

Porting mini-apps to ARM HPC systems

Brian J Gravelle, University of Oregon - Mentors: Dave Nystrom, Howard Pritchard

HPC-ENV, Application Readiness Team

Presented at the HPC 2019 Student Mini-Showcase, Aug. 1, 2019

Introduction

ARM in HPC

- HPC has long been dominated by x86 and IBM
- ARM offers an efficient alternative

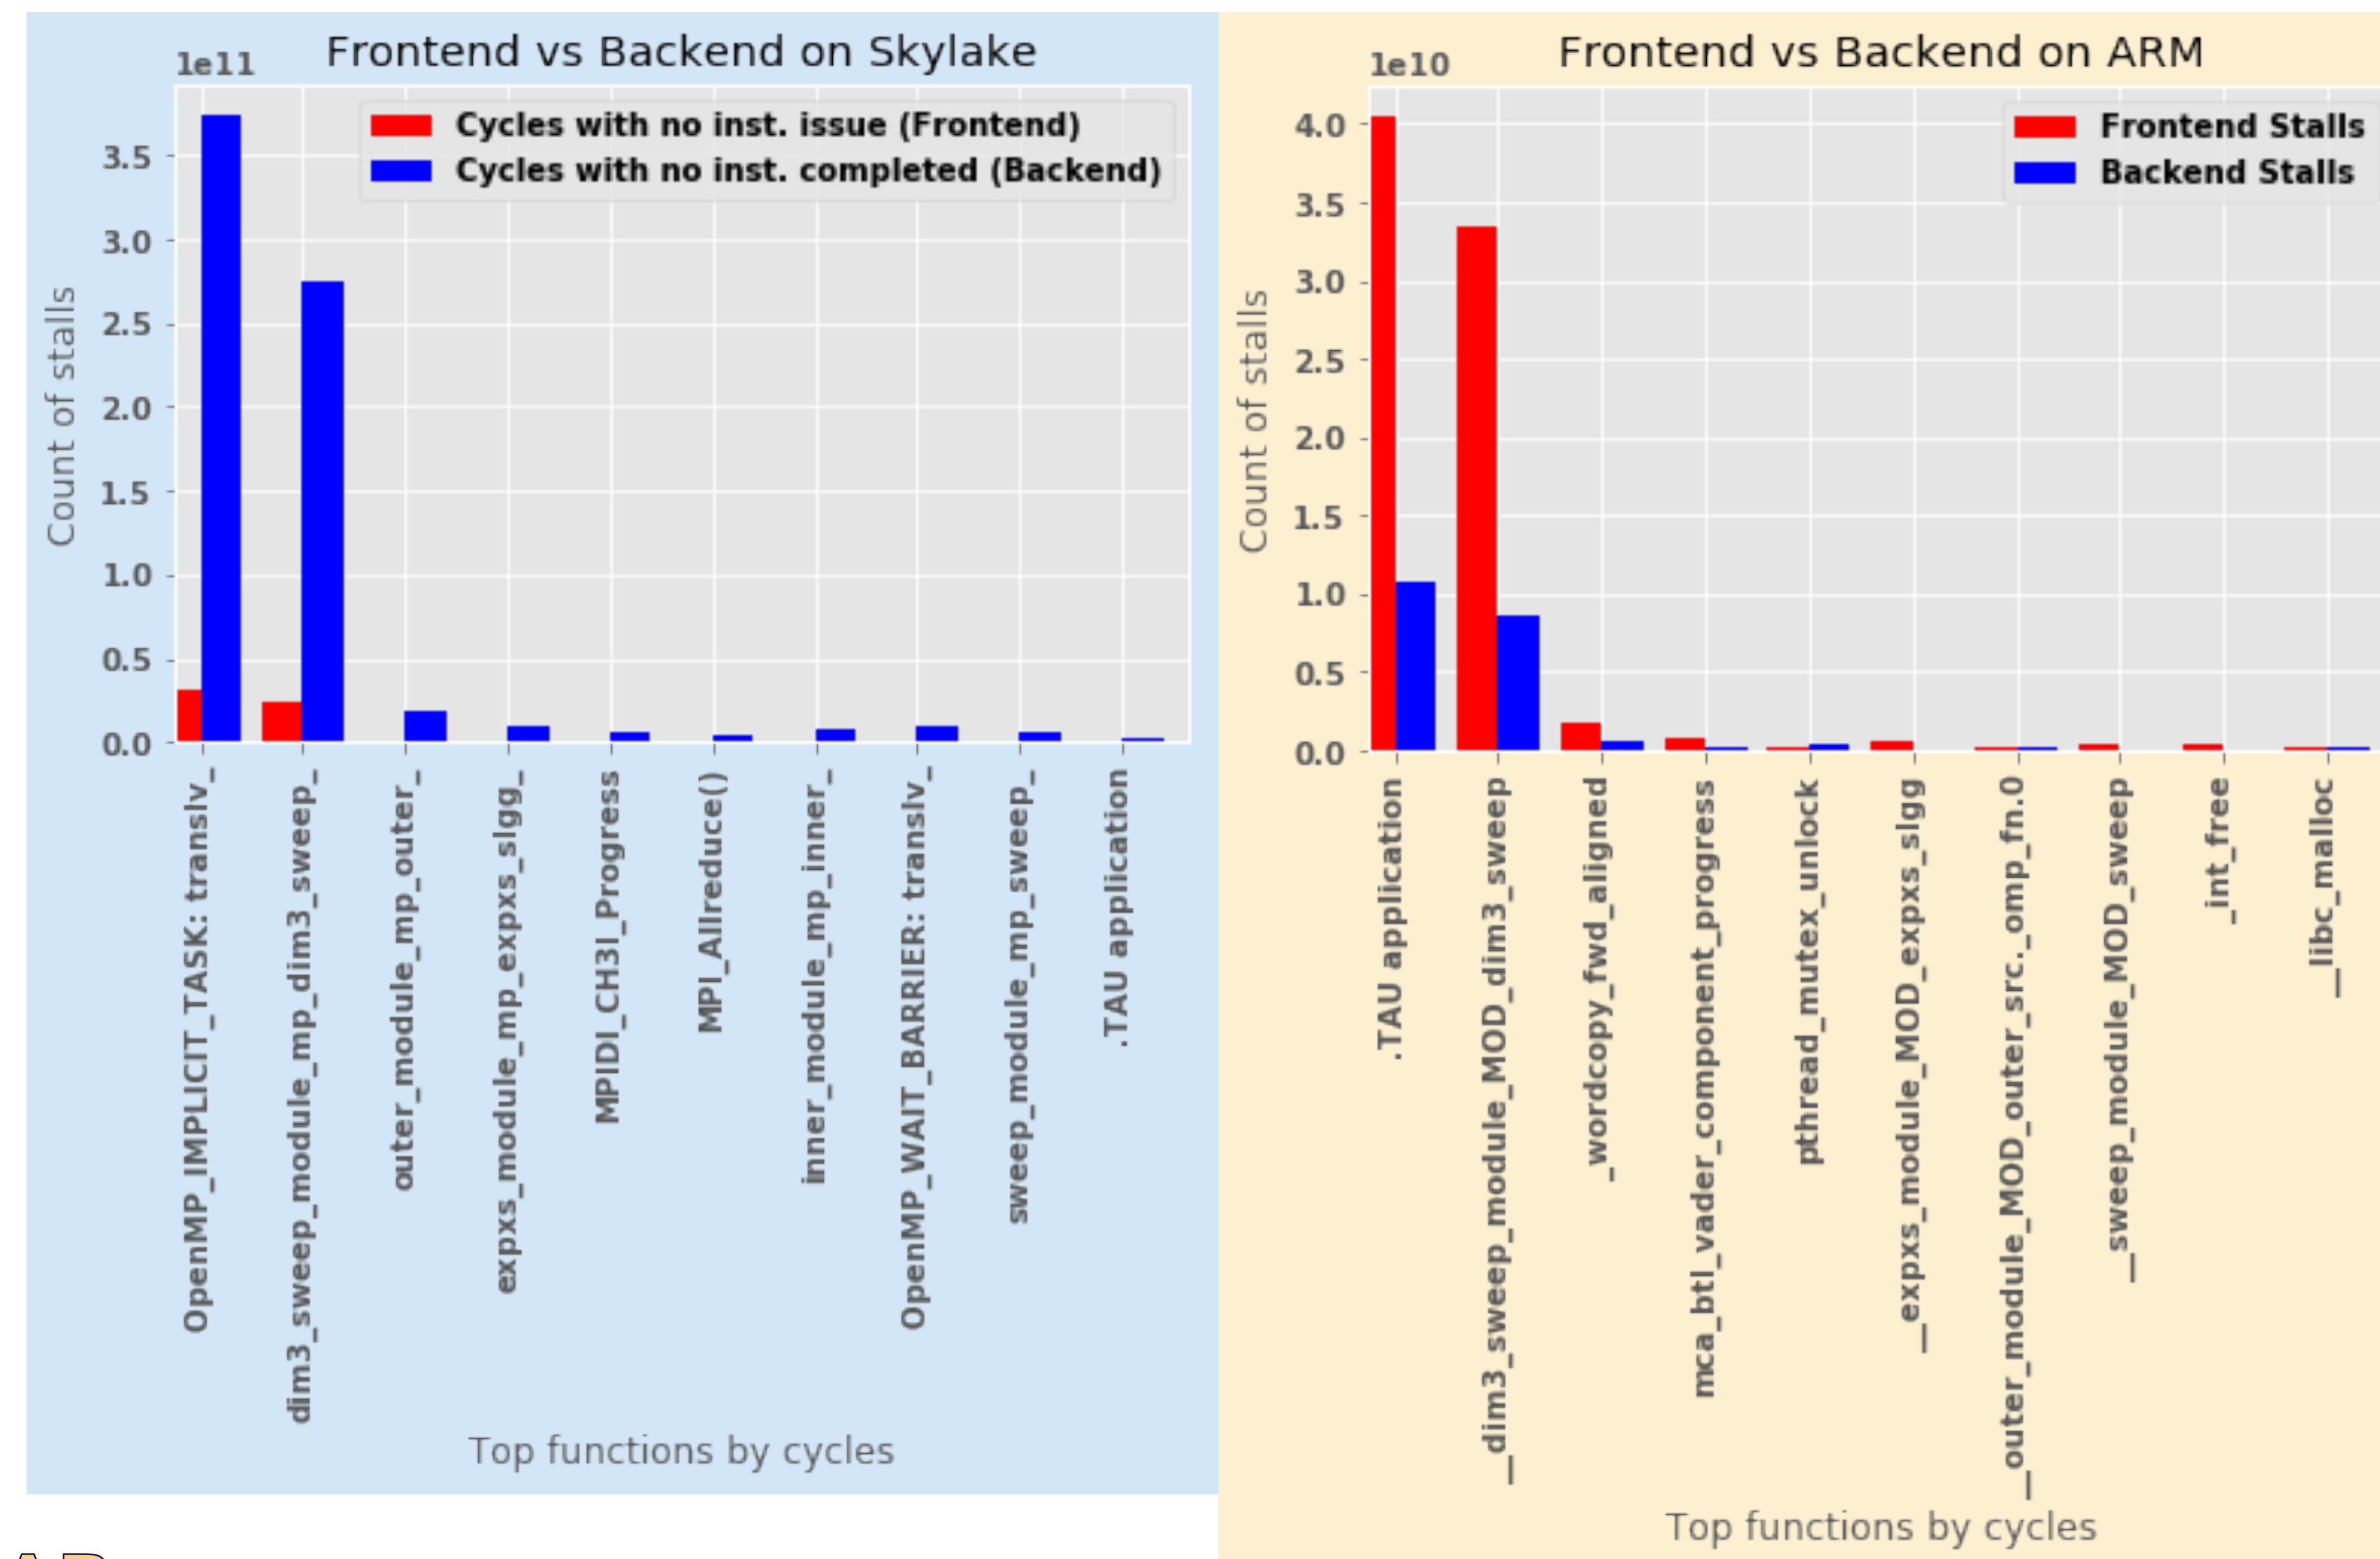
Methodology

- SNAP - OpenMP, MPI, Structured mesh
- Branson - MPI, Monte Carlo
- Performance sampling with TAU

Systems*

	Skylake Gold 6152	ThunderX2-B1
Cores	44	56
Threads per core	2	4
Clock	2.1GHz	2.0GHz
L1 data cache	32K	32K
L2 cache	1024K	256K
L3 cache	30976K	32768K
SIMD instructions	Up to 512 bit	128 bit NEON

Comparison of Frontend and Backend stalls for SNAP



Discussion of Results

Instruction mix

- ARM spends far more time branching
- Perhaps aggravated by lack of gather/scatter SIMD operations
- ARM uses fewer SIMD instructions

