Tensor Computing for Post-Quantum Encryption

Jonathan Mullin
Chief Scientist
DCI Solutions

To achieve highest efficiency for computation specialized hardware namely Fixed Programmable Gate Arrays (FPGA) or Application-Specific Integrated Circuits (ASICs) are used to reduce the SWAP of a system. While this hardware has been effective, the consumer market products have encroached on their performance per watt with the rise of high-performance Single Instruction Multiple Data (SIMD) architectures. This coupled with the ease of programing SIMD architectures make them a viable platform for development of a new low SWAP encryption platforms. Additionally, recent advances in tensor-based error correcting algorithms have pushed the boundary of compute per watt. These advances allow for a tuning of hardware and software meeting the forthcoming NIST-approved quantum resistant algorithms. Our interest lies in particular with The Learning with Errors (LWE) algorithms.

LWE based schemes are appealing in theory and for practical reasons. On the theoretical side, LWE based schemes offer a very strong security guarantee. The LWE problem is equivalent to the problem of decoding random linear codes, a problem that has been extensively studied in the last half century. The fastest known algorithms run in exponential time and unlike most number-theoretic problems used in cryptography, the LWE problem does not succumb to known quantum algorithms. On the practical side, LWE based schemes are often extremely simple and efficient in terms of code-size as well as time and memory requirements. This makes them prime candidates for light-weight/low power devices like RFID tags, which are too weak to implement standard cryptographic primitives like the AES block-cipher. This proven ability in low power systems makes LWE an excellent candidate for low SWAP embedded encryption devices.

With the hardware advances these algorithms are well suited to use the latest C based compilers for SIMD architecture optimization. Our initial test bed will most likely be an NVIDIA system, however the offerings from AMD and Intel are compelling options. One benefit of the SIMD architecture is the rapid development of the competing platforms. With the current race for high end HPC dominance in the exascale computing initiatives, the efficiency of these chips is constantly improving. As lower end models become available in the following years they benefit from the enormous R&D efforts yielding remarkably efficient embedded devices. An example is the model board we initially have been using is based on a six-year-old architecture, slightly larger models using the latest architecture have a 4x improvement on performance/Watt. We are just starting to explore the use of tensor cores for this area after recent advances yielding double the throughput of matrix-matrix multiplications at half the watt/s. By focusing on power efficiency and tunable error tolerance allowed by the LWE approach we believe this coupling of hardware and software can yield efficient post quantum cryptographic devices on commodity hardware. NVIDIA has taken the lead in the embedded space thus far, and has had several platforms hardened for use in currently deployed DoD systems.