

# Optimization of GENESIS/QSimulate on Fugaku for High-Performance QM/MM-MD Simulation using Current LLVM/Clang++ Compiler

162s1

Oshingo Ito<sup>1</sup>, Chigusa Kobayashi<sup>1</sup>, Kiyoshi Yagi<sup>2</sup>, Yuji Sugita<sup>1,3,4</sup>

<sup>1</sup> Computational Biophysics Research Team, RIKEN Center for Computational Science

<sup>2</sup> Department of Chemistry, Institute of Pure and Applied Sciences, University of Tsukuba

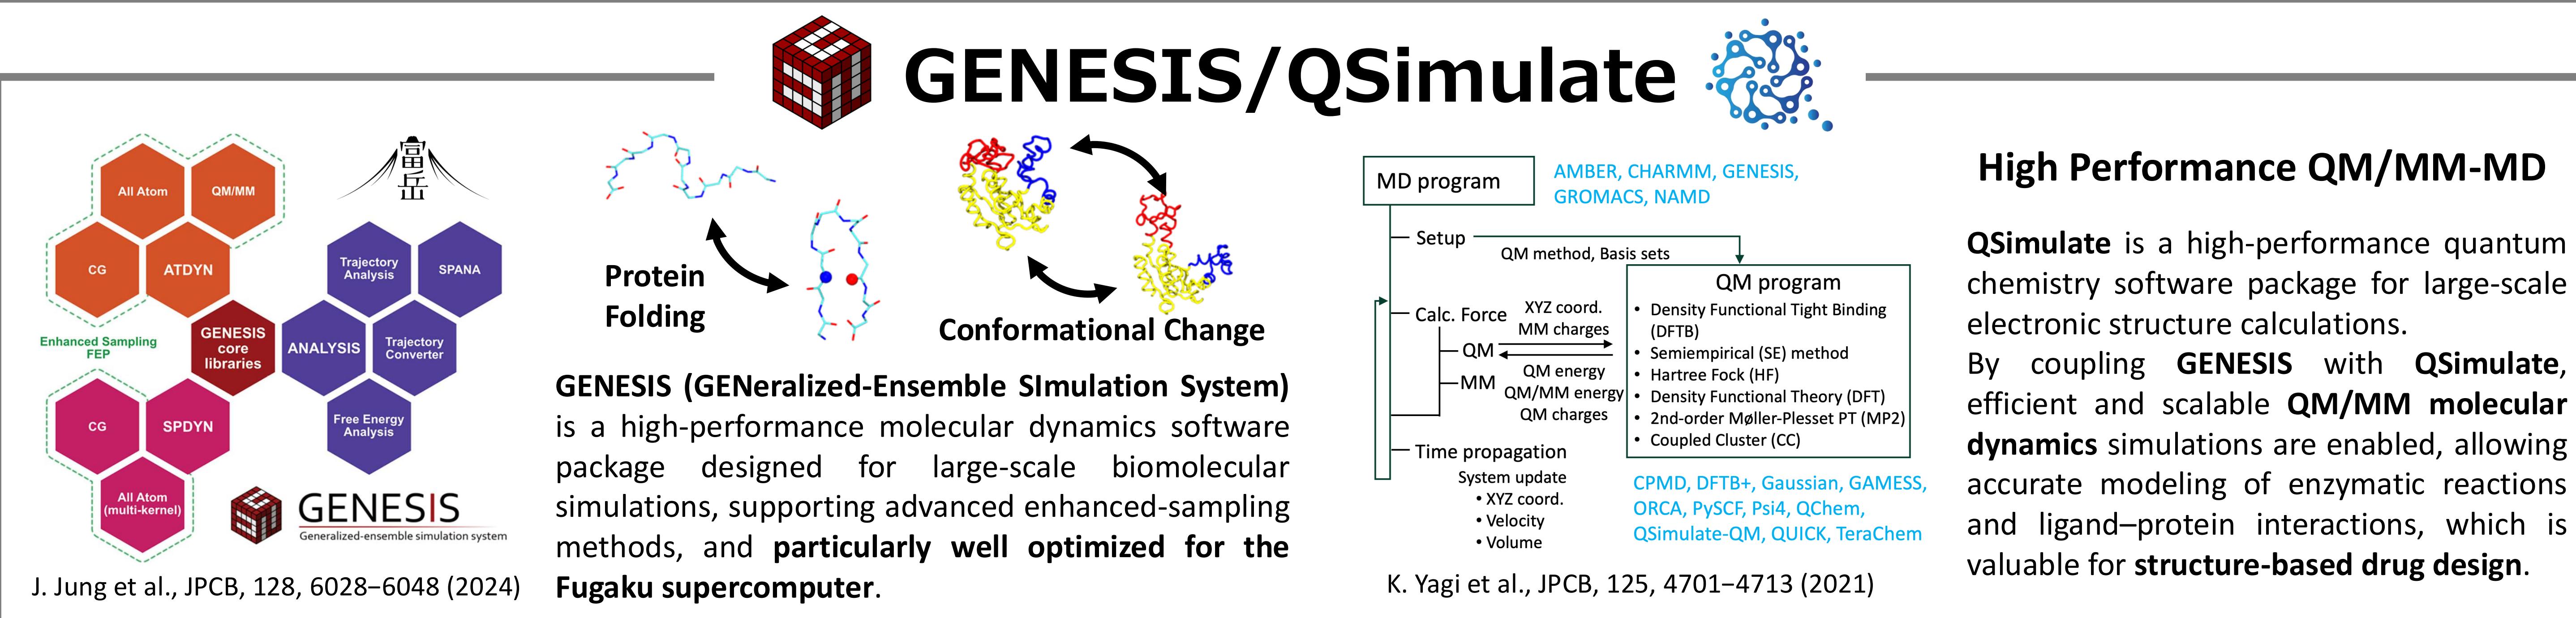
<sup>3</sup> Theoretical Molecular Science Laboratory, RIKEN Pioneering Research Institute

