Time Is Money, Friend! Timing Side-channel Attack against Garbled Circuit Constructions

Mohammad Hashemi[1], Domenic Forte[2], and Fatemeh Ganji[1]

[1] Worcester Polytechnic Institute, Worcester, MA 01609, USA
[2] University of Florida, Gainesville FL 32611, USA
dforte@ece.ufl.edu {mhashemi,fgangi}@wpi.edu

1 Abstract

Secure function evaluation (SFE) enables parties to process confidential data collaboratively without revealing any information about parties' input. Among SFE's arsenal, garbled circuits (GCs), introduced by Yao [7], have become pivotal. They ensure two parties can securely compute functions without compromising privacy and have been widely used in areas such as secure multi-party computation, quantum circuits [2], and advanced fields like genomics and AI [6].

Robust proofs have historically backed the security and privacy of these protocols.Yet, despite rigorous enhancements to increase their practicality [8], it must be ensured that these refinements do not compromise the security and privacy of SFE protocols.

Concurrently, side-channel attacks (SCAs) have proven their mettle in extracting secrets from cryptographic implementations [1,3]. The intersection of these SCAs with GCs remains under-explored. A case in point: Levi et al. [5] have launched a power side-channel attack on garbling schemes to recover a global secret on which all garbling/evaluation processes depend. Leveraging the non-secret input of the garbler, such as the circuit description, the attack can extract global secrets with high success rates from circuits with only a few non-linear gates. This raises the question: *Can one expose parties' inputs by analyzing the timing information leaking during a GC operation?* We answer affirmatively, especially for optimization techniques, such as free-XOR and half-gate [4,8].

In this regard, our contributions are:

1. We introduce *Goblin*, the first non-profiling, single-trace timing SCA that successfully extracts the garbler's input, which, by definition, should have been kept secret. To better demonstrate the power of our attack, we compare it with the recent attack in [5]. The power SCA in [5] has successfully extracted the global secret used in free-XOR optimization, whereas Goblin focuses entirely on recovering the garbler's input. Needless to say, even with the help of the disclosed secret, the garbler's input could not be fully recovered following Levi's attack. Moreover, in contrast to [5], Goblin's effectiveness is limited to neither circuits with a minimum number of input gates nor gate types (XOR or AND).

2. Goblin is machine-learning assisted in disclosing the garbler's input, regardless of its size. For this purpose, k-means clustering is applied, where no manual tuning or heuristic leakage models are needed. It is, of course, advantageous to the attacker and allows for scalable and efficient attacks.
3. Lastly, we highlight the vulnerabilities of various garbling tools to timing SCAs, advocating for further exploration and design enhancement in the GC domain.

2 Acknowledgement

This work has been supported partially by Semiconductor Research Corporation (SRC) under Task IDs 2991.001 and 2992.001 and NSF under award number 2138420.

3 Presenter's biography.

Mohammad Hashemi received his B.Sc. degree in electrical engineering from Islamic Azad University, Tehran, Iran, in 2017, and his M.Sc. degree from the University of Tehran, Tehran, in 2021. He is currently pursuing his Ph.D. degree in electrical engineering at the Vernam Applied Cryptography and Secure Embedded Systems Laboratory, Worcester Polytechnic Institute (WPI), Worcester, MA, USA. His current research interests include hardware security, field programmable gate array (FPGA)-based accelerators, and machine learning.

References

1. Dhem, J.F., Koeune, F., Leroux, P.A., Mestré, P., Quisquater, J.J., Willems, J.L.: A practical implementation of the timing attack. In: International Conference on Smart Card Research and Advanced Applications. pp. 167–182. Springer (1998)
2. Garg, S., Srinivasan, A.: Garbled protocols and two-round mpc from bilinear maps. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS). pp. 588–599. IEEE (2017)
3. Kocher, P., Jaffe, J., Jun, B.: Differential power analysis. In: Annual international cryptology conference. pp. 388–397. Springer (1999)
4. Kolesnikov, V., Schneider, T.: Improved garbled circuit: Free xor gates and applications. In: Intrl. Colloquium on Automata, Languages, and Programming. pp. 486–498. Springer (2008)
5. Levi, I., Hazay, C.: Garbled-circuits from an sca perspective: Free xor can be quite expensive... Cryptology ePrint Archive (2022)
6. Mohassel, P., Zhang, Y.: Secureml: A system for scalable privacy-preserving machine learning. In: 2017 IEEE symposium on security and privacy (SP). pp. 19–38. IEEE (2017)
7. Yao, A.C.C.: How to generate and exchange secrets. In: 27th Annual Symp. on Foundations of Computer Science (sfcs 1986). pp. 162–167. IEEE (1986)
8. Zahur, S., Rosulek, M., Evans, D.: Two halves make a whole. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 220–250. Springer (2015)