International Journal of Research
Available at https://edupediapublications.org/journals

e-ISSN: 2348-6848
p-ISSN: 2348-795X
Volume 04 Issue 17
December 2017

Service Placement for Detecting and Localizing Failures End-to-End Path Measurements

N.Savitha
Assistant Professor, Department of Computer Science University College for Women, Koti, Hyderabad.

Abstract: *Based on novel performance measures capturing the coverage, the identifiability, and the distinguishability in monitoring disasters, we formulate the service placement problem as a hard and fast of combinatorial optimizations with those measures as goal features. In particular, we display that maximizing the distinguishability is equivalent to minimizing the uncertainty in failure localization. We show that each one those optimizations are NP-difficult. However, we display that the goals of insurance and distinguishability have an appropriate asset that let in them to be approximated to a consistent aspect through a greedy set of rules.*

Keywords- ing

I. INTRODUCTION

In spite of the implicit excess in server farm systems, execution issues or disappointments in the system can prompt userperceived benefit interferences. In this way, deciding and restricting client affecting accessibility and execution issues in the system in close constant is critical. Restricting system issues is finished by proficient system activities groups who work with related accessible if the need arises architects to determine issues progressively with the assistance of checking information. Such an approach can be tedious, repetitive, and is additionally exacerbated by checking commotion and the expanding size of the system. As the size of the system develops, mechanized blame restriction turns out to be progressively vital since it can decrease mean-time-to-recuperation and administration interruption.

System observing is at the core of disappointment restriction and is partitioned into two classifications: uninvolved and dynamic checking. The uninvolved approach regularly includes surveying the system gadgets intermittently to gather different telemetry information about their wellbeing and the movement that cruises by. The framework will then break down the neighborhood telemetry information and raise accessibility and execution cautions at the level of individual gadgets and connections when it recognizes any variations from the norm.

The dynamic approach depends on the capacity to infuse test movement into the system and to screen the stream of that activity. Specifically, server farms regularly run a ping administration that creates a lot of pings between sets of end has in the system and goes about as an intermediary for client saw organize accessibility and inertness [15, 16]. In this setting, ping disappointments give a solid flag that there is in fact an issue in the fundamental system and, all the more particularly, that there is an issue some place along the way from the source to the goal of the ping. In any case, pings don't pinpoint the correct gadget or connection that has made the pings come up short since real system courses are commonly obscure.

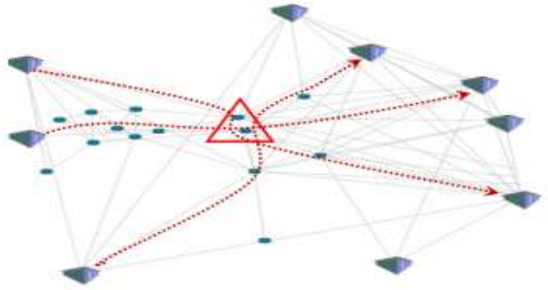

Figure 1: Failure localization approach: The successful, failed, and delayed ping data are overlaid on top of the network topology, and a statistical data mining approach is used to triangulate (i.e., localize) issues in the network

Even though a single ping failure cannot help in identifying the culprit, the conjecture is that the combination of ping data between multiple sources and destinations will allow us to triangulate the location of an existing issue. Figure 1 provides a simple visual representation of the proposed approach, which uses statistical data mining techniques to localize user-perceived network failures based on ping data.

In this capturing, we revisit the provider placement hassle from the perspective of monitoring disasters primarily based on binary service layer observations. Our contributions are:

1) We propose a set of performance measures that capture key factors of failure tracking, together with the competencies of detecting disasters (coverage), uniquely figuring out node states (identifiability), and distinguishing between candidate sets of failure locations (distinguishability).

2) Using the proposed measures as objective capabilities, we formulate the service placement hassle as a combinatorial optimization of maximizing a specific goal concern to QoS constraints. We display that the top-rated placement is NP-tough to compute for all of the above targets.

3) We show that underneath the insurance or distinguishability objective, our hassle can be solid as a submodular maximization beneath matroid constraints, which can be approximated to a factor of half of by way of a greedy set of rules. We show that although the identifiability objective isn't always submodular, it can be approximated via the distinguishability-based totally placement in the high-identifiability regime.