<sup>4</sup> Graduate School of Science, University of Tokyo



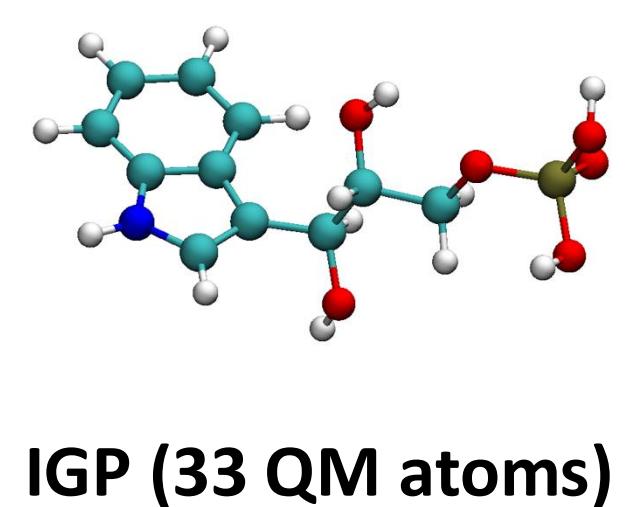
## Background

The GENESIS/QSimulate, combined molecular dynamics (MD) program, GENESIS, and Quantum Mechanics/Molecular Mechanics (QM/MM) program, QSimulate-QM (<https://qsimulate.com/academic>), demonstrates excellent parallel scalability for QM/MM-MD simulations on multi-node supercomputers. High-performance QM/MM-MD simulations using this program have elucidated the mechanisms of biological functions, including enzymatic reactions, on Intel-based systems. However, its performance on the Fugaku, which employs Arm CPUs and the Fujitsu compiler, remains suboptimal due to insufficient code optimization.



