Regaining Lost Seconds: Efficient Page Preloading for SGX Enclaves

Ximing Liu, Wang Lizhi,
Xiaoli Gong*, Ziyi Zhao
Nankai University

Wenwen Wang
Univeristy of Georgia

Pen-Chung Yew
Univeristy of Minnesota

Intel(R) SGX Comes with Significant Overheads

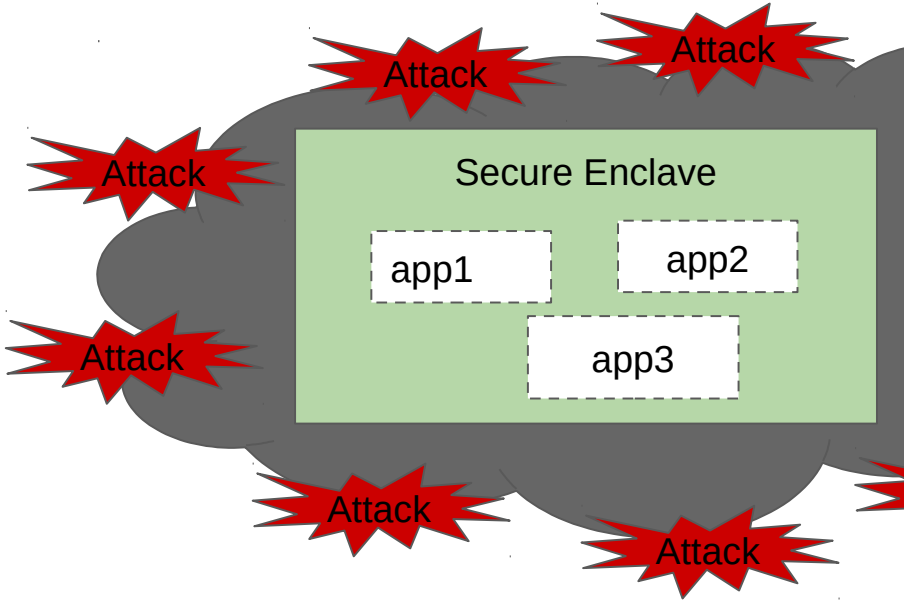

