

Anomaly Characterization Problems

Romaric Ludinard¹, Emmanuelle Anceaume², Yann Busnel³,
Erwan Le Merrer⁴, Jean-Louis Marchand⁵, Bruno Sericola¹, Gilles Straub⁴

¹*INRIA Rennes Bretagne-Atlantique, firstname.name@inria.fr*

²*CNRS UMR 6074 IRISA, emmanuelle.anceaume@irisa.fr*

³*LINA, Université de Nantes, yann.busnel@univ-nantes.fr*

⁴*Technicolor Rennes, firstname.name@technicolor.com*

⁵*Ecole Normale Supérieure de Rennes, jean-louis.marchand@ens-rennes.fr*

The context of this work is the online characterization of anomalies in large scale systems. In particular, we address the following question: Given two successive configurations of the system, can we distinguish massive anomalies from isolated ones, the former ones impacting a large number of nodes while the second ones affect solely a small number of them, or even a single one? The rationale of this question is twofold. First, from a theoretical point of view, we characterize anomalies with respect to their neighborhood, and we show that there are anomaly scenarios for which isolated and massive anomalies are indistinguishable from an omniscient observer point of view. We then relax this problem by introducing *unresolved* configurations, and exhibit necessary and sufficient conditions that allow any node to determine the type of anomaly it has been impacted by. This condition only depends on the close neighborhood of each node and thus is locally computable. From a practical point of view, distinguishing isolated anomalies from massive ones is of utmost importance for networks providers. For instance, Internet service providers (ISPs) would be interested to deploy procedures that allow gateways to self distinguish whether their dysfunction is caused by network-level anomalies or by their own hardware or software, and to notify the ISP only in the latter case.

Keywords: Network monitoring, anomaly detection, diagnosis.

We study the online monitoring problem in large scale distributed systems. This problem deals with the capability of collecting and analyzing relevant information provided by monitored devices so as to make the monitoring application continuously aware of the state of the system. Actually, standardized procedures exist at devices level to autonomously trigger investigations in presence of errors or networks events. However, these procedures are never used for practical reasons. Indeed if the cause of a QoS (quality of service) variation lies in the network itself – due to routing loops, router dysfunctions, or configuration errors – this may impact a very large number of devices (more precisely, impact services consumed by these devices), and thus letting thousands of impacted devices reporting the problem to the operator may quickly become a disaster. It is thus of utmost importance to minimize the overall pressure put on the service operator, by giving each device the capability to locally detect whether the local QoS degradation is also observed at many other devices or not, so that only isolated errors or events are reported on the fly by the devices experiencing them. In both cases, the key point is to provide each monitored device a way to estimate the impact on other devices of a locally perceived QoS degradation. The approach we propose boils down for a device to locally detect the presence of similarity features in the abnormal behavior of other devices. This is achieved by modeling the QoS of the different services accessed by a device by a point in a QoS space E , and the temporal evolution of its QoS by a trajectory in E . A trajectory is abnormal if the predicted values of the QoS differ from the observed ones. The problem we tackle amounts for a device to locally identify all the abnormal trajectories that are *close* to its own one, to determine how dense they are, to finally decide whether its services have been impacted by an isolated event or a network one.

1 System Model

We consider a set of n monitored devices, such that each one consumes d services s_1, \dots, s_d . At any discrete time k , the QoS of each service s_i at device j is locally measured with an end-to-end performance measurement function $q_{i,k}(j)$, whose range of values is $[0, 1]$. Measurement functions reflect errors (or failures) occurring on the chain of equipments and network links from the providers of consumed services to the monitored devices. We model the QoS of monitored devices at discrete time k as a set S_k of n points in a space $E = [0, 1]^d$, with $d \geq 1$, called the QoS space. The position of device j at time k is represented by point $p_k(j) = (q_{1,k}(j), \dots, q_{d,k}(j))$. The state S_k of the system at discrete time k is $S_k = (p_k(1), \dots, p_k(n))$.

Definition 1 (*r*-consistent motion) For any $r \in [0, 1/4)$, a subset $B \subseteq \llbracket 1, n \rrbracket$ has an *r*-consistent motion in the time interval $[k-1, k]$ if $\forall (i, j) \in B^2, \|p_k(i) - p_k(j)\| \leq 2r$ and $\|p_{k-1}(i) - p_{k-1}(j)\| \leq 2r$. Moreover, a subset $B \subseteq \llbracket 1, n \rrbracket$ has a maximal *r*-consistent motion in the time interval $[k-1, k]$ if B has an *r*-consistent motion in the time interval $[k-1, k]$ and $\forall j \in \llbracket 1, n \rrbracket \setminus B, B \cup \{j\}$ does not have an *r*-consistent motion in the time interval $[k-1, k]$.