## Optimization of GENESIS/QSimulate on Fugaku

### Single-Point Calculations



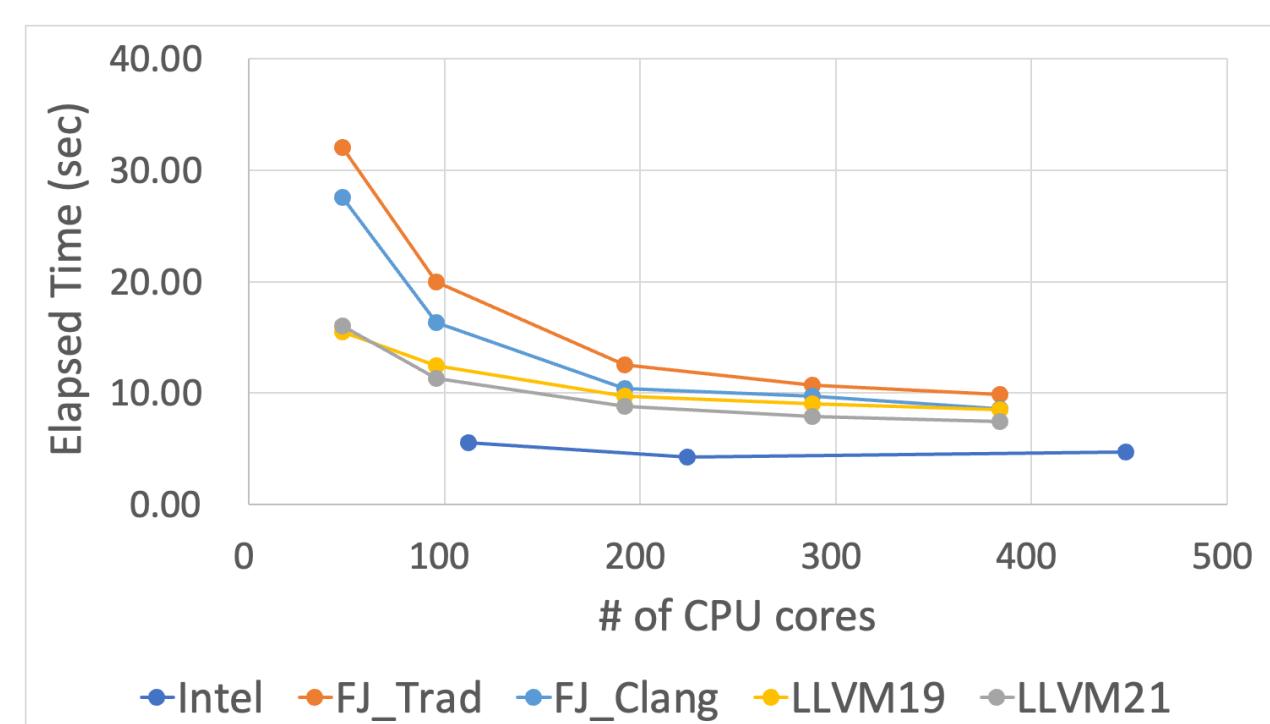
IGP (33 QM atoms)

First, we evaluated the performance of QSimulate on Fugaku (CPU=Intel Xeon Max 9480) and HOKUSAI BigWaterfall2 (so-called "HBW2", CPU=A64FX) using different compilers, **based on single-point calculations for small systems**.

