

# STRESS-SGX: Load and Stress your Enclaves for Fun and Profit

Sébastien Vaucher, Valerio Schiavoni, and Pascal Felber

University of Neuchâtel, Neuchâtel, Switzerland  
{sebastien.vaucher, valerio.schiavoni, pascal.felber}@unine.ch

**Abstract.** The latest generation of Intel processors supports Software Guard Extensions (SGX), a set of instructions that implements a Trusted Execution Environment (TEE) right inside the CPU, by means of so-called enclaves. This paper presents STRESS-SGX, an easy-to-use stress-test tool to evaluate the performance of SGX-enabled nodes. We build on top of the popular STRESS-NG tool, while only keeping the workload injectors (*stressors*) that are meaningful in the SGX context. We report on several insights and lessons learned about porting legacy code to run inside an SGX enclave, as well as the limitations introduced by this process. Finally, we use STRESS-SGX to conduct a study comparing the performance of different SGX-enabled machines.

**Keywords:** Intel SGX · Load · Stress · Benchmark

## 1 Introduction

The latest generation of Intel processors (starting from the Skylake microarchitecture) features a new set of instructions: Software Guard Extensions (SGX). This instruction set allows programs to execute securely inside hardware *enclaves*, hence creating a Trusted Execution Environment (TEE). Specifically, SGX enclaves protect the code from external threats, including privileged system software. Given the novelty of the technology (the first compatible CPU was released in August 2015) and the lack of in-depth literature, it is still challenging to validate the performance of code running inside enclaves. It becomes even harder to study the problem under varying conditions, such as different hardware revisions or workloads. Moreover, as we show in this paper, microcode updates issued by CPU vendors can introduce performance degradations that are difficult to promptly detect.

The main contribution of this paper is STRESS-SGX, a stress tool capable of artificially putting SGX enclaves under high load. It supports workloads of different nature, as explained in Section 2. As far as we know, STRESS-SGX is the first tool that can be leveraged to induce hybrid workloads (with and without SGX). We believe that many researchers can benefit from using our tool, simplifying design plans for their evaluation settings. STRESS-SGX is GPL-licensed free software<sup>1</sup>.

<sup>1</sup> <https://github.com/sebva/stress-sgx>

**Motivating scenarios.** We introduce three use-cases for STRESS-SGX and how researchers could leverage this tool. First, there exist several real-world cluster traces readily available for research, such as the ones released by Google Borg [7] or Microsoft Azure [2]. However, these traces do not have the same properties as typical SGX workloads. Due to anonymization, hardware characteristics of the original machines remain undisclosed. Finally, the nature of jobs deployed on these clusters is also confidential. STRESS-SGX allows to directly map behaviors described in such traces onto similar, SGX-specific ones.

A second motivating scenario stems from the availability of a range of SGX-enabled CPUs on the market, each with their own characteristics and associated performance. Apart from the expected differences due to hardware revision, cache size, available instruction sets or frequency, one still needs to hand-craft micro-benchmarks in order to evaluate SGX-specific performance. STRESS-SGX facilitates this process by providing the same interface as STRESS-NG [4] to inject load on a CPU under controlled conditions.

A third motivating scenario directly derives from the necessity to measure the electric energy consumption of code executing inside enclaves, by means of software [1] or hardware power meters. STRESS-SGX offers a common code-base for both SGX and native contexts, making it easier to isolate the energy requirements of SGX-specific workloads.

**Roadmap.** The remainder of the paper is organized as follows: Section 2 describes some implementation details and lessons learned. Section 3 presents our preliminary evaluation. Finally, Section 4 discusses future work and concludes.

## 2 Implementation

STRESS-SGX is implemented as a fork of STRESS-NG (version 0.09.10). It directly reuses the same compilation and runtime foundations. From a user perspective, SGX-enabled stressors (specific pieces of executable code that exercise a given CPU functionality or low-level operation) are selected by specifying command-line options starting with `--sgx`. We concentrate on porting CPU stressing methods to run in an enclave in a way that makes both native and SGX versions comparable from a performance standpoint. All features offered by the CPU stressors shipped with STRESS-NG are supported by STRESS-SGX, with the exception of partial load. Specifying a fixed load percentage is not possible (*i.e.*, it is locked to 100%), because this feature is based on precise timing, which is currently not available inside the enclave. Our STRESS-SGX prototype supports 54 enclave-enabled stress methods out of the 68 currently shipped with STRESS-NG. The exhaustive list is presented in Table 1. We note that support for each stressor depends on the availability of its required functionalities within the enclave, as well as their support by the Software Development Kit (SDK).