Frontend vs Backend

- ARM saturates the instruction pipelines
- Not the case when ARM uses 1 thread per MPI process
- Shows that optimizing the original code will be very different for each architecture

Energy Use

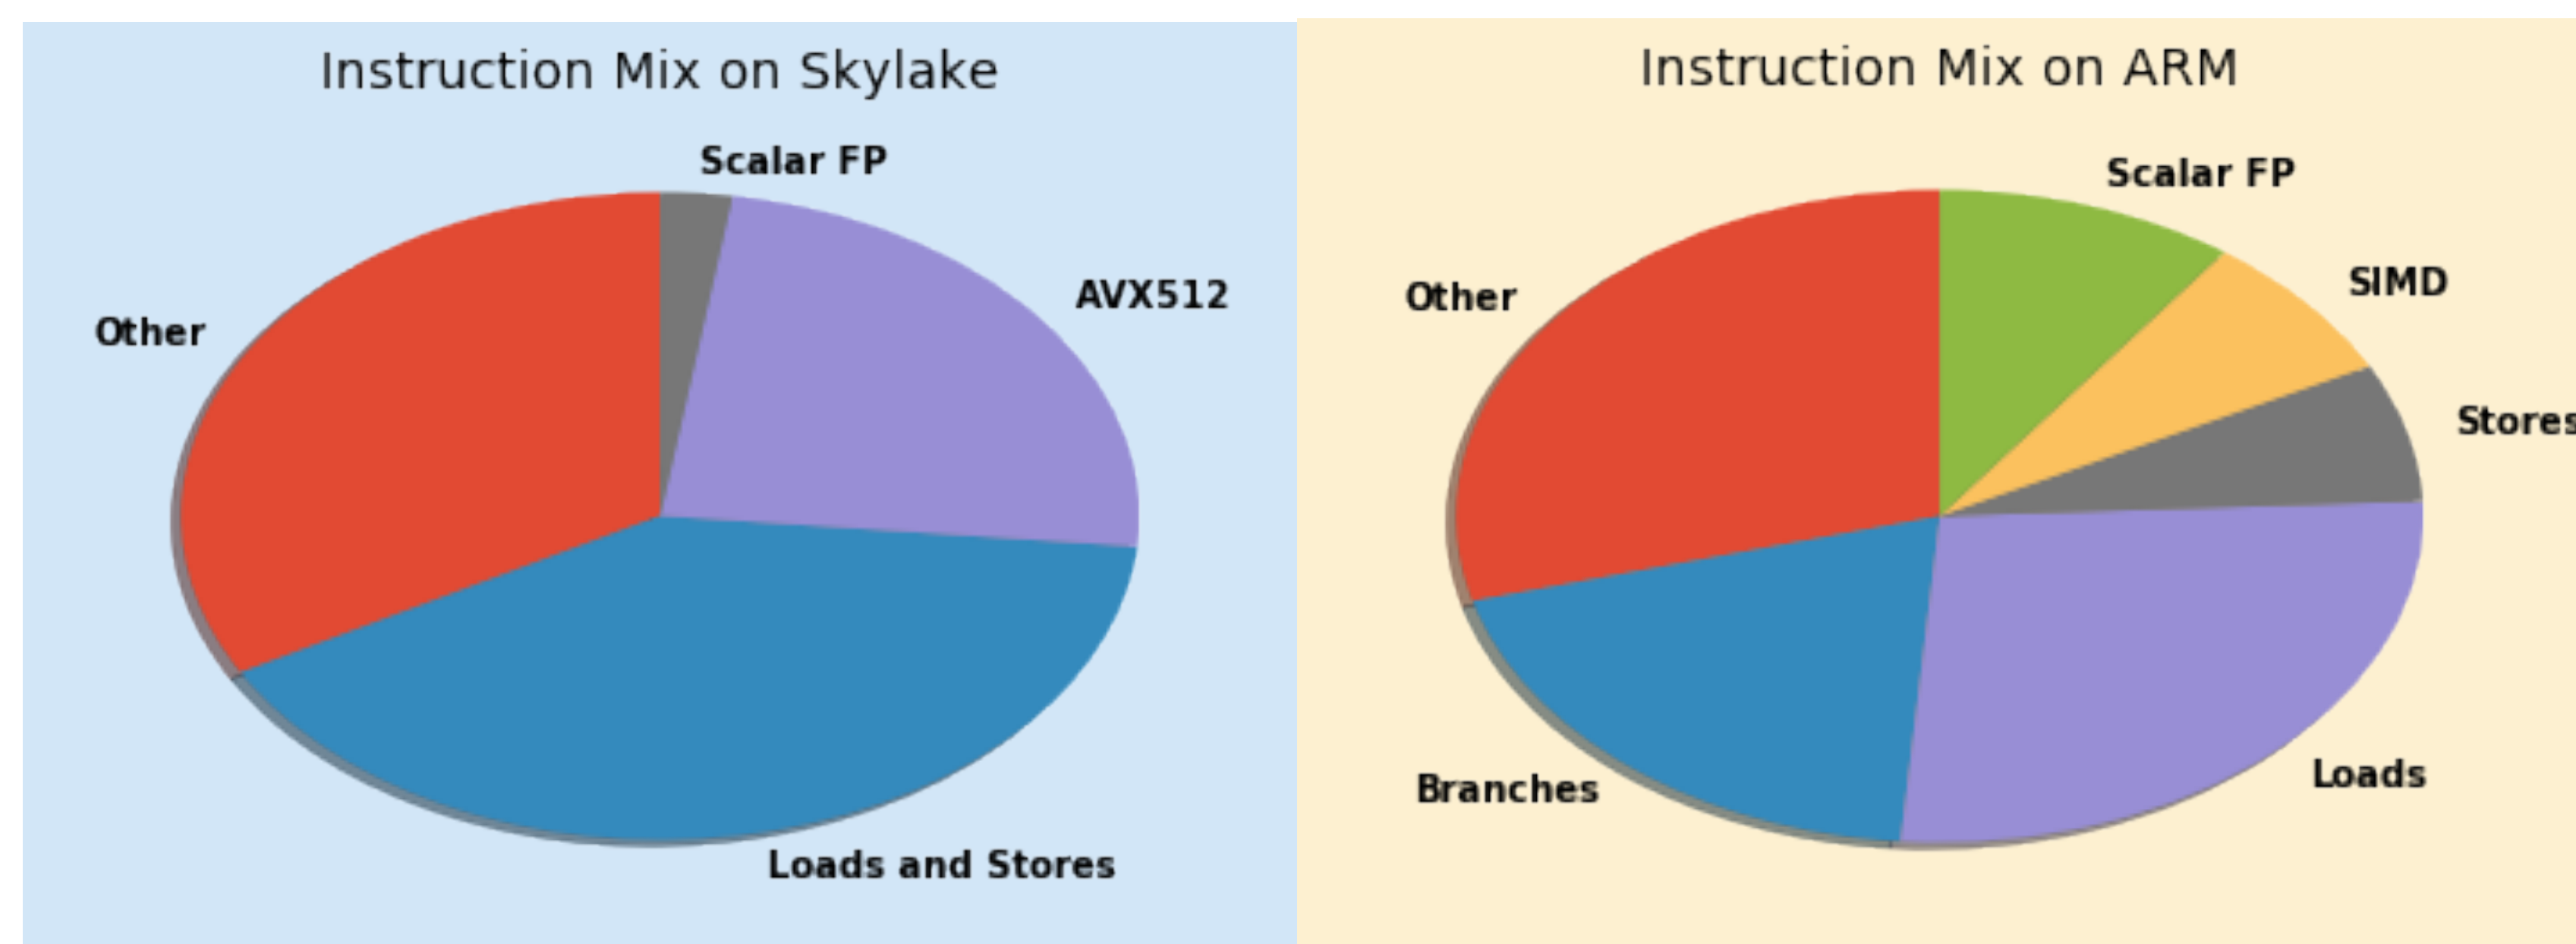
- ARM used significantly more energy
- Intel operates at a higher average power
- Finding ways to accelerate the ARM version could significantly improve energy consumption

Conclusion

In conclusion, the supposed advantages of ARM systems (energy efficiency and high thread counts) failed to materialize when compared to an Intel Skylake-Gold. However, no improvements have been made to fit the code to ARM, so minor adjustments may change this analysis.

Additionally, the next generation of ARM-based HPC systems is expected to have more advanced SIMD instructions (SVE) which should improve the energy efficiency and thread performance of the systems.

Instruction Mix comparison for SNAP



Energy use comparison for SNAP