Note that if B has an *r*-consistent motion in the time interval $[k-1, k]$, either B has a maximal *r*-consistent motion or there exists $B' \subseteq \llbracket 1, n \rrbracket, B \subseteq B'$ such that B' has a maximal *r*-consistent motion.

Each device j consumes d services, and for each of them, periodically computes an end-to-end quality of service which is used to feed an error detection function $a_k(j)$. If the variation of quality is considered as abnormal, this function returns `true`. The set of devices having an abnormal trajectory in the time interval $[k-1, k]$ is denoted by $A_k = \{j \in \llbracket 1, n \rrbracket \mid a_k(j) = \text{true}\}$.

Given the position of each device in the QoS space E at each time k , one can construct several plausible scenarios of errors that would explain the trajectories of each device. For instance if a group of points follow the same abnormal trajectories at different observations, it should be caused by the same error. Similarly, if some point shows an abnormal trajectory that moves it away from its previous neighbors it should be due to some isolated error. On the other hand, there are scenario of errors that cannot be captured by periodic snapshots, as for example the fact that some device has been hit by simultaneous or temporally close errors between two successive snapshots. We encapsulate these indistinguishable scenarios of errors by imposing the following restrictions on the impact of errors on devices QoS. First, in the time interval $[k-1, k]$, the abnormal trajectory of each device $j \in A_k$ is due to a single error (R1). An error has a similar effect on the abnormal trajectories of all impacted devices. In particular if a set of devices that are at no more than $2r$ from each other in E at time $k-1$ are impacted by a given error in the time interval $[k-1, k]$ then all these devices will undergo the same abnormal trajectories and thus by Definition 1 will follow the same *r*-consistent motion in $[k-1, k]$ (R2). Finally, if at least τ devices have suffered from isolated errors (possibly different ones) then they cannot form a consistent motion (R3). Note that a single error can impact devices whose QoS can be arbitrarily different.

Restrictions R1, R2 and R3 are taken into account by partitioning the set of devices in A_k . This partitioning of A_k is formally defined as follows.

Definition 2 (Anomaly partition \mathcal{P}_k) For any $k \geq 1, \tau \in \llbracket 1, n-1 \rrbracket, r \in [0, 1/4)$, the partition \mathcal{P}_k of A_k is said to be an anomaly partition at time k if it is made of non-empty and disjoint *r*-consistent motions B_1, \dots, B_ℓ that verify conditions C1 and C2 below. Subsets B_1, \dots, B_ℓ are called anomalies.

C1: $\forall B \subseteq \bigcup_{|B_i| \leq \tau} B_i$, either B has an *r*-consistent motion with $|B| \leq \tau$ or B has not an *r*-consistent motion,

C2: $\forall B \subseteq \bigcup_{|B_i| \leq \tau} B_i, \forall i \in \llbracket 1, \ell \rrbracket, B_i$ has an *r*-consistent motion with $|B| > \tau \Rightarrow B \cup B_i$ has not an *r*-consistent motion.

By extension, for any point $j \in A_k, \mathcal{P}_k(j)$ represents the (unique) subset $B \in \mathcal{P}_k$ such that $j \in B$. In spite of the apparent complexity of Definition 2, given A_k, S_{k-1}, S_k, τ and r , there always exists at least one anomaly partition. Finally, according to the number of devices belonging to each B_1, \dots, B_ℓ of \mathcal{P}_k , we differentiate between *isolated anomalies* and *massive anomalies*. Specifically,

Definition 3 (Massive / Isolated Anomalies) Let \mathcal{P}_k be an anomaly partition. An element $B \in \mathcal{P}_k$ is called a massive anomaly in the time interval $[k-1, k]$ if $|B| > \tau$. Otherwise it is called an isolated anomaly. The

Anomaly Characterization Problems

set of devices impacted by a massive anomaly in the time interval $[k-1, k]$ is denoted by $M_{\mathcal{P}_k}$. We have $M_{\mathcal{P}_k} = \{j \in A_k \mid |\mathcal{P}_k(j)| > \tau\}$. Similarly, the set of devices impacted by an isolated anomaly in the time interval $[k-1, k]$ is denoted by $I_{\mathcal{P}_k}$. We have $I_{\mathcal{P}_k} = \{j \in A_k \mid |\mathcal{P}_k(j)| \leq \tau\}$.

To summarize, if \mathcal{P}_k is an anomaly partition, then we have $A_k = M_{\mathcal{P}_k} \cup I_{\mathcal{P}_k}$ and $M_{\mathcal{P}_k} \cap I_{\mathcal{P}_k} = \emptyset$.