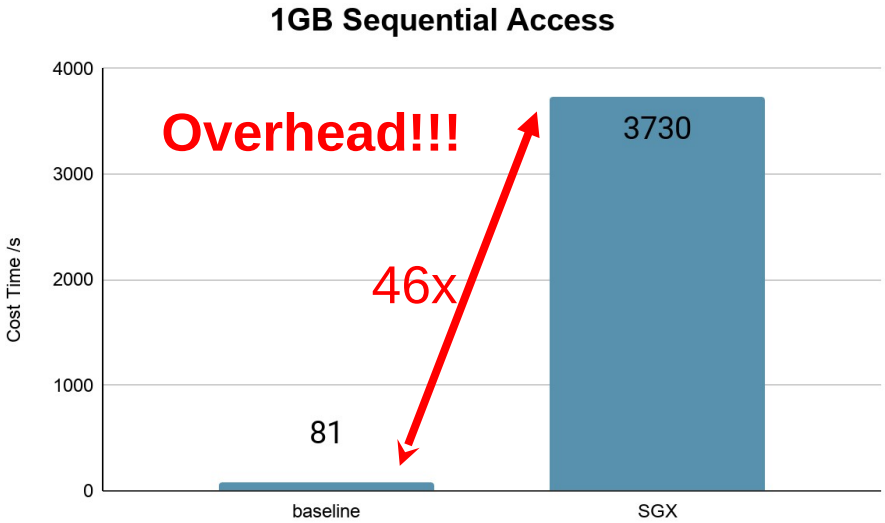

Enclave Page Cache (EPC) Crucial to Performance

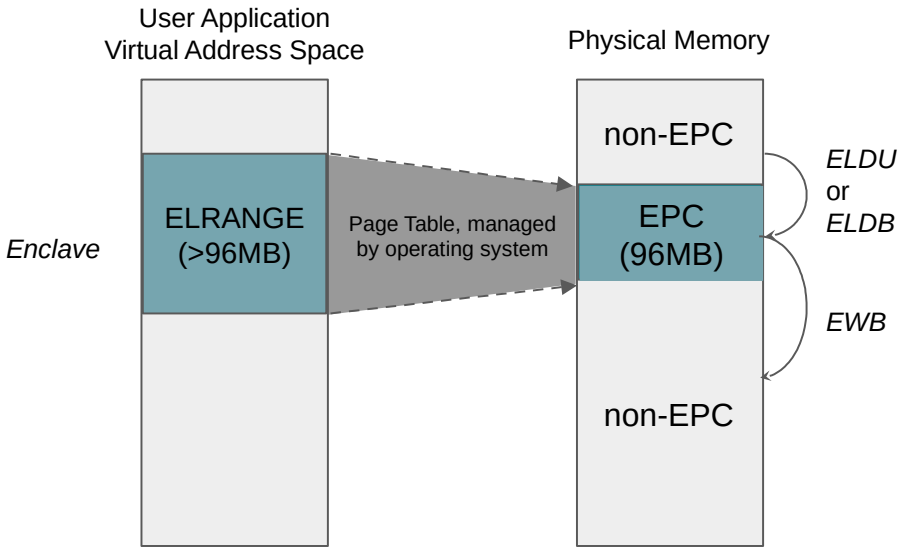

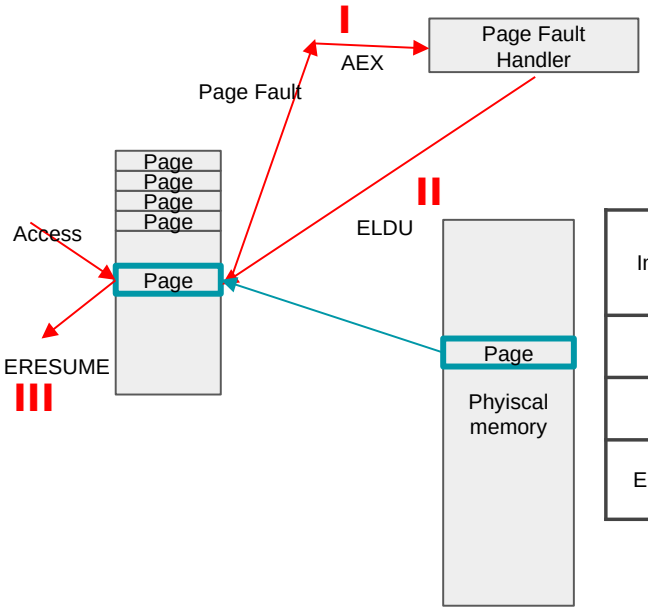

Instruction	Time/ **cycles**
AEX	10000
ELDU	44000
ERESUME	10000

Can We Hide SGX Page Fault Latency Using Data Prefetching?

