

抽取引起系统性能异常波动的关键方法调用。

目前,该方法还存在如下待改进的问题:首先,动态代码注入的方式可能会造成关键方法的遗漏,后续工作将静态程序分析方法引入到插桩点定位;其次,监测方法的性能开销与插桩方法数量成正比,过多插桩会对复杂服务产生较大的性能开销,我们下一步工作将对插桩点的选取及采样频率进行研究,以较小监测开销获得较高的故障诊断准确性。

References:

- [1] Namiot D, Sneps-Snepe M. On micro-services architecture. *Int'l Journal of Open Information Technologies*, 2014,2(9):24–27.
- [2] Newman S. *Building Microservices*. Sebastopol: O'Reilly Media, Inc., 2015. 280.
- [3] Erl T. *Service-Oriented architecture: Concepts, technology and design*. <http://www.soabooks.com>.
- [4] Chandola V, Banerjee A, Kumar V. Anomaly detection: A survey. *ACM Computing Surveys*, 2009,41(3):75–79. [doi: 10.1145/1541880.1541882]
- [5] Sigelman BH, Barroso LA, Burrows M, Stephenson P, Plakal M, Beaver D, Jaspan S, Shanbhag C. *Dapper, a large-scale distributed systems tracing infrastructure*. Google Technical Report, 2010.
- [6] Chen MY, Kiciman E, Fratkin E, Fox A, Brewer E. Pinpoint: Problem determination in large, dynamic internet services. In: *Proc. of the 2002 Int'l Conf. on Dependable Systems and Networks (DSN 2002)*. IEEE, 2002. 595–604. [doi: 10.1109/DSN.2002.1029005]
- [7] Sambasivan RR, Zheng AX, De Rosa M, Krevat E, Whitman S, Stroucken M, Wang W, Xu L, Ganger GR. Diagnosing performance changes by comparing request flows. In: *Proc. of the NSDI*. 2011.
- [8] Fu Q, Lou JG, Wang Y, Li J. Execution anomaly detection in distributed systems through unstructured log analysis. In: *Proc. of the ICDM*. 2009. 149–158. [doi: 10.1109/ICDM.2009.60]
- [9] ASM. 2016. <http://asm.ow2.org/>
- [10] Instrumentation (Java Platform SE 7). <https://docs.oracle.com/javase/7/docs/api/java/lang/instrument/Instrumentation.html>
- [11] Kopetz H, Ochsenreiter W. Clock synchronization in distributed real-time systems. *IEEE Trans. on Computers*, 1987,C-36(8): 933–940. [doi: 10.1109/TC.1987.5009516]
- [12] Yao XJ, Gong DW, Li B. Evolutional test data generation for path coverage by integrating neural network. *Ruan Jian Xue Bao/Journal of Software*, 2016,27(4):828–838 (in Chinese with English abstract). <http://www.jos.org.cn/1000-9825/004973.htm> [doi: 10.13328/j.cnki.jos.004973]
- [13] Qian ZS, Miao HK. Specification-Based logic coverage testing criteria. *Ruan Jian Xue Bao/Journal of Software*, 2010,21(7): 1536–1549 (in Chinese with English abstract). <http://www.jos.org.cn/1000-9825/03615.htm> [doi: 10.3724/SP.J.1001.2010.03615]
- [14] Bille P. A survey on tree edit distance and related problems. *Theoretical Computer Science*, 2005,337(1):217–239. [doi: 10.1016/j.tcs.2004.12.030]
- [15] King JR, Jackson DA. Variable selection in large environmental data sets using principal components analysis. *Environmetrics*, 1999,10(1):67–77. [doi: 10.1002/(SICI)1099-095X(199901/02)10:1<67::AID-ENV336>3.0.CO;2-0]
- [16] Jolliffe I. *Principal Component Analysis*. Wiley Online Library, 2002.
- [17] TPC-W. <http://www.tpc.org/tpcw/default.asp>
- [18] Zhang W, Wang S, Wang W, Zhong H. Bench4Q: A QoS-oriented e-commerce benchmark. In: *Proc. of the 35th Annual Computer Software and Applications Conf. IEEE*, 2011. 38–47. [doi: 10.1109/COMPSAC.2011.14]
- [19] About the TPC. <http://www.tpc.org/information/about/abouttpc.asp>
- [20] Casale G, Mi N, Smirni E. Model-Driven system capacity planning under workload burstiness. *IEEE Trans. on Computers*, 2010, 59(1):66–80. [doi: 10.1109/TC.2009.135]
- [21] Ghanbari S, Amza C. Semantic-Driven model composition for accurate anomaly diagnosis. In: *Proc. of the Int'l Conf. on Autonomic Computing (ICAC 2008)*. 2008. 35–44. [doi: 10.1109/ICAC.2008.33]
- [22] Wang T, Wei J, Zhang W, Zhong H, Huang T. Workload-Aware anomaly detection for Web applications. *Journal of System Software*, 2014,89:19–32. [doi: 10.1016/j.jss.2013.03.060]