4) We examine the proposed algorithms on real community topologies, which shows that: (i) the proposed algorithms gain extensively better monitoring overall performance than the satisfactory-QoS placement, and (ii) the distinguishability objective ends in a carrier placement with the high-quality usual overall performance throughout all the 3 goals.

II. RELATED WORK

service placement has been considerably studied based at the facility place idea [6], where two of the most well-known formulations are: (1) the uncapacitated okay-median trouble, which optimizes the places for putting a fixed number of offerings to limit the space between clients and (closest) servers, and (2) the uncapacitated facility region hassle, which jointly optimizes the wide variety and the locations of services to limit the mixed price of website hosting and getting access to services. The traditional formulations had been extended in diverse directions, consisting of disbursed service placement in largescale networks [7] and iterative carrier placement/migration in cellular advert hoc networks [8]. the focus, however, stays on optimizing the QoS and provisioning value.

A parallel line of works take a look at the location of nodes devoted to network monitoring (known as video display units). beneath roundtrip probing (e.g., ping, traceroute) where most effective resources of probes want to be monitors, [9] suggests that the highest quality monitor placement is NP-hard and proposes a grasping approximation algorithm. under one-manner probing wherein each resources and locations need to be video display units, [10] and references therein advise polynomial-time algorithms to area a minimal quantity of video display units to uniquely localize a given number of disasters. Our hassle differs in that: (i) we only manage one endpoint (server) of every size course, and (ii) the carrier placement ought to satisfy QoS constraints. Given the states of a fixed of paths, there may be a couple of units of failure places that are steady with the course states. most current works cope with such ambiguity by nice-effort answers. as an instance, [9], [11] count on one failure at a time, [12], [4], [2] attempt to find a minimal set of screw ups that can explain all of the observed path states. Followup works attempt to enhance accuracy by means of looking for screw ups that occur with higher probabilities [13] or enhancing the accuracy in estimating direction states [3]. there may be, but, lack of expertise within the essential functionality of failure localization. lately, [5] proposes to model this capability by the most number of disasters that may be uniquely localized, known as most identifiability. It develops tight upper/lower bounds on the most identifiability and polynomial-time algorithms to compute the boundaries for numerous varieties of size paths (arbitrarily controllable, controllable however cycle-free, or uncontrollable). We look at, for the first time, the impact of provider placement at the capability of tracking disasters primarily based on end-toend measurements between customers and servers. We adopt the uncontrollable direction assumption in [5], however significantly generalize the performance measure in [5]: (i) whilst the most identifiability measure requires all the node states to be identifiable, our new identifiability measure captures cases wherein only a subset of node states are identifiable; (ii) we advocate novel measures capturing different components of failure monitoring, such as the capability of detecting screw ups and the functionality of decreasing uncertainty in failure places.

III. PROPOSED WORK

We model the service network as an undirected graph $G = (N, L)$, where N is the set of nodes, including client nodes, (candidate) server nodes, and communication nodes in between, and L is the set of communication links. Each node is associated with a binary state: normal or failed. Here a “client node” represents an access point for end-clients in the service network, and thus its state is also of interest to the service provider. We assume links do not fail, as link failures can be modeled by the failures of logical nodes that represent the links. Given a vector of node states, the set of all failed nodes is called a failure set, denoted by F. We assume that node states cannot be measured directly, but only indirectly via measurement paths. Let P be a given set of measurement paths comprising the paths between all servers and all clients interested in their services, with one path per client-server pair as determined by the underlying routing protocol employed by the network. Each path $p \in P$ is represented as a set of nodes traversed by a client-server connection, whose state is normal if and only if all the traversed nodes (including endpoints) are in normal states. We use $P_F \subseteq P$ to denote the subset of paths affected by a failure set F (i.e., traversing at least one node in F); in particular, $Pv \subseteq P$ denotes the subset of paths traversing node v.

B. Performance Measure of Network Monitoring

Given a set of measurement paths P, we quantify the value of P in monitoring node states as follows.

1) Coverage: A basic objective is to detect node failures from failures of measurement paths. Denote the set of covered nodes, i.e., nodes traversed by at least one path in P.