We consider in the following that all the errors or events that occur in the system respect restrictions R1, R2 and R3. In this (ideal) context, there exists an anomaly partition that reconstructs exactly what really happens in the system. In the following we denote by \mathcal{R}_k , $k \geq 1$, this real scenario of errors, and by respectively $M_{\mathcal{R}_k}$ and $I_{\mathcal{R}_k}$ the set of devices that have been involved in respectively massive and isolated anomalies.

2 The Addressed Problems

Consider an omniscient observer that is able to read, at any time k , the state vector S_k , and knows for any point $j \in \llbracket 1, n \rrbracket$ the output of the error detection function $a_k(j)$. Based on this knowledge, the goal of the omniscient observer is to infer the set of devices that have been involved in massive and isolated anomalies. The question that naturally crosses our mind is whether these inferred sets exactly match both $M_{\mathcal{R}_k}$ and $I_{\mathcal{R}_k}$. We reformulate this question as the Anomaly Characterization Problem (ACP). Specifically, for any $k \geq 1$, for any system states S_{k-1} and S_k , for any A_k , for any $r \in [0, 1/4]$ and $\tau \in \llbracket 1, n-1 \rrbracket$, let M_k and I_k be the two sets built by the omniscient observer that contained all the devices that have been impacted by respectively massive and isolated anomalies.

Problem 1 (Anomaly Characterization Problem (ACP)) *Is the omniscient observer always capable of building M_k and I_k such that $M_k = M_{\mathcal{R}_k}$ and $I_k = I_{\mathcal{R}_k}$ without knowing \mathcal{R}_k ?*

Unfortunately, there exist configurations that do not allow an omniscient observer to decide with certainty which devices have been impacted by massive anomalies and which ones have been impacted by isolated anomalies. Because of the existence of such configurations, Problem 1 is not solvable. We propose to relax this problem by partitioning A_k into three sets M_k , I_k and U_k such that M_k and I_k contain all the devices for which it is certain that these devices have been impacted by respectively massive and isolated anomalies. We have $I_k = \{\ell \in A_k \mid \forall \mathcal{P}_k, |\mathcal{P}_k(\ell)| \leq \tau\}$ and $M_k = \{\ell \in A_k \mid \forall \mathcal{P}_k, |\mathcal{P}_k(\ell)| > \tau\}$. Thus, whatever the anomaly partition \mathcal{P}_k , $M_k \subseteq M_{\mathcal{P}_k}$ and $I_k \subseteq I_{\mathcal{P}_k}$. In particular $M_k \subseteq M_{\mathcal{R}_k}$, $I_k \subseteq I_{\mathcal{R}_k}$. On the other hand, set U_k contains all the other devices $j \in A_k$ for which an omniscient observer cannot decide with certainty whether j belongs to a massive anomaly or an isolated one. This is formalized as follows.

Definition 4 (Unresolved configuration) *Any device $j \in A_k$ is in an unresolved configuration if there exist two anomaly partitions \mathcal{P}_k and \mathcal{P}'_k such that $j \in I_{\mathcal{P}_k}$ and $j \in M_{\mathcal{P}'_k}$. The set of devices belonging to an unresolved configuration in the time interval $[k-1, k]$ is denoted by U_k .*

We now formulate a relaxed version of ACP. Specifically, for any $k \geq 1$, for any system states S_{k-1} and S_k , for any A_k , and $\tau \in \llbracket 1, n-1 \rrbracket$, let M_k , I_k and U_k be respectively the set of devices involved in massive and isolated anomalies and those being in an unresolved configuration.

Problem 2 (Relaxed ACP) *Is the omniscient observer always capable of building M_k , I_k and U_k such that $M_k \subseteq M_{\mathcal{R}_k}$ and $I_k \subseteq I_{\mathcal{R}_k}$ and $M_k \cup I_k \cup U_k = A_k$ without knowing \mathcal{R}_k ?*

The following section presents necessary and sufficient conditions for any device $j \in A_k$ to belong to one of these three sets M_k , I_k and U_k .