Taking Advantage of Application Access Patterns

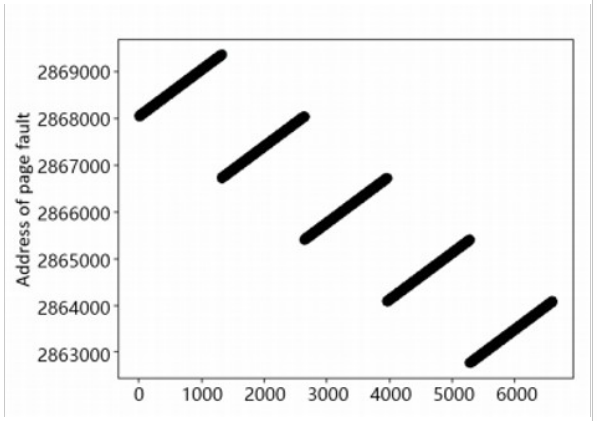

bwaves

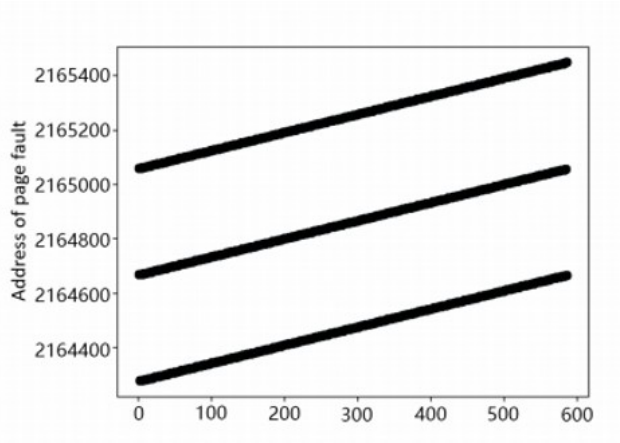

lbm

DFP - Dynamic Fault History-Based Preloading

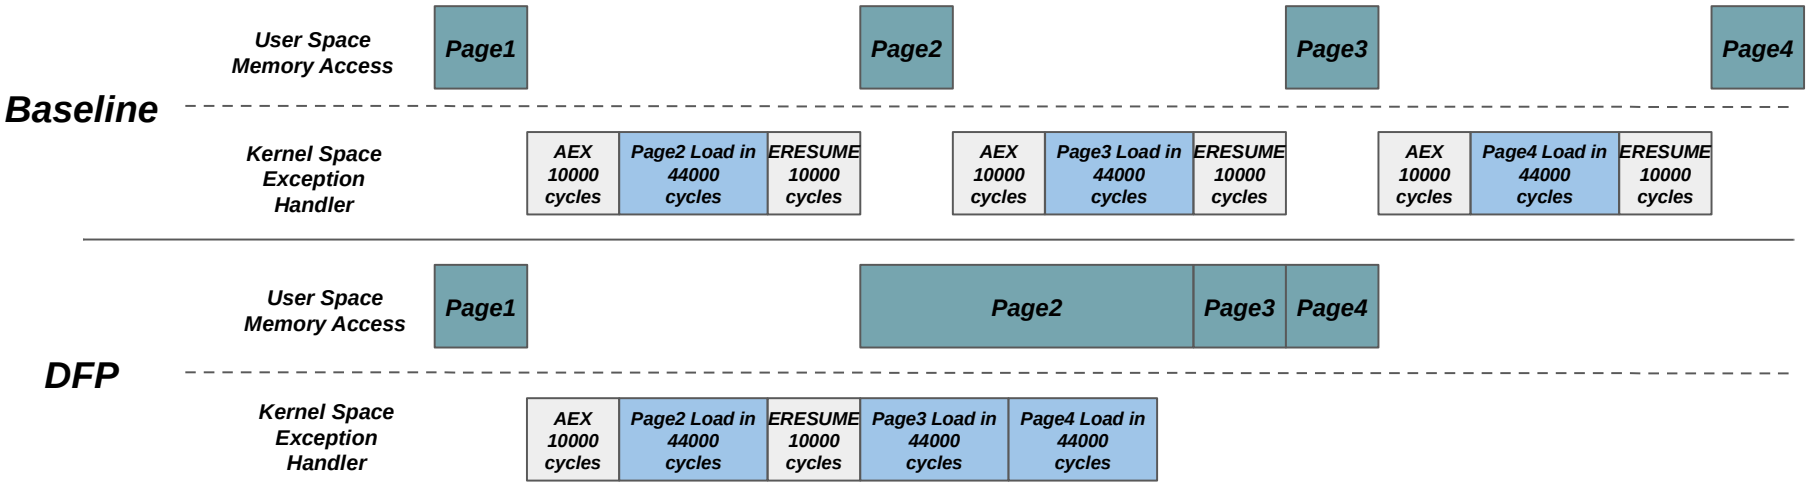

