Automatic Root Cause Analysis for large-scale dynamic networks

MOTIVATION

The digital era of the exponential growth of connections, services and communicating entities calls for an unprecedented need for understanding the causes of events that may impact performance and service availability in this complex ecosystem. The process of identifying root causes of faults or problems is called *Root Cause Analysis* (RCA). RCA is becoming more critical and challenging in the context of ad-hoc networks that are emerging in the presence of IoT, due to the lack of full knowledge of the network and to the number of potentially bad/faulty actors might be growing at a tremendous rate.

Traditional Root Cause Analysis approaches are manual, require domain knowledge and observability of faults and operate on computational expensive mathematical models which do not scale to our todays networks. Moreover, the current AI-based approaches focus on discovering event associations and patterns, instead of logically reasoning on the cause of an observation in the data.

We propose a Root Cause Analysis framework which is able to operate at scale. It is a model-based approach that learns generic fault behaviors from passive observations or controlled interventions. Scalable distributed algorithms are elaborated for fault localization and causal explanation. More specifically, our research focus covers:

- Stimulus generation framework for virtualized environments enabling the automated execution of controlled interventions on a running distributed application,
- Causal Model learning algorithms based on the analysis of stimulus-effect reactions and time analysis of intervention-free correlations;
- Model Inversion algorithms for causal probabilistic explanations given the observed alarms, based on Bayesian Network inference, hidden Markov models, probabilistic network unfolding.

FAULT MODEL-BASED DIAGNOSIS PROCESS

(Causal) Model discovery

Interventional framework for systematic fault injections and subsequent measurements (possible only in controlled environments)

Correlation analysis of observed data to reveal statistical dependencies between variables (with or without interventions)

Causal modelling for identifying causal relationships between state variables

Fault diagnosis by Model inversion

Causal Explanations: provide the most likely explanations given the observed alarms.

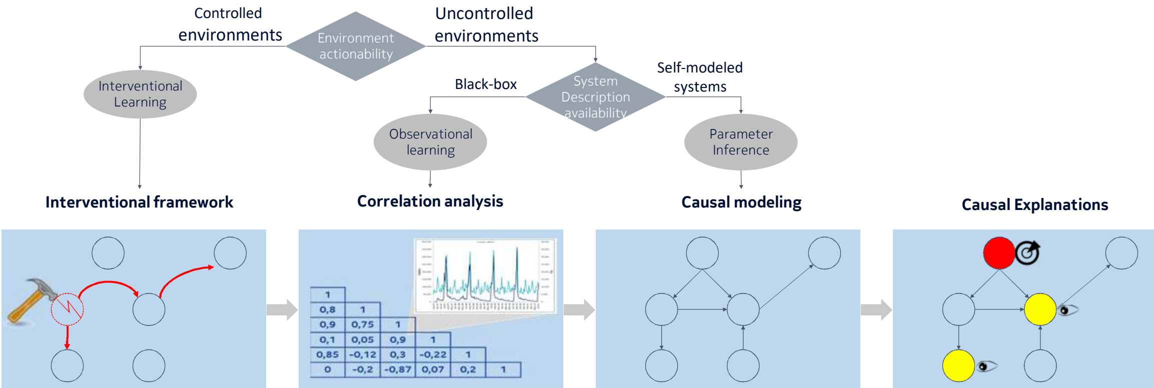

USE CASE 1

RCA for cloud applications

Stimulus creation: perturbation of available computing resources on each app node (CPU, RAM, Disk I/O, Ntwk I/O)

Correlation matrix: pairwise Pierson correlations between performance meters

Tile generation: elementary fault behaviors <pre-conditions, alarm, post-conditions>

Fault trajectories building: tile composition, Petri Net semantics

Stimulus creation | **Correlation matrix estimation** | **Tile generation (pieces of causal relationships)** | **Fault trajectories building (tile composition)**

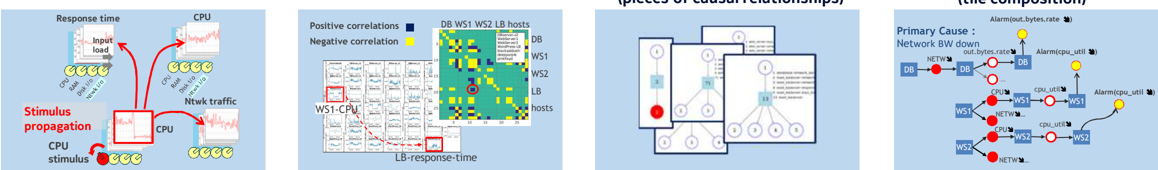

AUTOMATIC STIMULUS GENERATION FRAMEWORK

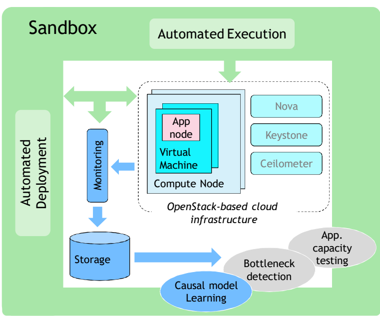

"Stimulus" creation: system-level fine-tuned control of available computing resources on a given app

- can be negative (restrict resources available to a node) or positive (launch "parasite" processes which consume resources)
- includes input load perturbations

In addition to causal model learning, can also be used for identifying resource bottlenecks, and optimizing deployment configurations

USE CASE 2

Fault localization in mobile network services

Using Bayesian Network (BN) for causal modelling and fault localization

Input data: end to end performance measurements of mobile flow records

Data model extraction: automatic discovery of structural cause-effect dependencies from mobile flow records

Parameter inference (learning conditional probabilities) using Maximum Likelihood Estimation, under a Noisy-OR gate hypothesis for multi-fault scenarios

Fault localization by computing a likelihood scores of fault hypothesis (Spark Scala)

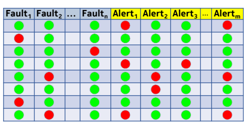

Data model extraction & Parameter inference

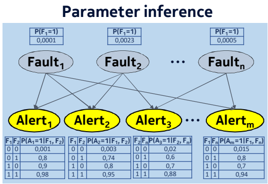

Fault localization

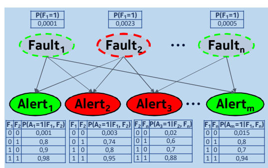

Evaluations:

>140 million rows

> 400 K possible fault locations

⇒ The BN-based algorithm successfully localizes the faults with the highest score

Fault Hypothesis Scoring

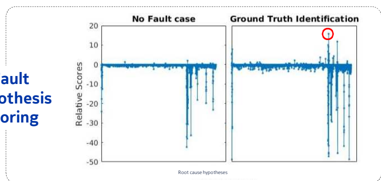

Hyper-scale Dynamic fault localization

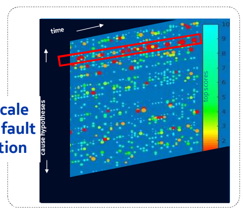

Contacts:

Armen Aghasaryan, Armen.Aghasaryan@nokia-bell-labs.com

Dimitre Kostadinov, Dimitre.Kostadinov@nokia-bell-labs.com

Makram Bouzid, Makram.Bouzid@nokia-bell-labs.com

NOKIA Bell Labs