3 Locally deciding whether one belongs to M_k , I_k , or U_k

A naive approach for device $j \in A_k$, $k \geq 1$, to decide whether it belongs to M_k , I_k or U_k consists in generating all admissible anomaly partitions and then in deciding whether it belongs to M_k , I_k , or U_k . Clearly this is impractical. We propose to solve the relaxed ACP through a cheaper and local computation which relies

uniquely on the knowledge of all the maximal r -consistent motions j is involved in. Theorem 1 provides a necessary and sufficient condition (NSC) for $j \in A_k$ to belong to I_k . Theorems 2 and 3 give respectively a sufficient condition and a NSC for $j \in A_k$ to belong to M_k . Finally, Corollary 4 exhibits a NSC for $j \in A_k$ to belong to U_k . We introduce the following two families.

$$\mathcal{W}_k(j) = \{B \subseteq A_k \mid j \in B, |B| > \tau, B \text{ has an } r\text{-consistent motion}\},$$

$$\overline{\mathcal{W}}_k(j) = \{B \subseteq A_k \mid j \in B, |B| > \tau, B \text{ has a maximal } r\text{-consistent motion}\}.$$

Theorem 1 For any $k \geq 1$, and for any $j \in A_k$, we have $\overline{\mathcal{W}}_k(j) = \emptyset \iff j \in I_k$.

This theorem illustrates the fact that if there are not enough other devices in the vicinity of a given device j exhibiting similar trajectories as j one, then j has necessarily been impacted by an isolated error. On the contrary, we denote by $D_k(j)$ the set of all devices having similar anomalous trajectories, and that belong to an element of $\overline{\mathcal{W}}_k(j)$. We have $D_k(j) = \bigcup_{B \in \overline{\mathcal{W}}_k(j)} B$. This set can be partitioned into two subsets $J_k(j)$ and $L_k(j)$ as follows.

$$J_k(j) = \{\ell \in A_k \mid \exists B \in \overline{\mathcal{W}}_k(j), \ell \in B \text{ and } \forall B' \in \overline{\mathcal{W}}_k(\ell), j \in B'\},$$

$$L_k(j) = \{\ell \in A_k \mid \exists B \in \overline{\mathcal{W}}_k(j), \ell \in B \text{ and } \exists B' \in \overline{\mathcal{W}}_k(\ell), j \notin B'\}.$$

Based on this neighborhood division, we enunciate the following theorems.

Theorem 2 For any time $k \geq 1$ and for any $j \in A_k$, $\exists B \in \overline{\mathcal{W}}_k(j)$ such that $|B \cap J_k(j)| > \tau \implies j \in M_k$.

Theorem 3 For any time $k \geq 1$ and for any $j \in A_k$, $j \in M_k$ if and only if $\overline{\mathcal{W}}_k(j) \neq \emptyset$ and for all collections C of pairwise disjoint sets defined by $C \subseteq \{B \in \mathcal{W}_k(\ell) \mid \ell \in L_k(j), j \notin B\}$, the following relation holds.

$$\left(\exists A \in \mathcal{W}_k(j) : A \subseteq D_k(j) \setminus \bigcup_{B \in C} B \right) \text{ or } \left(\exists B \in C : B \cup \{j\} \in \mathcal{W}_k(j) \right).$$

Corollary 4 For any time $k \geq 1$ and for any $j \in A_k$, $j \in U_k$ if and only if $\overline{\mathcal{W}}_k(j) \neq \emptyset$ and it exists a collection C of pairwise disjoint sets defined by $C \subseteq \{B \in \mathcal{W}_k(\ell) \mid \ell \in L_k(j), j \notin B\}$ such that the following relation holds.

$$\left(\forall A \in \mathcal{W}_k(j) : A \not\subseteq D_k(j) \setminus \bigcup_{B \in C} B \right) \text{ and } \left(\forall B \in C : B \cup \{j\} \notin \mathcal{W}_k(j) \right).$$

For space reasons, proofs of Theorems 1–3 and Corollary 4 are omitted from this paper but are presented in the companion paper [AB⁺14]. We have also described in [AB⁺14] the algorithms implementing these theorems, and evaluated their performance.

To summarize, we have derived conditions that allow any impacted device to decide whether many other devices have been impacted by the very same error or not. We have shown that the concomitance of errors may lead to unresolved scenarios that do not allow devices to distinguish which error they have been impacted by. Finally, we have shown that each device j only needs to know the trajectories of its neighbors (*i.e.*, the devices that belong to j maximal r -consistent motions), and possibly the trajectories of the neighbors of the devices that belong to $L_k(j)$. Thus j only needs to know the trajectories that are at no more than $4r$ from itself. A larger radius of knowledge – as the one got by an omniscient observer that samples at each time k the system state S_k – does not bring any additional information and thus does not provide a higher error detection accuracy.

References

- [AB⁺14] E. Anceaume, Y. Busnel, E. Le Merrer, R. Ludinard, J.L. Marchand and B. Sericola. Anomaly Characterization in Large Scale Networks. Research Report, hal-00948135, Feb. 2014.