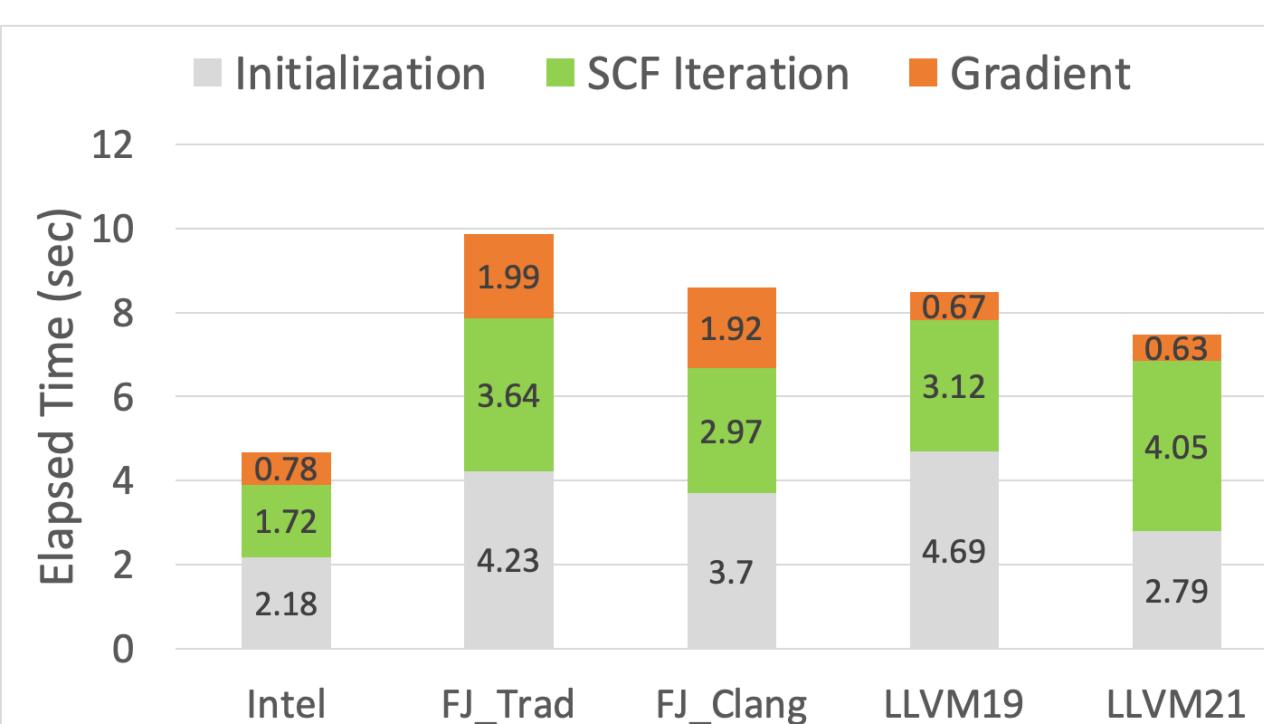
System: IGP

QM Level: B3LYP-D3/cc-pVDZ

Time vs. Cores



Computational Time Breakdown (448 cores (HBW2), 384 cores (Fugaku))