Multiple-Stream Predictor Algorithm in DFP

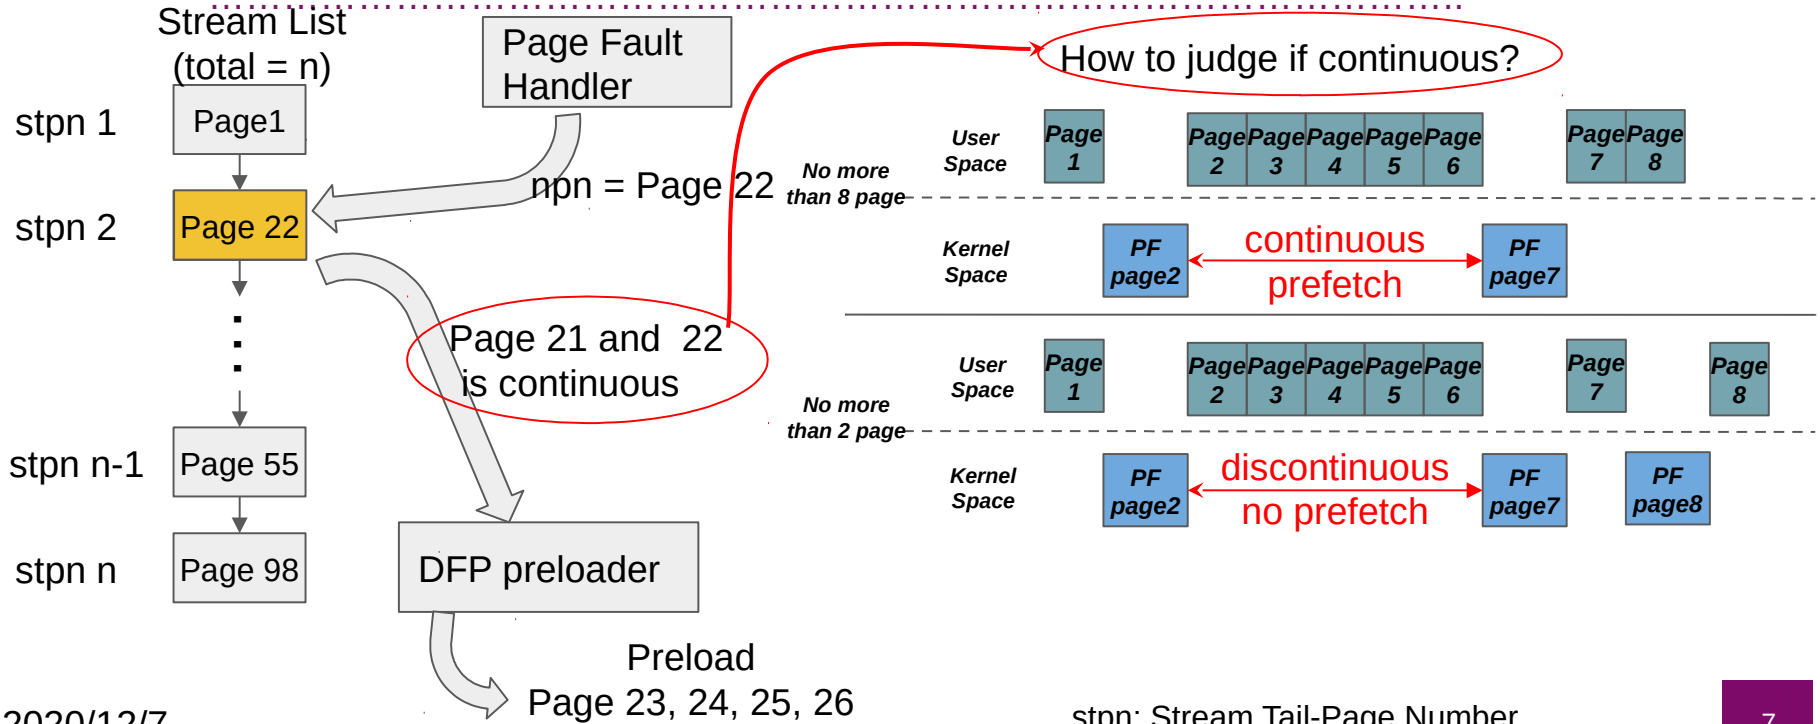

stpn: Stream Tail-Page Number
npn: New Page Number

Apps with Irregular/Unpredictable Access Patterns

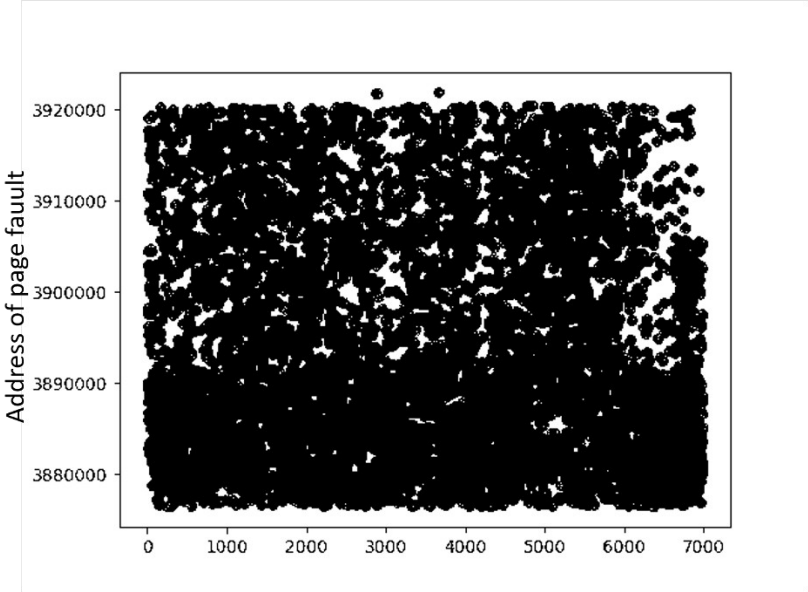

deepsjeng

