

Jiacheng Ma

jcma@umich.edu • +1 (734) 834-8682 • <https://jcma.me>

EDUCATION	University of Michigan , MI, United States ▪ Ph.D. Student in Computer Science ▪ Advisor: Prof. Baris Kasikci	Sep 2018 – Present
	Shanghai Jiao Tong University , Shanghai, P.R.China ▪ B.E. in Software Engineering ▪ Advisor: Prof. Zhengwei Qi	Sep 2014 – Jun 2018
PUBLICATIONS	[1] Jiacheng Ma , Xiao Zheng, Wentai Li, Yaozu Dong, Zhengwei Qi, Bingsheng He, Haibing Guan, “gMig: Efficient GPU Live Migration Optimized by Software Dirty Page for Full Virtualization” in <i>VEE</i> , Mar 2018. [2] Mochi Xue, Jiacheng Ma , Wentai Li, Kun Tian, Yaozu Dong, Jinyu Wu, Zhengwei Qi, Bingsheng He, Haibing Guan, “Scalable GPU Virtualization with Dynamic Sharing of Graphics Memory Space” in <i>TPDS</i> . [3] Mochi Xue, Kun Tian, Yaozu Dong, Jiacheng Ma , Jiajun Wang, Zhengwei Qi, Bingsheng He and Haibing Guan, “gScale: Scaling up GPU Virtualization with Dynamic Sharing of Graphics Memory Space” in <i>USENIX ATC</i> , Jun 2016.	
WORKING EXPERIENCE	Software Developer Intern in Intel Open Source Technology Center ▪ Advisor: Dr. Yaozu Dong ▪ Project: GPU Virtualization	Jul 2016 – Jun 2018
	Teaching Assistant in School of Software, Shanghai Jiao Tong University ▪ SE117 (Programming and Data Structure)	Spring 2016
PROJECT EXPERIENCE	Reproducing Bugs with Symbolic Execution ▪ This project intends to use symbolic execution to reproduce bugs for complex systems.	Aug 2019 – Present
	FPGA Virtualization ▪ This on-going project intends to isolate FPGAs and provide different virtualized accelerators to different virtual machines.	Apr 2018 – Present
	NVME/NVDIMM as Memory in KVM ▪ This work explores NVME/NVDIMM (Intel Optane) as main memory for KVM guests. This work can break the hardware limitation of host memory size and provide memory overbooking for guests with near-native performance.	Aug 2017 – Jun 2018
	vGPU Live Migration for Intel GVT-g ▪ Intel GVT-g is an open-source KVM/Xen based full GPU virtualization solution; however, it lacks the feature of live migration. This work enables live migrating vGPUs for cloud applications such as virtual desktops, cloud gaming farms, cloud transcoding services, etc.	Dec 2016 – Jul 2017
	gScale: Scaling up GPU Virtualization ▪ gScale scales up the maximum number of vGPUs in Intel GVT-g and minimizes the performance impact. In Haswell architecture, Intel GVT-g only supports 3 guests; gScale extends it for 5x. This work makes GVT-g more consolidated, since more VMs with vGPU can be deployed on one physical machine.	Jan 2016 – Mar 2017
SKILLS & INTERESTS	Language: C, C++, Verilog, System Verilog Interest: Operating System, System Virtualization, Heterogeneous System	