Blurring Faces for Anonymization on Intel SGX

Giovanni Luiz Zanetti[1†], Bogdan Tomoyuki Nassu, Marcelo de Oliveira Rosa, Keiko Verônica Ono Fonseca
Federal University of Technology, Curitiba, Paraná, Brazil
giovanni-zanetti@hotmail.com, bogdan, mrosa, keiko@utfpr.edu.br

Introduction A large amount of video data (around 25TB/day) is collected from cameras at bus stations by the public transportation system management from the city of Curitiba. The identity of citizens appearing in these images must be preserved although the video images should be used for public purposes like traffic management.

Our Aim Detect, track and blur faces in videos in order to make them available for public uses. Since there are privacy concerns, these procedures must be performed in a secure environment. Given the amount of data to be processed, cloud computing is required to use as a service, avoiding costs required to acquire and maintain computing servers for this task. In order to combine these two requirements (privacy and cloud computing), our proposed solution was implemented through a simple parallelizable image processing pipeline using the Intel Software Guard Extensions (SGX) [2].

The Processing Pipeline Our initial attempts consisted of porting an existing solution (based on the OpenCV library) to run using protected memory schemes (SGX enclaves). This task proved too difficult, due to a complex chain of compilation dependencies and use system-level calls by OpenCV. Therefore, we developed a proof-of-concept implementation using pure C++. The following operations are performed inside SGX enclaves:

- Face detection using a small convolutional network. Its inputs are image sub-regions sampled from an input frame.
- Motion detection using simple frame differencing. Results at static regions can be repeated from the previous frame.
- Tracking, to associate faces detected in previous frames with the current results. By keeping track of each face in the image, we can predict image positions where faces are likely to appear.
- Result buffering: results are presented with a short delay of around 0.67 seconds. Along with face tracking, this allow us to handle false positives and negatives, overriding the output from the face detector when needed.
- Face blurring, using a classical box (mean) filter.

Parallel processing The processing pipeline requires state keeping between frames, but larger video chunks can be processed in parallel. To parallelize their processing, the video files were decoded from MPEG format to raw format in order to encrypt their frames using Rijndael AES-GCM encryption with 128bit key [3]: this cipher method is implemented directly on Intel processor. All frames were arranged in blocks and sent to SGX-based cloud system in order to be processed by parallel sets of processing pipelines before been decrypted (these tasks were carried out inside SGX enclaves). The resulting blurred frames were reorganized in order to be encoded into MPEG video format. Although the encoding process is carried outside SGX enclaves, the sensible data (people's face) were already removed. Finally, ciphering keys are transferred to SGX enclaves running in unsecured cloud computers after remote attestation procedures.

Results Our implementation was tested on a machine with a 3.2GHz Intel I7-8700 processor, with 8GB DDR4 RAM and Ubuntu 16.04. Detection performance was close to that offered by built-in face detector from the OpenCV library. The times for processing one frame (in ms) were as follows:

Environment	Normal (CPU)	SGX
Mean	2400	2627
Standard deviation*	4222	4338
Minimum	255	356
Maximum	12973	13554

*The large variance is explained by the fact that faces are detected in the entire frame at least once per second, while in other occasions motion detection is used to considerably reduce the search area.

Acknowledgment This project has been receiving funds granted from the 3rd EU-BR Coordinated Call (Brazilian Ministry of Science, Technology and Innovation, MCTIC/RNP, BR grant agreements 2550, 2549, 2553, 2552 and 2568) and European Unions Horizon2020 research and innovation programme - EU grant agreement 690111). The project is also supported by the Swiss State Secretariat for Education, Research and Innovation (SERI).

REFERENCES

[1] C. Stergiou and D. Siganos. Neural Networks. https://www.doc.ic.ac.uk/ nd/surprise_96/journal/vol4/cs11/report.html. Accessed: 2019-01-29

[2] Intel. Take Control of Protecting Your Data https://www.intel.com.br/content/www/br/pt/architecture-and-technology/software-guard-extensions.html. Accessed: 2019-01-29

[3] Intel. Rijndael128GCM Encryption https://software.intel.com/en-us/sgx-sdk-dev-reference-sgx-rijndael128gcm-encrypt. Accessed: 2019-01-29

[1] Undergraduate student †Presenter