Title: Attacks and Defenses for Intel SGX

Speaker: Taesoo Kim

Abstract: The Intel Software Guard Extensions (SGX)---a game-changing feature introduced in the recent Intel Skylake CPU---is a new technology likely to make secure and trustworthy computing in a hostile environment practical. At a high level, SGX consists of a set of new instructions that can be used to create secure regions (i.e., enclaves) to defeat attacks that aim to steal or tamper with the data within an enclave. Without a doubt, we expect that SGX will allow developers to protect sensitive code and data from unauthorized access or modification by software running at higher privilege levels such as an OS or a hypervisor.

However, SGX is merely a set of instructions; it lacks support from the OS and libraries. These deficiencies allow programmers to easily introduce naive yet preventable bugs that often lead to critical security holes in an enclave program. Further, designing a correct and secure SGX infrastructure is also far from straightforward; enclave programs rely on the support of an underlying OS, but their security models exclude the OS from the TCB. This unconventional dependency makes various attack vectors, which are often considered impractical in a traditional setting, immediate and practical, especially in a cloud environment.

In this tutorial, I will first provide a security-focused introduction to Intel SGX, including its internal mechanisms to implement the security features. Then, I will explain known security concerns, including recent research outputs from the community: e.g., memory safety issues, cache/branch side-channel attacks, and rowhammer attacks. I will not just only demonstrate these attacks but also guide you to the proper defense mechanisms.

Bio: Taesoo Kim is a Catherine M. and James E. Allchin Early Career, Assistant Professor in the School Computer Science at Georgia Tech. He also serves as the director of the Georgia Tech Systems Software and Security Center (GTS3). He is genuinely interested in building a system that has underline principles for why it should be secure. Those principles include the design of a system, analysis of its implementation, and clear separation of trusted components. His thesis work, in particular, focused on detecting and recovering from attacks on computer systems. He has developed tools that would detect intrusion and discover which parts of

the operating system could have been affected, allowing a systems administrator to recover from an attack without excessive manual effort. His thesis work has been a foundation of a company, Nerati, where he has co-founded with colleagues during his graduate study. He holds a BS from KAIST (2009), a SM (2011) and a Ph.D. (2014) from MIT, all in CS.