2) Identifiability: Besides detection, it is also important to localize the failures, i.e., determine the failure set F from observed path states. Generally, there can be multiple failure sets that generate the same path states, leading to ambiguity. The extent to which we can overcome such ambiguity thus measures our capability of localizing failures.

3) Distinguishability: Another measure of the capability of failure localization is our ability to distinguish between candidate failure sets.

C. Monitoring-Aware Service Placement

It is clear that the performance of network monitoring depends on P, the set of paths connecting servers and clients, which are determined by the network topology, the locations of clients, the routing of service requests/responses, and the positioning of services. The last parameter, positioning of services, is of particular interest as it is controlled by the service placement algorithm. We now formally define the service placement problem when taking network monitoring performance into account.

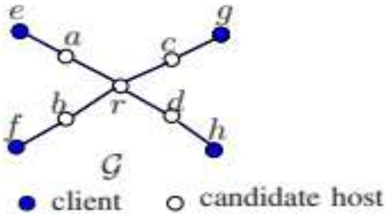

Fig. 2. Service placement example

We start by considering two prerequisites of solving problem (1): how to determine the candidate hosts and how to evaluate the objective function for a given service placement.

A. Computing Candidate Set

The goal of computing the candidate set H_s is to ensure a minimum QoS for all clients while maximizing the flexibility for service placement. As a concrete example, we consider latency as the QoS measure. Let $d(C_s, h)$ denote the maximum distance (in hop count) between node h and any client in C_s under a given routing protocol. Then a natural way of guaranteeing QoS is to impose an upper bound on $d(C_s, h)$ and only view nodes satisfying the upper bound as candidate hosts.

B. Computing Objective Function

Another prerequisite is an efficient method to evaluate the objective function in problem (1). For coverage, we can easily compute $|C(P)|$ by taking the union of all measurement paths. For distinguishability, we have to compute $|Dk(P)|$ by enumerating all pairs of failure sets and testing whether they affect the same set of measurement paths, which requires $O(|N|2k)$ tests, each of complexity $O(|P|)$ (assuming $k \ll |N|$).

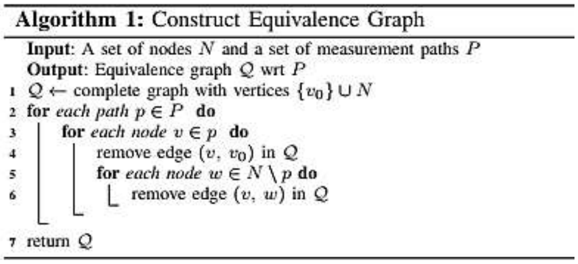

```
Algorithm 1: Construct Equivalence Graph
Input: A set of nodes $N$ and a set of measurement paths $P$
Output: Equivalence graph $Q$ wrt $P$
$Q \leftarrow$ complete graph with vertices $\{v_0\} \cup N$
for each path $p \in P$ do
    for each node $v \in p$ do
        remove edge $(v, v_0)$ in $Q$
        for each node $w \in N \setminus p$ do
            remove edge $(v, w)$ in $Q$
return $Q$
```

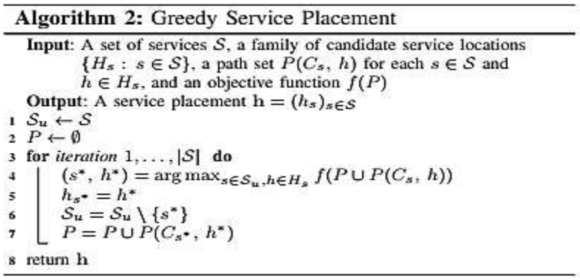

```
Algorithm 2: Greedy Service Placement
Input: A set of services $\mathcal{S}$, a family of candidate service locations
       $\{H_s : s \in \mathcal{S}\}$, a path set $P(C_s, h)$ for each $s \in \mathcal{S}$ and
       $h \in H_s$, and an objective function $f(P)$
Output: A service placement $\mathbf{h} = (h_s)_{s \in \mathcal{S}}$
$\mathcal{S}_u \leftarrow \mathcal{S}$
$P \leftarrow \emptyset$
for iteration $1, \ldots, |\mathcal{S}|$ do
    $(s^*, h^*) = \arg\max_{s \in \mathcal{S}_u, h \in H_s} f(P \cup P(C_s, h))$
    $h_{s^*} = h^*$
    $\mathcal{S}_u = \mathcal{S}_u \setminus \{s^*\}$
    $P = P \cup P(C_{s^*}, h^*)$
return $\mathbf{h}$
```

