





























- [39] Lena E. Olson, Jason Power, Mark D. Hill, and David A. Wood. 2015. Border Control: Sandboxing Accelerators. In *Proceedings of the 48th International Symposium on Microarchitecture (MICRO '15)*. 470–481.
- [40] Meni Orenbach, Pavel Lifshits, Marina Minkin, and Mark Silberstein. 2017. Eleos: ExitLess OS Services for SGX Enclaves. In *12th European Conference on Computer Systems (EuroSys '17)*. 238–253.
- [41] PCI-SIG. 2004. *PCI Local Bus Specification Specification, Revision 3.0*. Technical Report. PCI-SIG, Beaverton, OR, USA.
- [42] PCI-SIG. 2009. *Address Translation Services Specification, Revision 1.1*. Technical Report. PCI-SIG, Beaverton, OR, USA.
- [43] PCI-SIG. 2010. *PCI Express Base Specification Specification, Revision 3.0*. Technical Report. PCI-SIG, Beaverton, OR, USA.
- [44] Bharath Pichai, Lisa Hsu, and Abhishek Bhattacharjee. 2014. Architectural Support for Address Translation on GPUs: Designing Memory Management Units for CPU/GPUs with Unified Address Spaces. In *The 19th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '14)*. 743–758.
- [45] Roberto Di Pietro, Flavio Lombardi, and Antonio Villani. 2016. CUDA Leaks: A Detailed Hack for CUDA and a (Partial) Fix. *ACM Transactions on Embedded Computing Systems (TECS)* 15, 1, Article 15 (Feb 2016), 25 pages.
- [46] Phillip W. Rogaway. 2006. Method and Apparatus for Facilitating Efficient Authenticated Encryption. Patent No. U.S. 7,046,802, Filed July 30th., 2001, Issued May 16th., 2006.
- [47] Phil Rogers. 2013. Heterogeneous System Architecture Overview. In *A Symposium on High Performance Chips (Hot Chips '13)*. 1–41.
- [48] Nikolay Sakharnykh. 2017. Unified Memory on Pascal and Volta. <http://on-demand.gputechconf.com/gtc/2017/presentation/s7285-nikolay-sakharnykh-unified-memory-on-pascal-and-volta.pdf> GPU Technology Conference '17.
- [49] Darmawan Salihun. 2014. System Address Map Initialization in x86/64 Architecture Part 2: PCI Express-Based Systems. Retrieved Jan 2, 2019 from <http://resources.infosecinstitute.com/system-address-map-initialization-x86x64-architecture-part-2-pci-express-based-systems/>
- [50] Yusuke Suzuki, Shinpei Kato, Hiroshi Yamada, and Kenji Kono. 2014. GPUvm: Why Not Virtualizing GPUs at the Hypervisor?. In *2014 USENIX Annual Technical Conference (USENIX ATC '14)*. 109–120.
- [51] Giorgos Vasiliadis, Elias Athanasopoulos, Michalis Polychronakis, and Sotiris Ioannidis. 2014. PixelVault: Using GPUs for Securing Cryptographic Operations. In *ACM SIGSAC Conference on Computer and Communications Security (CCS '14)*. 1131–1142.
- [52] Stavros Volos, Kapil Vaswani, and Rodrigo Bruno. 2018. Graviton: Trusted Execution Environments on GPUs. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI '18)*. 681–696.
- [53] Samuel Weiser and Mario Werner. 2017. SGXIO: Generic Trusted I/O Path for Intel SGX. In *ACM Conference on Data and Application Security and Privacy (CODASPY '17)*. 261–268.
- [54] Sheng Yang. 2008. Extending KVM with new Intel Virtualization Technology. [https://www.linux-kvm.org/images/c/c7/KvmForum2008%24kdf2008\\_11.pdf](https://www.linux-kvm.org/images/c/c7/KvmForum2008%24kdf2008_11.pdf) KVM Forum.
- [55] Hangchen Yu and Christopher J. Rossbach. 2017. Full Virtualization for GPUs Reconsidered. In *14th Annual Workshop on Duplicating, Deconstructing, and Debunking (WDDD '17)*. 1–11.
- [56] Zhe Zhou, Wenrui Diao, Xiangyu Liu, Zhou Li, Kehuan Zhang, and Rui Liu. 2017. Vulnerable GPU Memory Management: Towards Recovering Raw Data from GPU. *Proceedings on Privacy Enhancing Technologies (PoPETs) 2017*, 2 (2017), 57–73.
- [57] Zongwei Zhou, Virgil D. Gligor, James Newsome, and Jonathan M. McCune. 2012. Building Verifiable Trusted Path on Commodity x86 Computers. In *Symposium on Security and Privacy (SP '12)*. 616–630.
- [58] Zhiting Zhu, Sangman Kim, Yuri Rozhanski, Yige Hu, Emmett Witchel, and Mark Silberstein. 2017. Understanding The Security of Discrete GPUs. In *Proceedings of the General Purpose GPUs (GPGPU '10)*. 1–11.