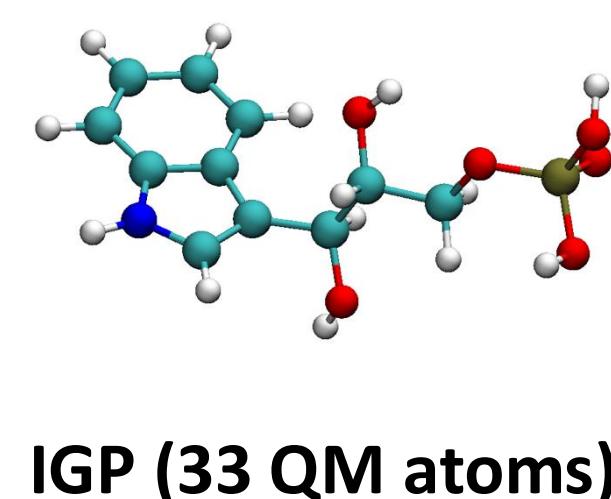


Efficiency Comparison with an Intel-Based Machine

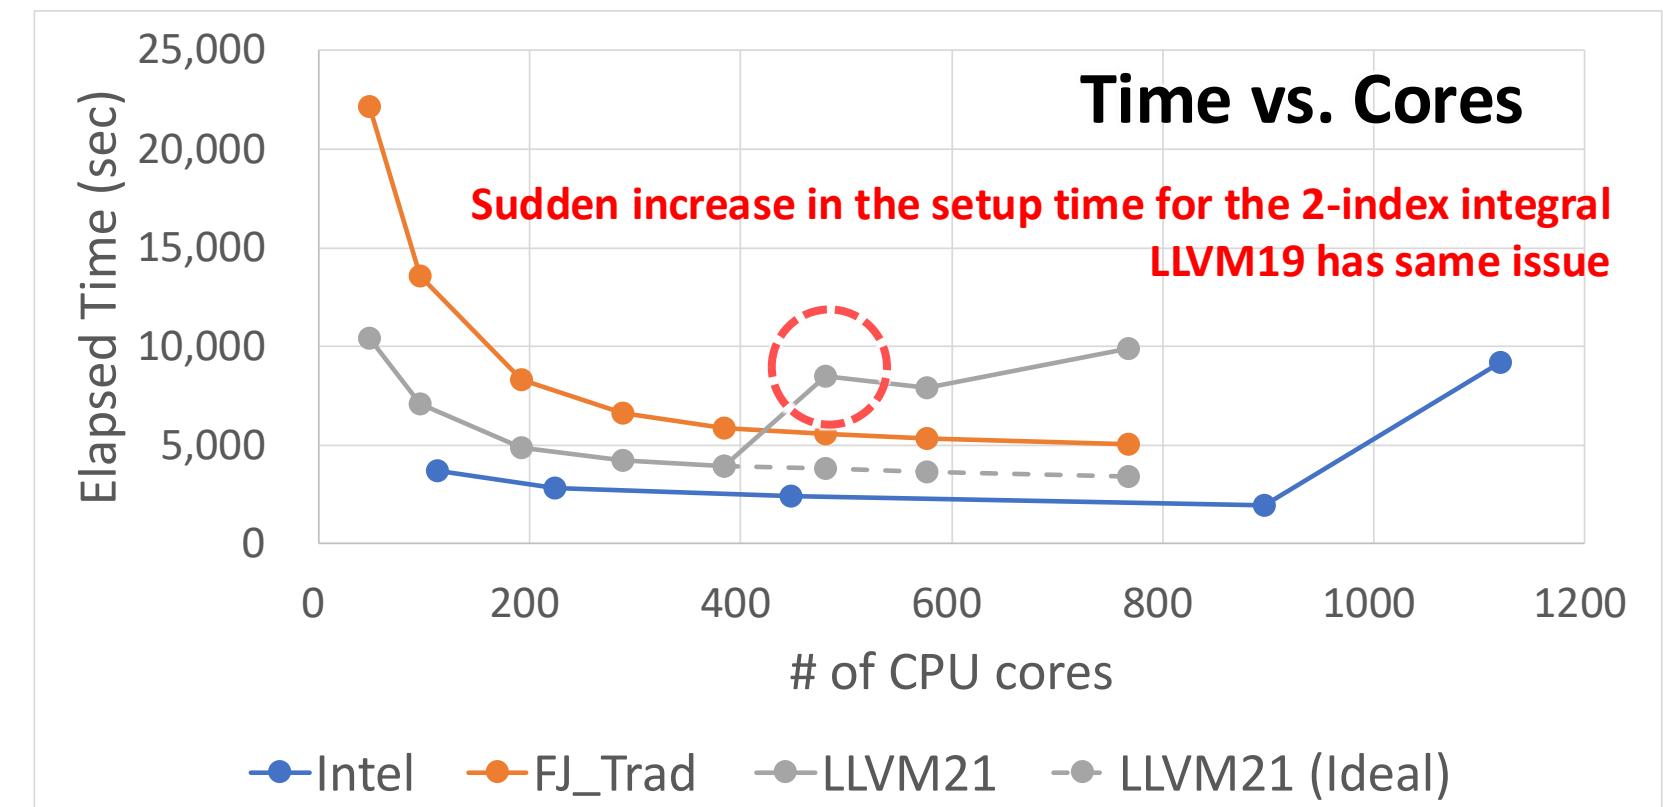
Machine	Compiler	CPU core	Elapsed Time (sec)	Runtime Ratio
HBW2	Intel	448	4.68	x 1.00
Fugaku	FJ (Traditional)	384	9.68	x 2.07
Fugaku	FJ (Clang)	384	8.59	x 1.84
Fugaku	LLVM19	384	8.48	x 1.81
Fugaku	LLVM21	384	7.47	x 1.60

The LLVM compilers (ver. 19 and 21) did not improve the computational time of the **SCF iterations**, which are one of the most time-consuming parts of QM calculations. In contrast, the LLVM compilers significantly reduced the computational time of the **gradient calculations**, another major time-consuming component.