- [23] Wang T, Wei J, Qin F, Zhang W, Zhong H, Huang T. Detecting performance anomaly with correlation analysis for Internetware. *Science China Information Sciences*, 2013,56(8):1–15. [doi: 10.1007/s11432-013-4906-6]
- [24] Namiot D, Sneps-Snepp M. On micro-services architecture. *Int'l Journal of Open Information Technologies*, 2014,2(9):4–8.
- [25] Balalaie A, Heydarnoori A, Jamshidi P. Microservices architecture enables DevOps: Migration to a cloud-native architecture. *IEEE Software*, 2016,33(3):42–52. [doi: 10.1109/MS.2016.64]
- [26] Dragoni N, Giallorenzo S, Lafuente AL, Mazzara M, Montesi F, Mustafin R, Safina L. Microservices: Yesterday, today, and tomorrow. arXiv preprint arXiv:160604036, 2016.
- [27] Microservices resource guide. <http://martinfowler.com/microservices/>
- [28] Kang H, Chen H, Jiang G. PeerWatch: A fault detection and diagnosis tool for virtualized consolidation systems. In: *Proc. of the 7th Int'l Conf. on Autonomic Computing*. Washington: ACM Press, 2010. 119–128. [doi: 10.1145/1809049.1809070]
- [29] Jiang G, Chen H, Yoshihira K, Saxena A. Ranking the importance of alerts for problem determination in large computer systems. *Cluster Computing*, 2011,14(3):213–227. [doi: 10.1007/s10586-010-0120-0]
- [30] Pham C, Wang L, Tak B, Baset S, Tang C, Kalbarczyk Z, Iyer R. Failure diagnosis for distributed systems using targeted fault injection. *IEEE Trans. on Parallel and Distributed Systems*, 2017,28(2):503–516. [doi: 10.1109/TPDS.2016.2575829]
- [31] Wang T, Zhang W, Ye C, Wei J, Zhong H, Huang T. FD4C: Automatic fault diagnosis framework for Web applications in cloud computing. *IEEE Trans. on Systems, Man, and Cybernetics: Systems*, 2016,46(1):61–75. [doi: 10.1109/TSMC.2015.2430834]
- [32] Chandola V, Banerjee A, Kumar V. Anomaly detection: A survey. *ACM Computing Surveys (CSUR)*, 2009,41(3):15. [doi: 10.1145/1541880.1541882]
- [33] Pertet S, Narasimhan P. Causes of Failure in Web Applications. Parallel Data Laboratory, Carnegie Mellon University, 2005. 48–54.
- [34] Kiciman E, Fox A. Detecting application-level failures in component-based internet services. *IEEE Trans. on Neural Networks*, 2005,16(5):1027–1041. [doi: 10.1109/TNN.2005.853411]
- [35] Xu W, Huang L, Fox A, Patterson D, Jordan MI. Detecting large-scale system problems by mining console logs. In: *Proc. of the ACM SIGOPS the 22nd Symp. on Operating Systems Principles*. ACM Press, 2009. 117–132. [doi: 10.1145/1629575.1629587]
- [36] Zou H, Hastie T, Tibshirani R. Sparse principal component analysis. *Journal of Computational & Graphical Statistics*, 2012, 2007(Special):1–30. [doi: 10.1198/106186006X113430]
- [37] Kubernetes—Production-Grade Container Orchestration. <http://kubernetes.io/>
- [38] Netflix Open Source Software Center. <https://netflix.github.io/>
- [39] Twitter's finagle library. <https://twitter.github.io/finagle/>
- [40] Rajagopalan S, Jamjoom H. App-Bisect: Autonomous healing for microservice-based apps. In: *Proc. of the Usenix Conf. on Hot Topics in Cloud Computing*. 2015.
- [41] Heorhiadi V, Rajagopalan S, Jamjoom H, Reiter MK, Sekar V. Gremlin: Systematic resilience testing of microservices. In: *Proc. of the 36th Int'l Conf. on Distributed Computing Systems (ICDCS)*. IEEE, 2016. 57–66. [doi: 10.1109/ICDCS.2016.11]
- [42] Aguilera MK, Mogul JC, Wiener JL, Reynolds P, Muthitacharoen A. Performance debugging for distributed systems of black boxes. *ACM SIGOPS Operating Systems Review*, 2003,37(5):74–89. [doi: 10.1145/1165389.945454]
- [43] Reynolds P, Wiener JL, Mogul JC, Aguilera MK, Vahdat A. WAP5: Black-Box performance debugging for wide-area systems. In: *Proc. of the 15th Int'l Conf. on World Wide Web*. ACM Press, 2006. 347–356. [doi: 10.1145/1135777.1135830]
- [44] Bahl P, Chandra R, Greenberg A, Kandula S, Maltz DA, Zhang M. Towards highly reliable enterprise network services via inference of multi-level dependencies. In: *Proc. of the ACM SIGCOMM Computer Communication Review*. ACM Press, 2007. 13–24. [doi: 10.1145/1282380.1282383]
- [45] Reynolds P, Killian CE, Wiener JL, Mogul JC, Shah MA, Vahdat A. Pip: Detecting the unexpected in distributed systems. In: *Proc. of the NSDI*. 2006. 115–128.
- [46] Gschwind T, Eshghi K, Garg PK, Wurster K. Webmon: A performance profiler for Web transactions. In: *Proc. of the 4th IEEE Int'l Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems (WECWIS 2002)*. IEEE, 2002. 171–176. [doi: 10.1109/WECWIS.2002.1021256]