SIP - Source-Level Instrumentation-Based Preloading

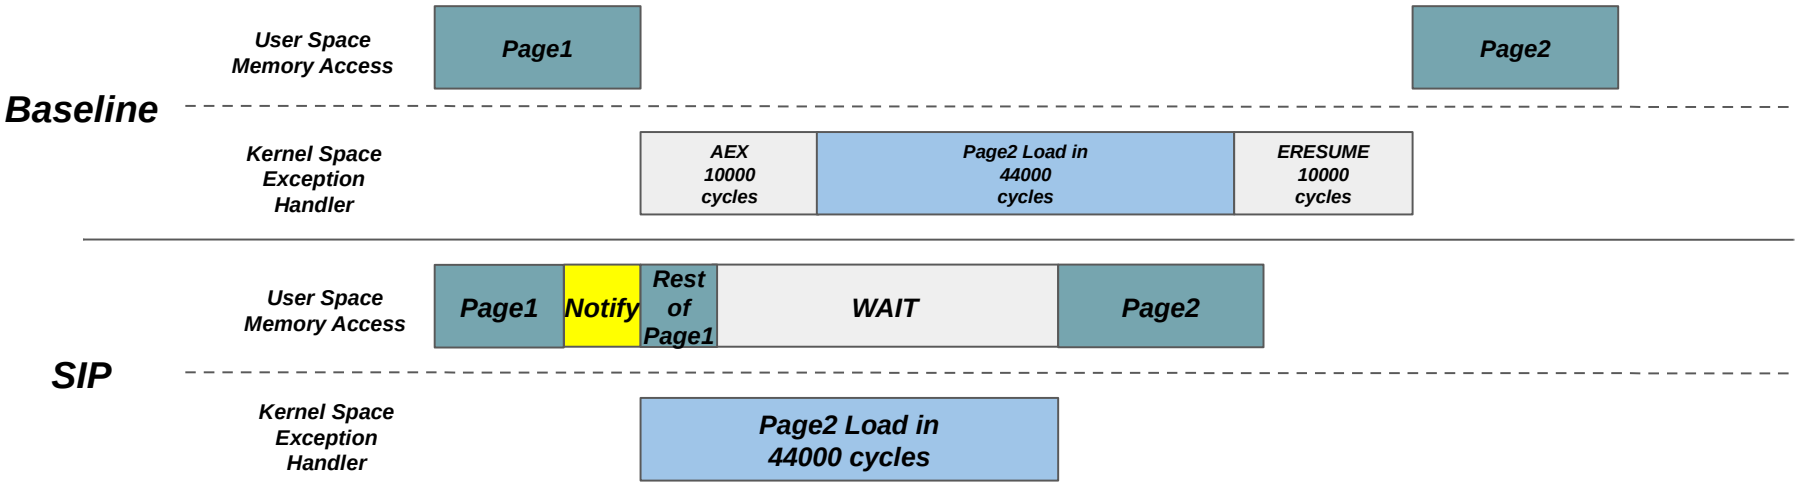

Profile-Guided Program Instrumention

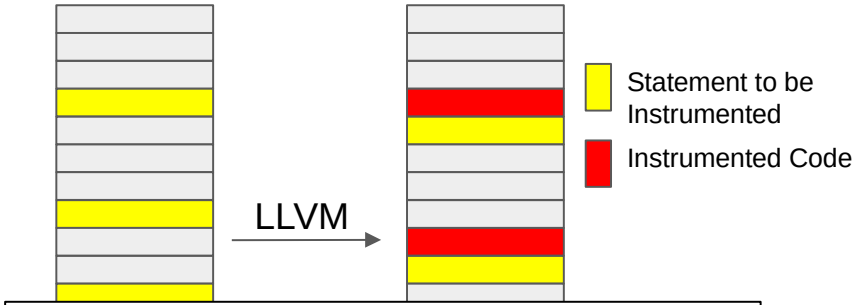

- During analysis phase:
  - Mark all the address accessed
  - Mark all locations of instrumentation

