Achieving 128-bits of Security on the Optalysys Accelerator for TFHE

Emmet Friel, Erin Hales, Florent Michel, and Thanh-Huyen Nguyen
Optalysys

Workflow

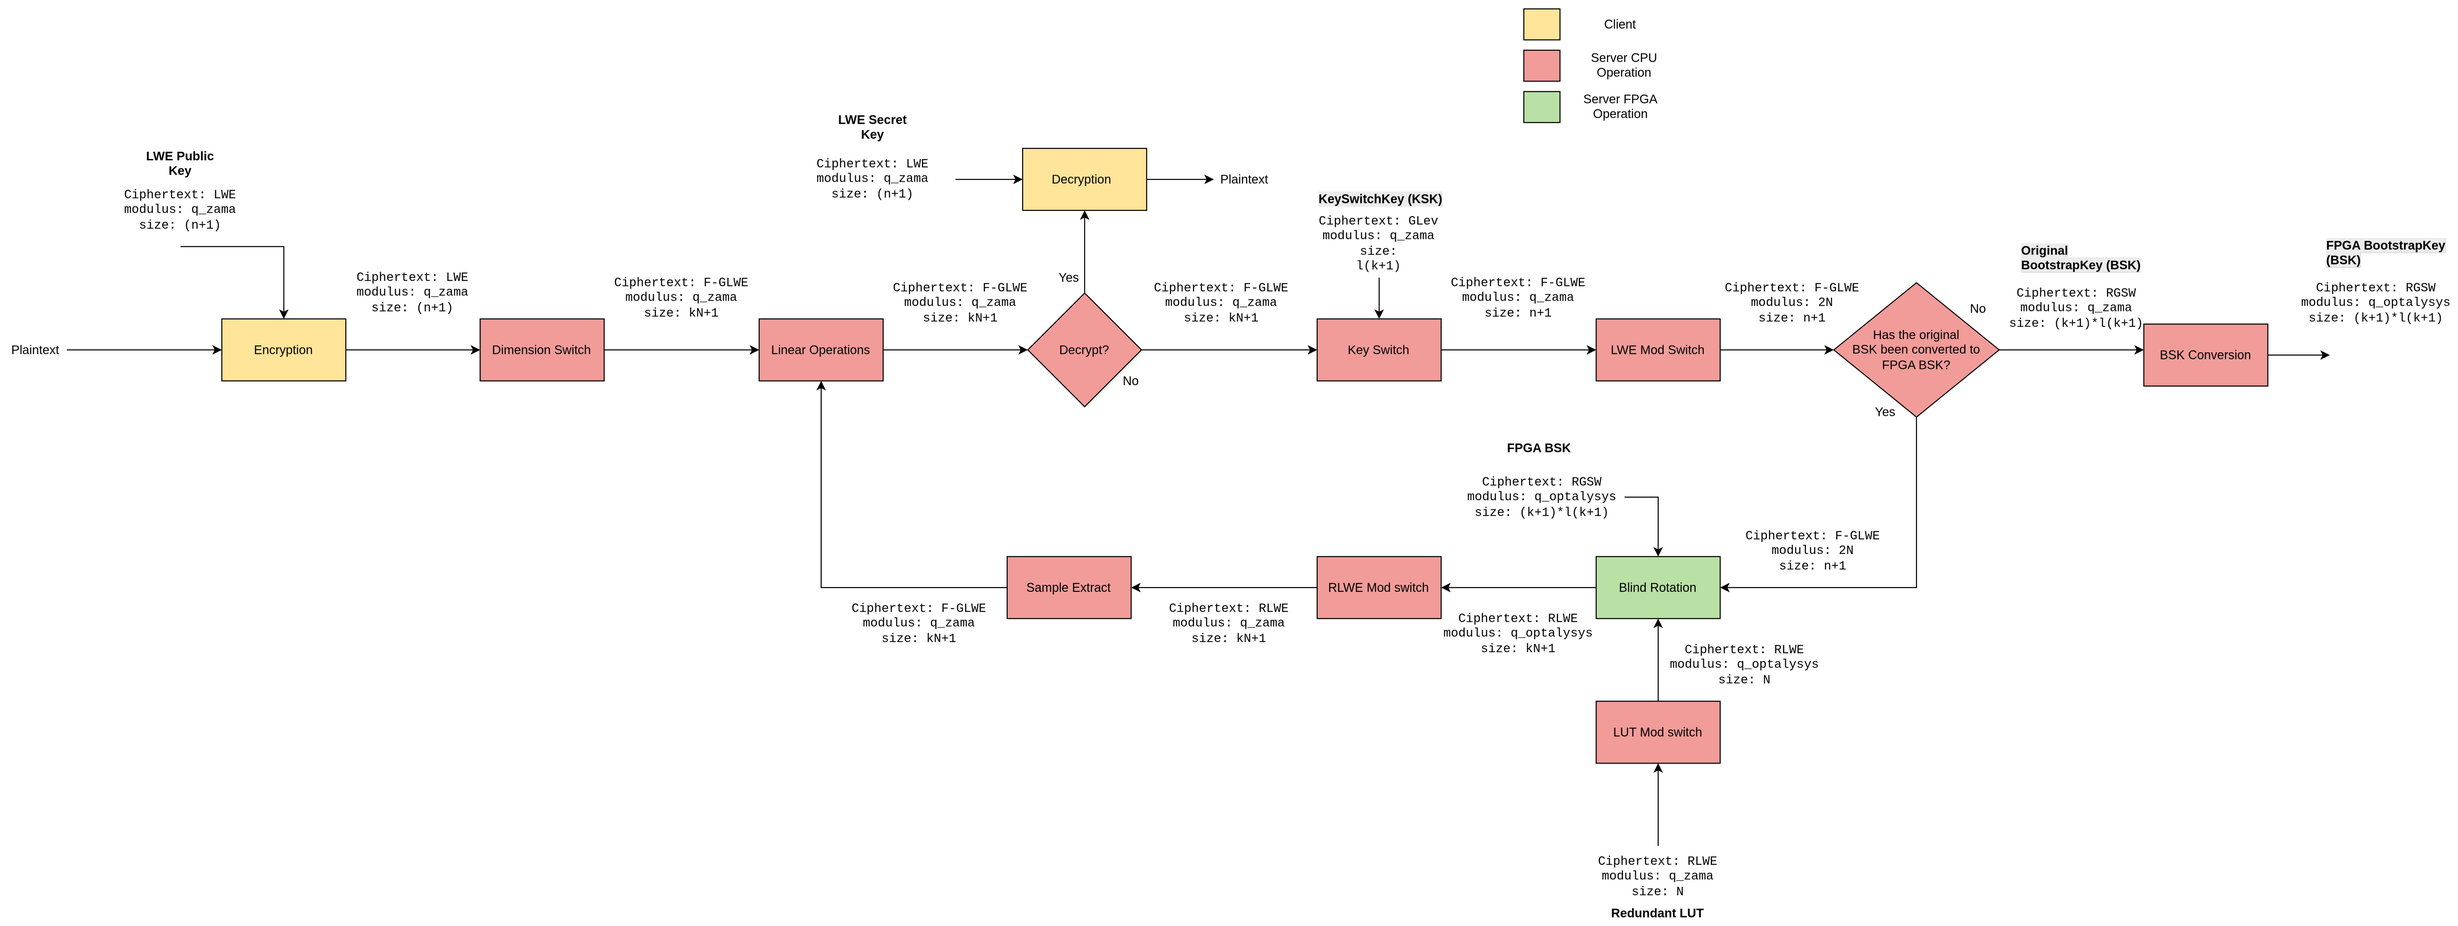

Motivation

Selecting secure, performant and correct parameters is not trivial. In order to select appropriate parameters, rigorous analysis of noise propagation and failure probability is required. We demonstrate this analysis for TFHE parameters in the context of practical deployment of FPGA acceleration. We particularly focus on additional noise introduced by pre- and post-processing stages required for compatibility with FPGA architecture. Our parameter sets target 128 bits of security with a decryption failure probability of 2^{-128}. Such a low failure probability is vital for the current application of the Optalysys FPGA accelerator to ERC20 smart contracts.