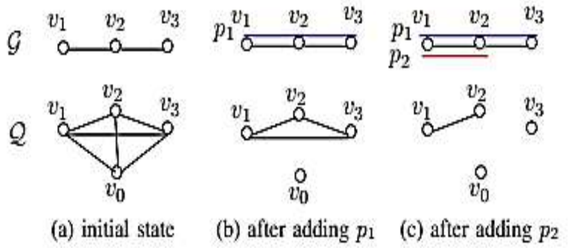

(a) initial state (b) after adding p_1 (c) after adding p_2

Fig. 3. Construct equivalence graph Q

IV. CONCLUSION

We don't forget monitoring-aware service placement, which places services inside QoS constraints such that in the face of disasters, node states may be most correctly decided from the states of quit-to-cease connections among customers and servers. Measuring overall performance by using the coverage, the identifiability, and the distinguishability in tracking failures, we solid the problem as a fixed of combinatorial optimizations, every maximizing one performance degree.

REFERENCES

[1] "Network functions virtualisation," ETSI White Paper, 2013. [Online]. Available: https://portal.etsi.org/NFV/NFV White Paper.pdf

[2] R. R. Kompella, J. Yates, A. G. Greenberg, and A. C. Snoeren, "Detection and localization of network black holes," in IEEE INFOCOM, May 2007.

[3] I. Cunha, R. Teixeira, N. Feamster, and C. Diot, "Measurement methods for fast and accurate blackhole identification with binary tomography," in ACM IMC, November 2009.

[4] N. Duffield, “Network tomography of binary network performance characteristics,” IEEE Transactions on Information Theory, vol. 52, pp. 5373–5388, 2006.

[5] L. Ma, T. He, A. Swami, D. Towsley, K. K. Leung, and J. Lowe, “Node failure localization via network tomography,” in ACM IMC, November 2014.

[6] G. Wittenburg and J. Schiller, Service Placement in Ad Hoc Networks. Springer Briefs in Computer Science, 2012.

[7] N. Laoutaris, G. Smaragdakis, and D. Oikonomou, “Distributed placement of service facilities in large-scale networks,” in IEEE INFOCOM, May 2007.

[8] G. Wittenburg and J. Schiller, “A survey of current directions in service placement in mobile ad-hoc networks,” in IEEE PerCom, March 2008.

[9] Y. Bejerano and R. Rastogi, “Robust monitoring of link delays and faults in IP networks,” in IEEE INFOCOM, March 2003.

[10] L. Ma, T. He, A. Swami, D. Towsley, and K. Leung, “On optimal monitor placement for localizing node failures via network tomography,” Elsevier Performance Evaluation, vol. 91, pp. 16–37, September 2015.

[11] J. D. Horton and A. Lopez-Ortiz, “On the number of distributed ´ measurement points for network tomography,” in ACM IMC, October 2003.

[12] N. Duffield, “Simple network performance tomography,” in ACM IMC, October 2003.

[13] H. Nguyen and P. Thiran, “The boolean solution to the congested IP link location problem: Theory and practice,” in IEEE INFOCOM, May 2007.

[14] T. He, N. Bartolini, H. Khamfroush, I. Kim, L. Ma, and T. L. Porta, “Service placement for detecting and localizing failures using end-to-end observations,” Technical Report, December 2015. [Online]. Available: http: //researcher.watson.ibm.com /researcher /files/ us-the/MaxIdentifiabilityServicePlacement report.pdf

[15] L. Ma, T. He, A. Swami, D. Towsley, and K. Leung, “Network capability in localizing node failures via end-to-end path measurements,” arXiv, 2015. [Online]. Available: http:// arxiv.org/abs /1509.06333

Author Profile

N.Savitha working as Assistant Professor, Department of Computer Science in University College for Women,Koti, Hyderabad.