### QM/MM-MD Simulations

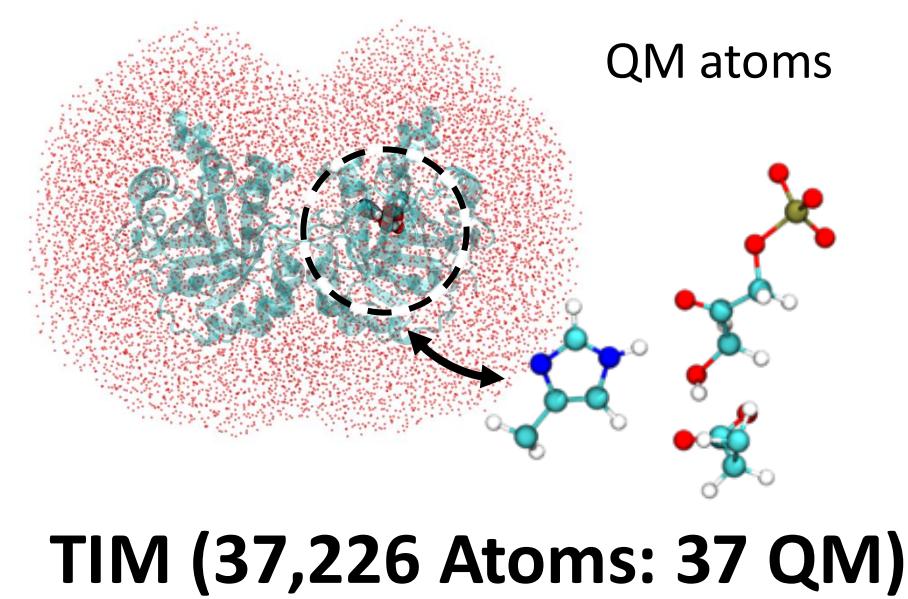


IGP (33 QM atoms)  
QM Level: B3LYP-D3/cc-pVDZ  
NVT, 300K, 1,000 MD steps

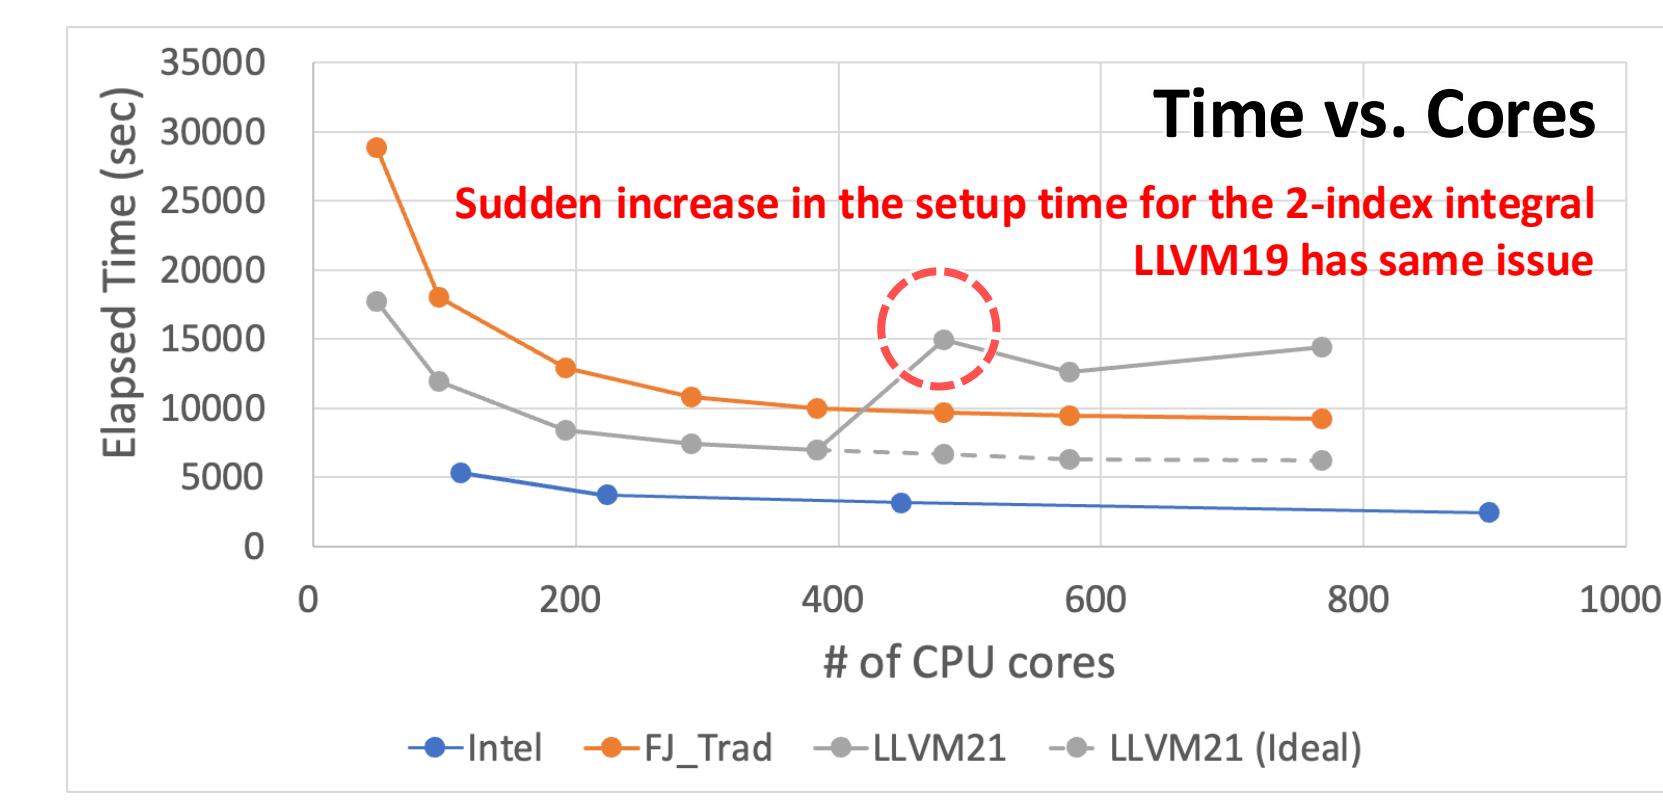


Efficiency Comparison with an Intel-Based Machine

Machine	Compiler	CPU core	SCF (sec)	Grad (sec)	Elapsed Time (sec)	Runtime Ratio
HBW2	Intel	896	1101.60	336.04	1937.14	x 1.00
Fugaku	FJ (Traditional)	768	1587.06	1230.21	5027.28	x 2.60
Fugaku	LLVM21	384	2050.27	592.72	3918.86	x 2.02
Fugaku	LLVM21 (Ideal)	768	1735.47	434.51	3413.26	x 1.76



TIM (37,226 Atoms: 37 QM)  
QM Level: B3LYP-D3/cc-pVDZ  
NVT, 300K, 1,000 MD steps



Efficiency Comparison with an Intel-Based Machine

Machine	Compiler	CPU core	SCF (sec)	Grad (sec)	Elapsed Time (sec)	Runtime Ratio
HBW2	Intel	896	1155.34	633.71	2462.81	x 1.00
Fugaku	FJ (Traditional)	768	2098.65	4132.13	9269.55	x 3.76
Fugaku	LLVM21	384	2715.94	2508.10	6982.46	x 2.84
Fugaku	LLVM21 (Ideal)	768	2278.13	2285.26	6250.00	x 2.54

## Conclusion & Future work

This work demonstrates the feasibility of performing long-time, high-accuracy QM/MM-MD simulations on the Fugaku using the recent LLVM/Clang++ compiler optimized for C++ code. The optimized code achieved performance close to that of programs running on Intel CPUs with the Intel compiler. Further optimization of the code for the recent LLVM/Clang++ compiler is expected to yield additional performance improvements.