FPGA implementation

Our novel FPGA-based accelerator for the TFHE scheme targets programmable bootstrapping (PBS) operations. The computationally intensive bootstrapping step is outsourced to dedicated hardware (AMD Alveo V80 FPGA cards) and other homomorphic operations are carried out on CPU. Client-side operations are handled by the TFHE-rs external library. FPGA implementation introduces two constraints that influence parameter selection.

- Some parameters are fixed by the design, and cannot be changed at run-time
- Using Number Theoretic Transform (NTT) for efficient polynomial multiplication, this imposes constraints on Q, the RLWE ciphertext modulus

Example Parameters

Parameter	Value	Parameter	Value
Q_{FPGA}	$2^{64} - 2^{32} + 1$	l_{PBS}	1
Q	2^{64}	l_{KS}	6
N	2^{11}	σ_{LWE}	4.06274×10^{13}
n	837	σ_{GGSW}	7.56745×10^{4}
k	1	carry modulus	4
$\log_2(\beta_{\mathsf{PBS}})$	23	message modulus	4
$\log_2(\beta_{\mathsf{KS}})$	3	padding bits	1

This is an example parameter set which provides 128 bit security and a probability of decryption failure below 2^{-128}.

Parameter Selection

Several sets of default TFHE-rs parameters were evaluated to assess their security level and probability of decryption failure. Inheriting default TFHE-rs CPU parameters gives a probability of decryption failure above 2^{-128} due to the pre- and post-processing stages required for FPGA acceleration.

We sketch the noise analysis required to calculate the corresponding probability of decryption error using an average-case heuristic, using the PBS workflow. Let ν be defined as the maximum value of the sum of squares of integer weights used in dot products before the PBS. Assuming all sources of noise are statistically independent, the error after PBS has variance given by:

$$\sigma^2_{\mathsf{PBS}} = \nu^2 \cdot (\sigma^2_{\mathsf{MSBSK}} + \sigma^2_{\mathsf{SGLWE_{LUT}}} + \sigma^2_{\mathsf{MSGLWE}}) + \sigma^2_{\mathsf{KS}} + \sigma^2_{\mathsf{MSLWE}},$$

where the following terms are specific to our FPGA implementation:

- $\sigma^2_{\mathsf{MSBSK}}$ is the noise variance introduced by bootstrapping key conversion
- $\sigma^2_{\mathsf{SGLWE_{LUT}}}$ is the variance of the error resulting from modulus switching applied to the GLWE ciphertext coefficients of the LUT
- $\sigma^2_{\mathsf{MSGLWE}}$ is the variance of the error introduced from modulus switching applied to the output GLWE ciphertext of Blind Rotate

Let p be the plaintext modulus, σ_{PBS} be the standard deviation of noise introduced by bootstrapping, and π the number of padding bits used. The maximum noise allowed for correct decryption is $(2 \cdot 2^{\pi} \cdot p)^{-1}$. Assuming that the bootstrapping error behaves as a continuous Gaussian then the probability of incorrect bootstrapping is given by:

$$p_{\mathsf{fail}} = \mathsf{erfc}((2^{\pi+3/2} \cdot p \cdot \sigma_{\mathsf{PBS}})^{-1}).$$

Outlook: Optical Acceleration of FHE

We are developing an integrated accelerator for FHE that combines ASIC and photonic technologies. The optical compute elements under development are designed to remove the dominant memory-access and polynomial▫arithmetic bottlenecks in FHE workflows. While the noise analysis presented here remains valid for the optical implementation, the substantial gains in speed and power efficiency enabled by photonics will broaden the space of practical parameters and unlock new concrete applications, specifically in the Machine Learning space.