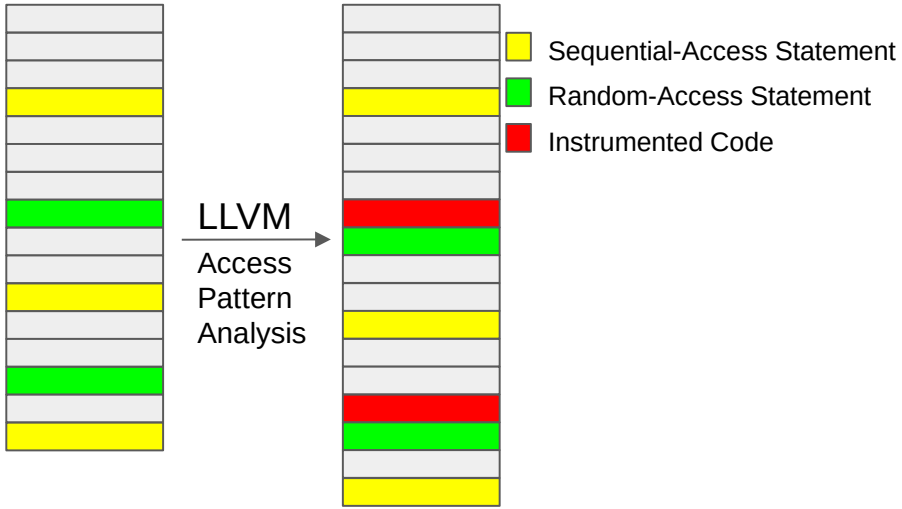

How To Integrate DFP And SIP In An Application?

