Towards Sustainable Software Infrastructures for Data-Intensive Systems

Guillaume Fieni, Romain Rouvoy, Lionel Seinturier

University of Lille / INRIA

Introduction

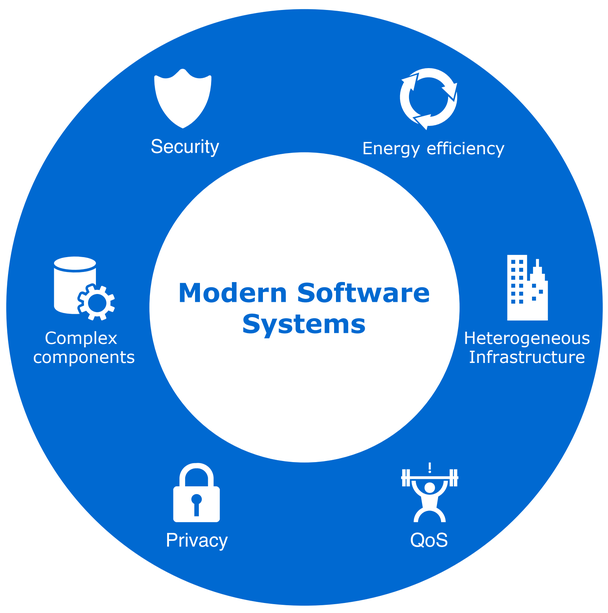

Figure 1:Modern software systems constraints

Context and challenges

Focus on data workflows architecture :

- Consume data produced by a multitude of devices
- Hosted in a cloud environments (shared ressources)
- Requires the execution of secured tasks (e.g., anonymization of attributes)

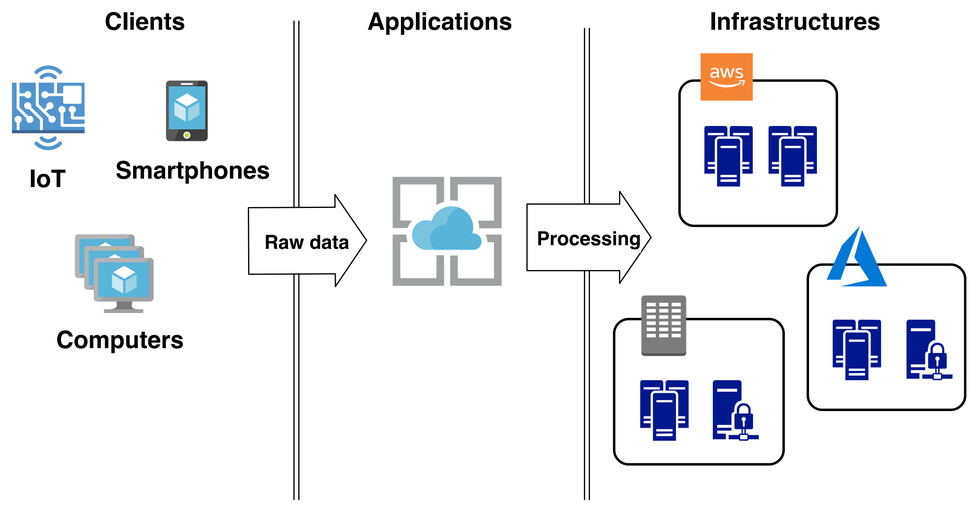

Figure 2:Example of Data-intensive infrastructure

Our research consists in providing solutions to :

- Account the power consumption beyond the scale of nodes
- Maximize the availability of the hosting infrastructure
- Minimize the energy footprint of secured or non-secured tasks

Research agenda

Main objectives of this PhD thesis :

- Exploring the trade-off between security and energy efficiency
- Optimization of the resource usage of data-intensive software infrastructures

Important steps

- Providing accurate and fine-grained power estimations
- Evaluating the energy impact of secure isolation mechanisms
- Optimization of the resources usage (including power consumption)

Preliminary results

SmartWatts, a software defined power meter based on *PowerAPI* to monitor the CPU and RAM power consumption of application containers in a datacenters.

Current approaches problems :

- Use an external power meter as reference for global power consumption
- Requires off-line generation of power models on a reference machine
- No fine-granularity DRAM power estimations

Our approach :

- Provides on-line per-container power estimations for CPU and DRAM
- Generation at run-time of the power model using active-learning methods
- *Hardware Performance Counters* (HPC) as ressource activity baseline
- Intel *Running Average Power Limit* (RAPL) as power usage reference

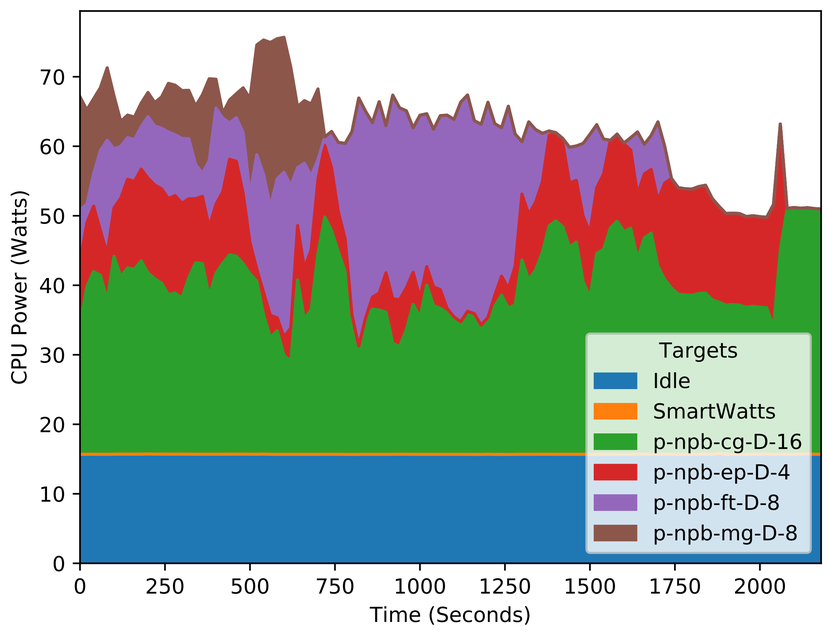

Figure 3:SmartWatts per-container CPU power usage estimations

Future work

- Provides power estimations for ARM CPUs
- Monitoring of Intel *Software Guard Extensions* (SGX) secure enclaves

Related Work

[1] S. Arnautov, B. Trach, F. Gregor, T. Knauth, A. Martin, C. Priebe, J. Lind, D. Muthukumaran, D. O'Keeffe, M. Stillwell, D. Goltzsche, D. M. Eyers, R. Kapitza, P. R. Pietzuch, and C. Fetzer.
SCONE: Secure Linux Containers with Intel SGX.
In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 689–703, Nov. 2016.

[2] M. Colmant, M. Kurpicz, P. Felber, L. Huertas, R. Rouvoy, and A. Sobe.
Process-level Power Estimation in VM-based Systems.
In *Proceedings of the 10th European Conference on Computer Systems (Eurosys)*, 2015.

[3] A. Havet, V. Schiavoni, P. Felber, M. Colmant, R. Rouvoy, and C. Fetzer.
GenPack: A Generational Scheduler for Cloud Data Centers.
In *5th IEEE International Conference on Cloud Engineering (IC2E)*, Apr. 2017.