- [47] Fonseca R, Porter G, Katz RH, Shenker S, Stoica I. X-trace: A pervasive network tracing framework. In: Proc. of the 4th USENIX Conf. on Networked Systems Design & Implementation. USENIX Association, 2007. 20.
- [48] Barham P, Isaacs R, Mortier R, Narayanan D. Magpie: Online modelling and performance-aware systems. In: Proc. of the HotOS. 2003. 85–90.
- [49] Barham P, Donnelly A, Isaacs R, Mortier R. Using magpie for request extraction and workload modelling. In: Proc. of the OSDI. 2004. 18–27.
- [50] Thereska E, Salmon B, Strunk J, Wachs M, Abd-El-Malek M, Lopez J, Ganger GR. Stardust: Tracking activity in a distributed storage system. In: Proc. of the ACM SIGMETRICS Performance Evaluation Review. ACM Press, 2006. 3–14. [doi: 10.1145/1140277.1140280]
- [51] Chow M, Meisner D, Flinn J, Peek D, Wenisch TF. The mystery machine: End-to-End performance analysis of large-scale Internet services. In: Proc. of the 11th USENIX Symp. on Operating Systems Design and Implementation (OSDI 2014). 2014. 217–231.
- [52] Mace J, Roelke R, Fonseca R. Pivot tracing: Dynamic causal monitoring for distributed systems. In: Proc. of the 25th Symp. on Operating Systems Principles. ACM Press, 2015. 378–393. [doi: 10.1145/2815400.2815415]
- [53] Ghanbari S, Hashemi AB, Amza C. Stage-Aware anomaly detection through tracking log points. In: Proc. of the 15th Int'l Middleware Conf. ACM Press, 2014. 253–264. [doi: 10.1145/2663165.2663319]
- [54] Traeger A, Deras I, Zadok E. DARC: Dynamic analysis of root causes of latency distributions. In: Proc. of the ACM SIGMETRICS Performance Evaluation Review. ACM Press, 2008. 277–288. [doi: 10.1145/1375457.1375489]

附中文参考文献:

- [12] 姚香娟, 巩敦卫, 李彬. 融入神经网络的路径覆盖测试数据进化生成. 软件学报, 2008, 19(7): 1565–1580. <http://www.jos.org.cn/1000-9825/004973.htm> [doi: 10.13328/j.cnki.jos.004973]
- [13] 钱忠胜, 缪淮扣. 基于规格说明的若干逻辑覆盖测试准则. 软件学报, 2010, 21(7): 1536–1549. <http://www.jos.org.cn/1000-9825/03615.htm> [doi: 10.3724/SP.J.1001.2010.03615]



王子勇(1989—),男,山东临沂人,硕士生,主要研究领域为分布式软件与理论,软件故障检测.



陈宁江(1975—),男,博士,副教授,CCF 高级会员,主要研究领域为网络分布计算,软件工程,云计算.



王焘(1982—),男,博士,助理研究员,CCF 专业会员,主要研究领域为软件运行时监控,软件故障检测,自主计算,云计算.



左春(1959—),男,研究员,主要研究领域为分布式软件与理论,云计算,软件工程.



张文博(1976—),男,博士,研究员,博士生导师,CCF 专业会员,主要研究领域为分布式计算,云计算.