**Porting CPU stressors to run inside an enclave.** The Intel SGX SDK is designed in a way that allows existing code to be ported to run in an enclave

**Table 1.** List of stressors supported by STRESS-SGX. “X” indicates a stressor available in STRESS-NG that is not SGX-compatible. f=float, d=double, ld=longdouble.

ackermann	✓
bitops	✓
calltime	X
complex[f,d,ld]	✓
correlate	✓
crc16	X
decimal[32,64,128]	✓
dither	✓
djb2a	✓
double	✓
euler	✓
explog	✓
factorial	✓
fft	✓
fibonacci	✓
float	✓
fnv1a	✓
gamma	✓
gcd	✓
gray	✓
hamming	✓
hanoi	✓
hyperbolic	✓
idct	✓
int8	✓
int16	✓
int32[f,d,ld]	✓
int64[f,d,ld]	✓
int128[f,d,ld]	X
jenkin	✓
jmp	✓
ln2	✓
longdouble	✓
loop	✓
matrixprod	X
nsqrt	✓
odds	✓
omega	✓
parity	✓
pi	✓
pi	✓
piw	✓
prime	✓
psi	✓
queens	✓
rand	✓
rand48	X
rgb	✓
sdbm	✓
sieve	X
stats	✓
sqrt	✓
trig	✓
union	✓
zeta	✓

with reasonable engineering efforts [3]. We take advantage of this fact to port the CPU stressors of STRESS-NG in our own STRESS-SGX fork. We detail here the particular complications that we encountered throughout this porting effort, and the specific adjustments applied to solve each of them.

After creating an enclave using the template given by the SDK, we copy the relevant source code for the CPU stressors in the aforementioned template. STRESS-NG defines several macros in its global header file. It is not possible to include the code *verbatim* as it depends on numerous system-specific features. Our solution is to define the needed symbols on a case-by-case basis.

The next obstacle is the need to split the code in *trusted* and *untrusted* parts. We decompose the code in a way that limits the number of enclave transitions (*i.e.*, entering and exiting the enclave) required to run a stressor. Transitions are costly [6], so it is crucial that none happen while stressing is in progress, to ensure a consistent behavior.

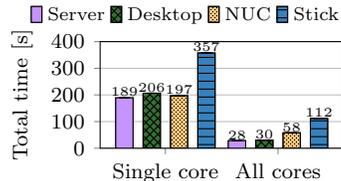
The user can gracefully abort the execution of STRESS-SGX using standard Linux signals such as SIGINT. STRESS-NG includes a mechanism to catch the majority of signals and react accordingly for its built-in stressors. We leverage the fact that SGX enclaves can access the *untrusted* memory of their enclosing process to pass a pointer to the `g_keep_stressing_flag` variable to the enclave. This variable is later used to indicate when to stop the execution of stressors. Code running inside the enclave periodically polls the flag, and stops the execution if asked by the user. The same flag is also used to make a stressor run for a given duration. Timekeeping is done outside the enclave, with the indication to stop the execution notified by changing the value of the variable.

**Ensuring byte-per-byte equivalence of native and SGX code.** During the initial testing phase of STRESS-SGX, we observed vast differences in performance for the same stressor executed in native and enclave modes. As a matter of fact, while the source code was identical in both instances, the resulting compiled binaries slightly differ. We believe that these slight differences are inevitable, as different linking rules are needed by each execution mode.

Conveniently, an enclave compiled using the official SGX SDK will produce a statically-compiled shared library. We leverage this aspect to guarantee that both native and SGX versions of a stressor execute a perfectly identical binary by dynamically linking this shared object. Choosing this optional approach limits STRESS-SGX to stressors that are available in enclave mode.

**Table 2.** Hardware characteristics of our test machines. All processors are made by Intel.

Category	Model	Processor	Cores	Freq. [GHz]	
				Base	Max.
Server	Supermicro 5019S-M2	Xeon E3-1275 v6	4	3.8	4.2
Desktop	Dell Optiplex 7040	Core i7-6700	4	3.4	4.0
NUC	Intel NUC7i7BNHX1	Core i7-7567U	2	3.5	4.0
Stick	Intel STK2m3W64CC	Core m3-6Y30	2	0.9	2.2