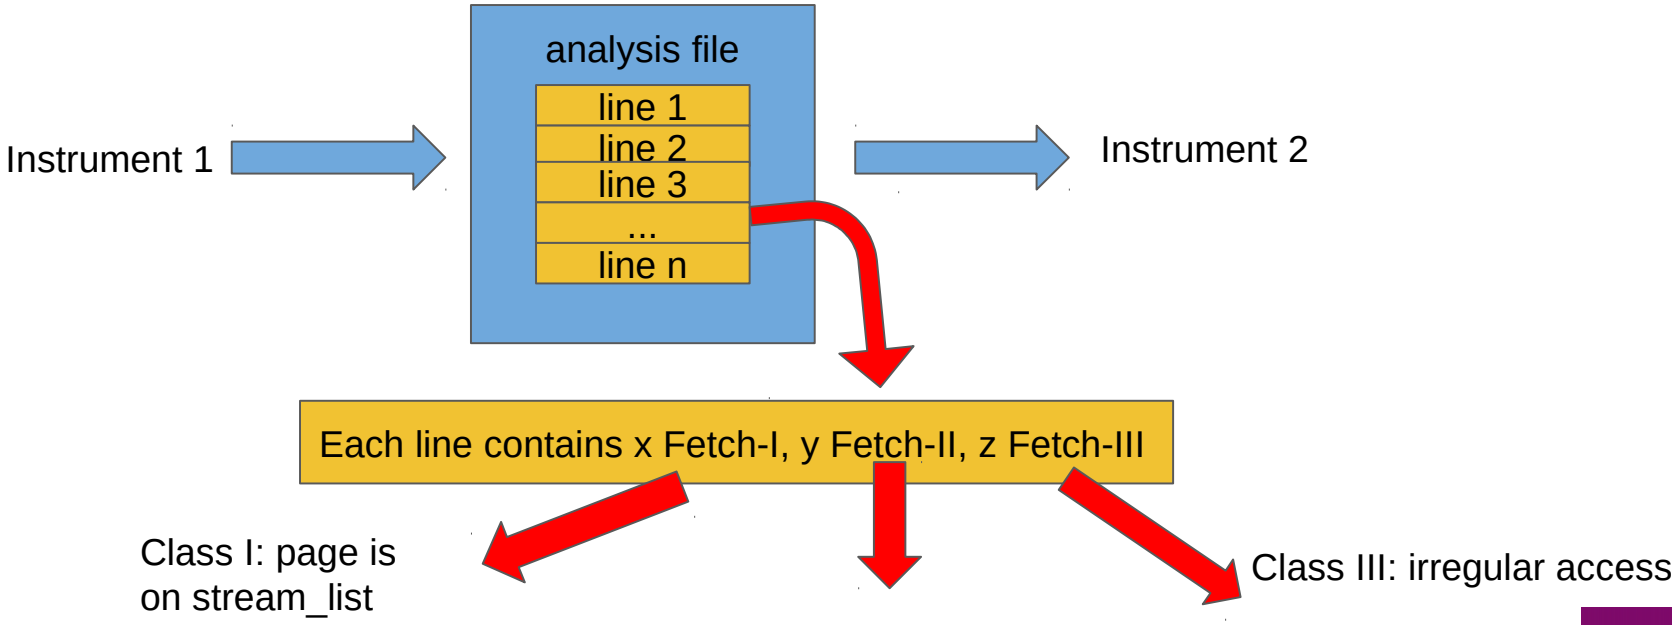

DFP Performance on SPEC2017

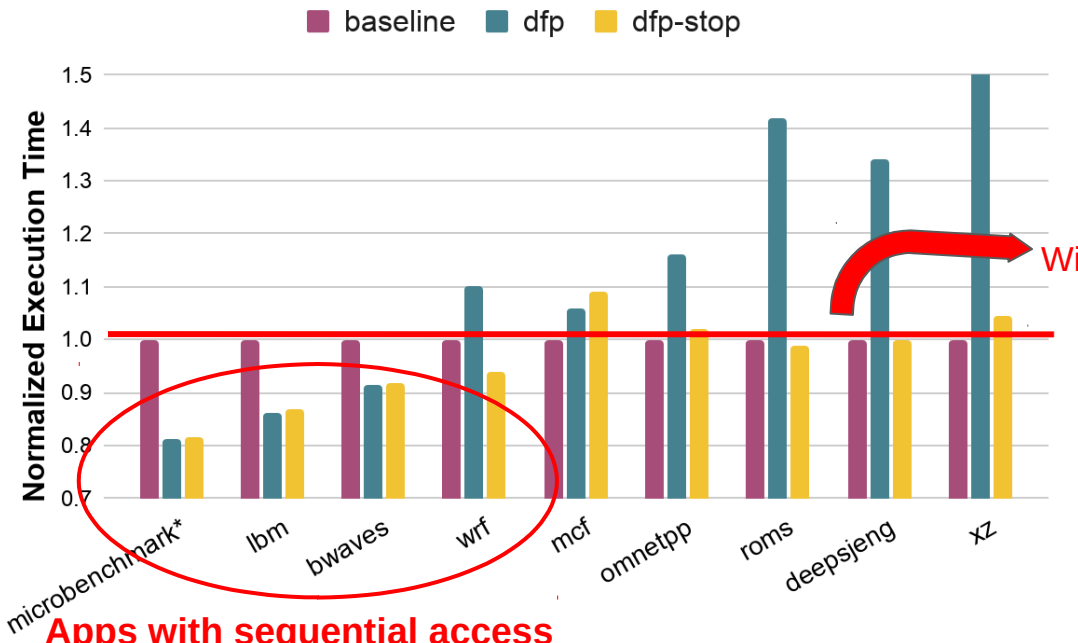

*microbenchmark: 1GB sequential access app

SIP Performance on SPEC2017

Benchmark	Instrumentation Points
mcf.2006	114
mcf	99
xz	46
deepsjeng	35
lbm	0
microbenchmark	0

