

课题报告

- 报告组成（软件+计算机）：
 - 研究背景 + 动机 + 提出方法 + 结果验证 + 结论
 - 篇幅：推荐不超过5页。
 - 提交时间：2018/05/01
 - 提交方法：邮件laiping@tju.edu.cn.
- 讲演（计算机学院）：每人十分钟
 - 10 分钟 = 8 分钟讲演+ 2 分钟问答

研究主题

- **包括**但不仅限于以下主题：
 - **云计算**应用开发: 任何基于云计算系统的应用开发工作
 - 数据处理/安全应用/人工智能/专业领域云应用/游戏等。
 - **论文**阅读 (仅限CCF-A/B类期刊会议英文论文) :
 - 云计算服务的响应时间、尾延迟优化;
 - 云计算系统的利用效率优化;
 - 云服务性能评测;
 - 云资源调度;
 - etc.

- [1] Virajith Jalaparti, Peter Bodik, Srikanth Kandula, et al. 2013. Speeding up distributed request-response workflows. SIGCOMM '13. ACM, 219-230.
- [2] Lalith Suresh, Marco Canini, Stefan Schmid, et al., 2015. C3: cutting tail latency in cloud data stores via adaptive replica selection. NSDI'15. USENIX Association, 513-527.
- [3] Darja Krushevskaja, Mark Sandler. 2013. Understanding latency variations of black box services. WWW'13, ACM, 703-714.
- [4] Sriram Govindan, Jie Liu, Aman Kansal, et al. 2011. Cuanta: quantifying effects of shared on-chip resource interference for consolidated virtual machines. SOCC '11, 22:1-22:14.
- [5] Xiangping Bu, Jia Rao, and Cheng-zhong Xu. 2013. Interference and locality-aware task scheduling for MapReduce applications in virtual clusters. HPDC '13, 227-238.
- [6] Kit Eaton. 2018. <https://www.fastcompany.com/1825005/how-one-second-could-cost-amazon-16-billion-sales>.
- [7] Jialin Li, Naveen Kr. Sharma, Dan R. K. Ports, and Steven D. Gribble. 2014. Tales of the Tail: Hardware, OS, and Application-level Sources of Tail Latency. SoCC '14. 9:1-9:14.
- [8] Fei Xu, Fangming Liu, Linghui Liu, et al., 2014. iAware: Making Live Migration of Virtual Machines Interference-Aware in the Cloud. IEEE Trans. Comput. 63(12): 3012-3025.
- [9] Sebastian Angel, Hitesh Ballani, Thomas Karagiannis, et al. 2014. End-to-end performance isolation through virtual datacenters. OSDI'14. USENIX Association, 233-248.
- [10] R. M. Smith, K. S. Trivedi and A. V. Ramesh, 1988. Performability analysis: measures, an algorithm, and a case study. IEEE Transactions on Computers, 37(4): 406-417.
- [11] Jeffrey Dean and Luiz André Barroso. 2013. The tail at scale. Commun. ACM, 56(2):74-80.
- [12] Christina Delimitrou and Christos Kozyrakis. 2014. Quasar: resource-efficient and QoS-aware cluster management. ASPLOS '14, 127-144.
- [13] L. Barrosa, J. Clidaras, U. Holzle, The Datacenter as a Computer (2nd Edition), July, 2013.
- [14] Abhishek Verma, Luis Pedrosa, Madhukar Korupolu, et al.. 2015. Large-scale cluster management at Google with Borg. EuroSys'15, 18:1-18:17.
- [15] Krste Asanović, 2014. FireBox: A Hardware Building Block for 2020 Warehouse-Scale Computers. FAST'14, Usenix.
- [16] Jiuyue Ma, Xiufeng Sui, Ninghui Sun, et al., 2015. Supporting differentiated services in computers via programmable architecture for resourcing-on-demand (pard). ASPLOS'15, ACM SIGPLAN Notices, 131-143.
- [17] Hailong Yang, Alex Breslow, Jason Mars, and Lingjia Tang. 2013. Bubble-flux: precise online QoS management for increased utilization in warehouse scale computers. ISCA '13, 607-618.
- [18] K. LaCurts, J. C. Mogul, H. Balakrishnan, and Y. Turner. Cicada: Introducing predictive guarantees for cloud networks. HotCloud'14, 1-14.
- [19] Chuanxiong Guo, Guohan Lu, Helen J. Wang, et al. 2010. SecondNet: a data center network virtualization architecture with bandwidth guarantees. Co-NEXT '10, 15:1-15:12.
- [20] Hitesh Ballani, Paolo Costa, Thomas Karagiannis, and Ant Rowstron. 2011. Towards predictable datacenter networks. SIGCOMM'11, 242-253.
- [21] Di Xie, Ning Ding, Y. Charlie Hu, et al. 2012. The Only Constant is Change : Incorporating Time-Varying Network Reservations in Data Centers. Sigcomm'12, 199-210.
- [22] Lenar Yazdanov, Maxim Gorbunov, and Christof Fetzer. 2015. EHadoop: Network I/O Aware Scheduler for Elastic MapReduce Cluster. IEEE CLOUD '15, 821-828.
- [23] Katrina LaCurts, Shuo Deng, Ameesh Goyal, and Hari Balakrishnan. 2013. Choreo: network-aware task placement for cloud applications. IMC '13, 191-204.
- [24] Peng Qin, Bin Dai, Benxiong Huang, Guan Xu. 2017. Bandwidth-Aware Scheduling with SDN in Hadoop: A New Trend for Big Data. IEEE Systems Journal, 11(4): 2337-2344 .
- [25] Haishan Zhu and Mattan Erez. 2016. Dirigent: Enforcing QoS for Latency-Critical Tasks on Shared Multicore Systems. ASPLOS '16, 33-47.
- [26] Ripal Nathuji, Aman Kansal, and Alireza Ghaffarkhah. 2010. Q-clouds: managing performance interference effects for QoS-aware clouds. EuroSys '10, 237-250.