**Fig. 1.** Time needed to perform 100 000 000 enclave transitions.

### 3 Evaluation

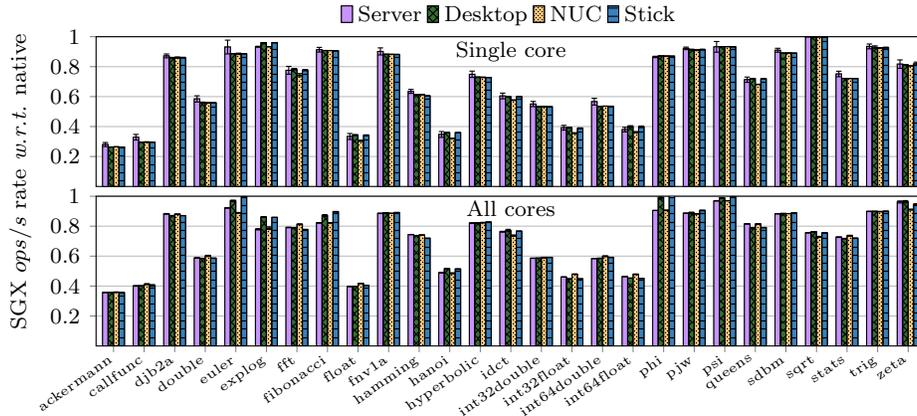
This section presents our preliminary evaluation of STRESS-SGX. Table 2 lists the SGX machines used for our experiments. We chose these machines due to their different form factors, hardware features, and widespread market availability. We expect them to represent a meaningful sample of SGX machines that are in use nowadays. These machines are configured to run Ubuntu Linux 17.10, along with v2.0 of the Intel SGX SDK, Platform Software (PSW) and driver.

**Cost of enclave transitions.** When programming for SGX, it is important to keep in mind that the cost required to enter and exit an enclave is significant [6]. Figure 1 presents the time needed to perform 100 million enclave transitions on a single core at a time (left) and on all available cores (right). As expected, the server-class machine is consistently faster than the other ones, while the Intel Compute Stick performs the worst. We also observe marginal differences in single-core performance, in which case the processor can run at its maximal turbo frequency. Given the notable price difference between the various machines, the Intel NUC offers the best price-performance ratio.

**Cost of latest microcode update.** The second suite of microbenchmarks highlights a rather surprising effect of a recent microcode update, recently issued by Intel (on 2018-01-08) to mitigate the Spectre attack [5]. To observe the performance impact of this microcode update, we execute all 54 supported stressors before and after the microcode update. Using the previous microcode, all stressors display the same performance in SGX and native modes. Under the updated microcode, on the other hand, we observe a significant difference in SGX *versus* native performance. Figure 2 presents the results for the 27 tests for which SGX performance is affected. We measure slowdowns up to  $3.8\times$  (ackermann running on a single core). Given the undisclosed nature of microcode updates, it is difficult to identify the root cause of this performance degradation.

### 4 Conclusion and Future Work

The expanding availability of SGX-enabled machines calls for new tools to evaluate the performance of *secure clusters*. The current lack of ready-to-deploy SGX



**Fig. 2.** SGX performance compared to native on different types of computers, using January 2018 microcode (previous microcode had performance similar to native).

applications usually forces researchers to implement single-use workloads. This paper presented STRESS-SGX, an easy-to-use tool capable of stressing SGX enclaves and report on diverse metrics. We plan to extend our prototype along the following directions: first, we intend to evaluate the performance of reading and writing encrypted memory using the memory stressors of STRESS-NG, ported in STRESS-SGX (not presented in this paper due to lack of space). Second, we will integrate these stressors within a large-scale monitoring and performance framework for containerized microservices, to easily monitor performance regressions.

**Acknowledgement.** The research leading to these results has received funding from the European Union’s Horizon 2020 research and innovation programme under the LEGaTO Project (legato-project.eu), grant agreement No 780681.

## References

1. Bourdon, A., Noureddine, A., Rouvoy, R., Seinturier, L.: PowerAPI: A software library to monitor the energy consumed at the process-level. ERCIM News (92) (2013)
2. Cortez, E., Bonde, A., Muzio, A., Russinovich, M., Fontoura, M., Bianchini, R.: Resource central: Understanding and predicting workloads for improved resource management in large cloud platforms. In: SOSP ’17. pp. 153–167 (2017)
3. Intel Corporation: Intel Software Guard Extensions SDK Developer Reference for Linux OS (Nov 2017), <https://frama.link/sgx20>
4. King, C.I.: Stress-ng, <http://kernel.ubuntu.com/~cking/stress-ng/>
5. Kocher, P., Genkin, D., Gruss, D., Haas, W., Hamburg, M., Lipp, M., et al.: Spectre attacks: Exploiting speculative execution. arXiv preprint arXiv:1801.01203 (2018)
6. Orenbach, M., Lifshits, P., Minkin, M., Silberstein, M.: Eleos: ExitLess OS services for SGX enclaves. In: EuroSys ’17. pp. 238–253 (2017)
7. Wilkes, J.: More Google cluster data. Google research blog (Nov 2011), <http://googleresearch.blogspot.com/2011/11/more-google-cluster-data.html>