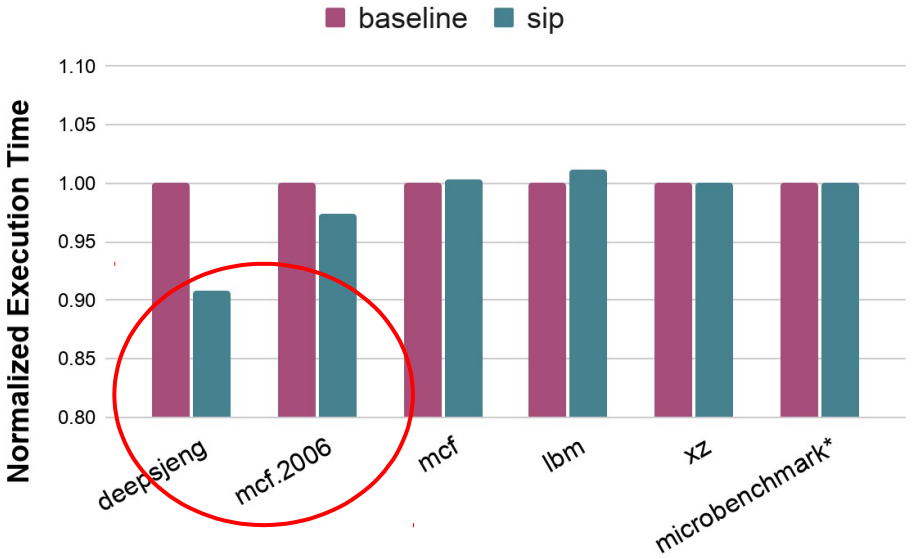

*microbenchmark: 1GB sequential access app

Performance Using SIP + DFP

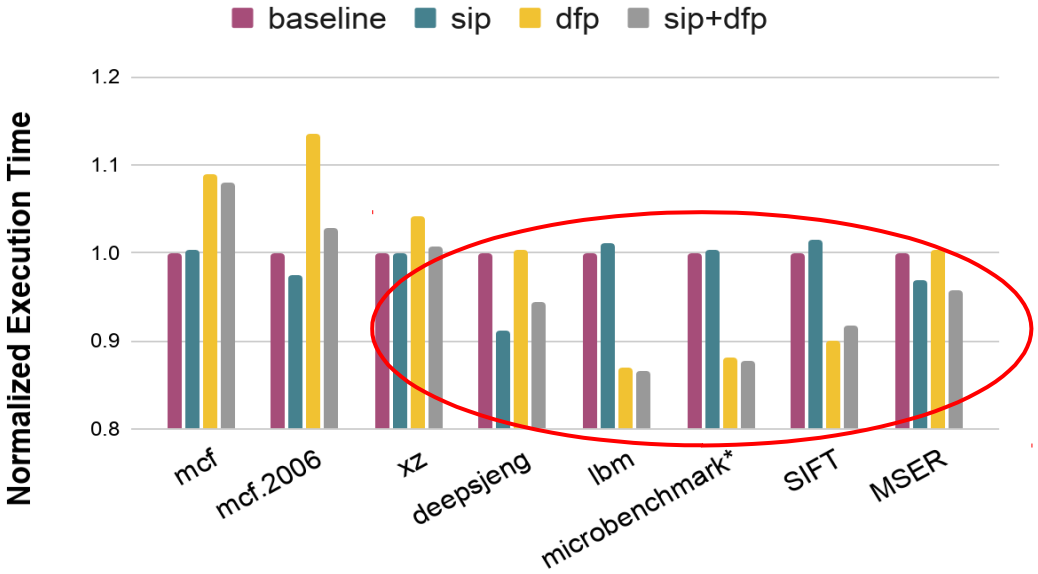

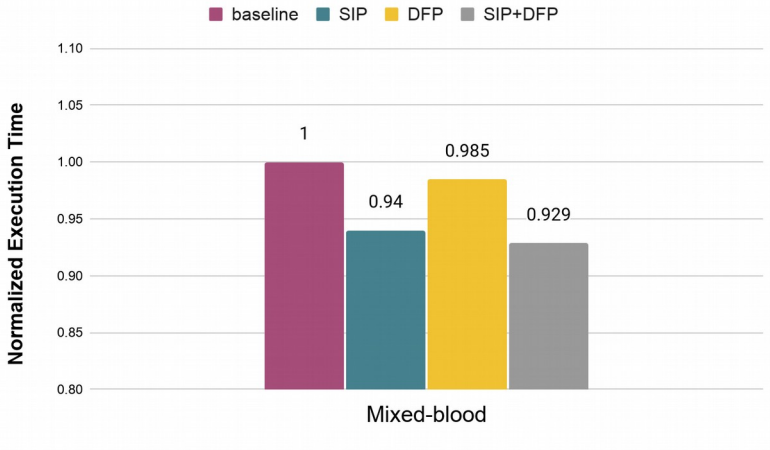

*microbenchmark: 1GB sequential access app

Conclusion

- Intel SGX offers **security** but also **overhead**s caused by **page faults**.
- We propose two page-preloading mechanisms **DFP** and **SIP** to improve **sequential** and **random** memory accesses in applications.
- Evaluation on SPEC2017, some real-world applications and a micro-benchmark program shows these two preloading mechanisms achieve an average of **11.4%** and **7.0%** performance improvement, respectively

Thanks !
Q&A