have $x_1 = b_1$, $x_2 = b_2$ and $x_\Xi = \xi$ for some letter $\xi \in \Xi$. Clearly $\mathbf{tr} \notin \varphi_1, \varphi_2, \varphi$; $\text{Repair}_{\mathbf{tr}}(C_1) = \varphi_1$, and $\text{CFac}_{\mathbf{tr}}^{\hat{f}}(C_2) = \{b_2\} \cdot O_2^* \cup \varphi_2$. Thus, the

set $\mathcal{C}_1 = \{C_1\}$ is fault mitigation-capable w.r.t. \hat{f} iff

$$\begin{aligned} \varphi_1 \parallel (\{b_2\} \cdot O_2^*) \subseteq \mathcal{L}(\mathcal{A}) \parallel (\{a_1\} \cdot O_1^*) \parallel O_2^* \\ \text{iff} \\ \mathcal{L}(\mathcal{B}) \parallel (\{a_1\} \cdot O_1^*) \parallel (\{b_2\} \cdot O_2^*) \subseteq \mathcal{L}(\mathcal{A}) \parallel (\{a_1\} \cdot O_1^*) \parallel O_2^* \\ \text{iff} \\ \mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{A}). \end{aligned}$$

Similarly, the set $\{C_2\}$ is *not* fault manifestation-capable iff $\varphi_1 \parallel (\{b_2\} \cdot O_2^*) \subseteq \mathcal{L}(\mathcal{A}) \parallel (\{a_1\} \cdot O_1^*) \parallel O_2^*$ iff $\mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{A})$.

Given the automata \mathcal{A} and \mathcal{B} , in time polynomial in their size we can construct NFAs for φ_1 and φ by extending their alphabet by $O_1 \parallel O_2$. The NFAs for φ_2 , $f_1(\cdot)$ and $f_2(\cdot)$ are of constant size. Thus, we can reduce $\mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{A})$ to checking fault mitigation/ fault manifestation for a suitable fault model.

F Additional Related Work

In this section we discuss some of the most prominent definitions of causes from the literature and highlight connections to our definitions of causal sets from Equations 1 through 6.

Causality for Structural Equations. The paper [12] gives a definition of “actual cause” in a setting of structural equations over Boolean variables, where the structural equations describe the causal dependencies between these variables. Actual causality is based on counterfactual as well as on contingency dependency: only contingencies that “do not interfere with active causal processes” are considered. The major difference from our (and related work) on concurrent reactive systems is that [12] assumes that a full model of the system is known. In contrast, our work deals with (concurrent) systems for which we are only given a set of correct behaviors, and *one single incorrect behavior*. In addition, most of the work of [12] deals with acyclic variable dependencies, while concurrent reactive systems are usually cyclic.

Recently, the structural equation model approach by Halpern and Pearl was extended to reason about models of event-based concurrent systems [17, 1] in which temporal logic formulas are used to describe the causal process of a violation. Similarly to the original approach, these methods do not face the challenge of constructing counterfactual executions for black-box systems, as they work with a given system model. In fact, they integrate causality checking in the model-checking process.

Necessary Causes. The work of [7], building upon their earlier work [8], presents a causality definition that takes into account the effect of changing the behavior of one component on others in a reactive setting. In [6] the reasoning based on necessary causes (fault mitigation-capable sets, in our terms) is extended to the real-time setting; here the definition of counterfactual traces requires that for each component the difference between the local counterfactual traces and the observed local trace is minimal. This difference is minimized locally for each individual component, which, a careful analysis shows, can construct local traces that are not composable. The analysis in [22] can also yield inconsistent counterfactual traces, leading to erroneous results. The fault ascription analysis in [10] is parameterized by a counterfactual operator, but only a single concrete definition based on the idea of closeness of counterfactuals to the observed behaviours is provided. In contrast, we propose and discuss several fault models.

Sufficient causes. The notion of sufficient causes has not been explored much in the literature. Definitions based on universal quantification have been studied in [8] and in [7]. Our definition of fault manifesting sets presented in Subsection 3.4, on the other hand, is an existential one and requires that some resultant behavior is faulty.

Contributory causes. [21] studies a conservative version of the so called *contributory causes* (where one is interested in the ratio between the number of traces that satisfy φ after replacing all components in

\mathcal{C} with correct ones and the total number of such alternative traces). The analysis is focused on a special case where a set of components is determined to be a cause (called culprit) iff this ratio is 1, which is essentially the same as the classical definition of necessary cause. The key difference between [21] and previous work is in the way the set of counterfactual traces is defined: it assumes that faulty components in $\bar{\mathcal{C}}$ will produce the same output letter sequence as in the observed trace \mathbf{tr} *even if their inputs change*, while components that have not violated their local properties on \mathbf{tr} will react correctly to changed input. The set of counterfactual traces obtained using our fault model $F2$ is larger, and therefore it is more difficult to characterize a set as a cause based on $F2$ than based on [21].

Root causes. The notion of *root causes* for contract violations in message-based systems has been studied in [11]. The problem studied there is different from ours, as it does not capture mitigation. The root cause is determined by the shortest non-compliant prefix of the error trace, regardless of whether this violation could have been mitigated by another component. As a consequence, there is exactly one root cause, while there can be multiple mitigation-capable sets, which will not be discovered by the algorithm in [11].

Counterexample analysis. In [20] the authors perform causality analysis of counterexample traces, relying on a single error trace and the program source, without considering counterfactual traces. The key difference to our work is that they do not reason about concurrent reactive systems, but work on the level of variables in a single program, over which they compute weakest preconditions. Other works [15, 18] derive fault trees from probabilistic counterexamples employing counterfactuals in the flavour of the notion by Pearl and Halpern. Since they also work in the single-component setting, they do not face the challenges we address.