

## **IEEE Copyright Notice**

© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Accepted to be Published in: Proceedings of the 33rd IEEE International Parallel & Distributed Processing Symposium, May 20-24, 2019 Rio de Janeiro, Brazil

# Double-precision FPUs in High-Performance Computing: an Embarrassment of Riches?

Jens Domke<sup>\*,§</sup>, Kazuaki Matsumura<sup>†</sup>, Mohamed Wahib<sup>‡</sup>, Haoyu Zhang<sup>†</sup>, Keita Yashima<sup>†</sup>,  
Toshiki Tsuchikawa<sup>†</sup>, Yohei Tsuji<sup>†</sup>, Artur Podobas<sup>†,§</sup>, Satoshi Matsuoka<sup>§,†</sup>

<sup>\*</sup>Global Scientific Information and Computing Center, Tokyo Institute of Technology

<sup>†</sup>Department of Mathematical and Computing Science, Tokyo Institute of Technology

<sup>‡</sup>AIST-TokyoTech Real World Big-Data Computation Open Innovation Laboratory, Tokyo, Japan

<sup>§</sup>RIKEN Center for Computational Science (R-CCS), RIKEN, Japan

**Abstract**—Among the (uncontended) common wisdom in High-Performance Computing (HPC) is the applications’ need for large amount of double-precision support in hardware. Hardware manufacturers, the TOP500 list, and (rarely revisited) legacy software have without doubt followed and contributed to this view.

In this paper, we challenge that wisdom, and we do so by exhaustively comparing a large number of HPC proxy applications on two processors: Intel’s Knights Landing (KNL) and Knights Mill (KNM). Although similar, the KNL and KNM architecturally deviate at one important point: the silicon area devoted to double-precision arithmetics. This fortunate discrepancy allows us to empirically quantify the performance impact in reducing the amount of hardware double-precision arithmetic.

Our analysis shows that this common wisdom might not always be right. We find that the investigated HPC proxy applications do allow for a (significant) reduction in double-precision with little-to-no performance implications. With the advent of a failing of Moore’s law, our results partially reinforce the view taken by modern industry (e.g., upcoming Fujitsu ARM64FX) to integrate hybrid-precision hardware units.

## I. INTRODUCTION

It is becoming increasingly clear that the road forward in High-Performance Computing (HPC) is one full of obstacles. With the ending of Dennard’s scaling [1] and the ending of Moore’s law [2], there is today an ever-increasing need to oversee how we allocate the silicon to various functional units in modern many-core processors. Amongst those decisions is how we distributed the hardware support for various levels of compute-precision.

Historically, most of the compute silicon has been allocated to double-precision (DP; 64-bit) compute. Nowadays – in processors such as the forthcoming AA64FX [3] and NVIDIA Volta [4] – the trend, mostly driven by market/AI demands, is to replace some of the double-precision units with lower-precision units. Lower-precision units occupy less area (up to  $\approx 3x$  going from double- to single-precision Fused-Multiply-Accumulate [5]), leading to more on-chip resources (more instruction-level parallelism), potentially lowered energy consumption, and a definitive decrease in external memory bandwidth pressure (i.e., more values per unit bandwidth). The gains – up to four times over their DP variants with little loss in accuracy [6] – are attractive and clear, but what is the impact on performance (if any) on existing HPC

applications? What performance impact can HPC users expect when migrating their code to future processors with a different distribution in floating-point precision support? Finally, how can we empirically quantify this impact on performance using existing processors in an apples-to-apples comparison on real-life use cases without relying on tedious, slow, and potentially inaccurate simulators?

The Intel Xeon Phi was supposed to be the high-end for many-core processor technology for nearly a decade (Knights Ferry was announced in 2010), and has changed drastically since its first released. The latest (and also last) two revisions – the Knights Landing and Knights Mill – are of particular importance since they arguable reflect two different ways of thinking. Knights Landing has relatively large support for double-precision (64-bit) computations, and follows a more traditional school of thought. While Knights Mill follows a different direction, which is the replacement of double-precision compute units with lower-precision (single-precision, half-precision, and integer) compute capabilities.

In the present paper, we quantify and analyze the performance and compute bottlenecks of Intel’s Knights Landing [7] and Knights Mill [8] architectures – two processors with identical micro-architecture where the main difference is in the relative allocation of double-precision units. We stress both processors with numerous realistic benchmarks from both the Exascale Computing Project (ECP) proxy applications [9] and RIKEN R-CCS Fiber Miniapp Suite [10] – benchmarks used in HPC system acquisition. Through an extensive (and robust) performance measurement process (which we also open-source), we empirically show the architecture’s relative weaknesses. In short, the contributions of the present paper are:

- 1) An empirical performance evaluation of the Knights Landing and Mill family of processors – both proxies for previous and future architectural trends – with respect to benchmarks derived from realistic HPC workloads,
- 2) An in-depth analysis of results, including identification of bottlenecks for the different application/architecture combinations, and
- 3) An open-source compilation of our evaluation methodology, including our collected raw data.

TABLE I

DETAILED COMPUTE NODE HARDWARE INFORMATION; DIFFERENCES BETWEEN KNIGHTS LANDING & MILL HIGHLIGHTED IN BOLD; SHOWN BANDWIDTH (BW) MEASURED WITH BABELSTREAM (SEE SEC.II-B); NUMBERS FOR DUAL-SOCKET REFERENCE SYSTEM ACCUMULATED

Feature	KNL	KNM	Broadwell-EP
CPU Model	<b>7210F</b>	<b>7295</b>	2x E5-2650v4
#{Cores} (HT)	<b>64</b> (4x)	<b>72</b> (4x)	24 (2x)
Base Frequency	<b>1.3 GHz</b>	<b>1.5 GHz</b>	2.2 GHz
Max Turbo Freq.	<b>1.5 GHz</b>	<b>1.6 GHz</b>	2.9 GHz
CPU Mode	Quadrant	Quadrant	N/A
TDP	<b>230 W</b>	<b>320 W</b>	210 W
DRAM Size	96 GiB	96 GiB	256 GiB
↔ Triad BW	<b>71 GB/s</b>	<b>88 GB/s</b>	122 GB/s
MCDRAM Size	16 GiB	16 GiB	N/A
↔ Triad BW	<b>439 GB/s</b>	<b>430 GB/s</b>	N/A
MCDRAM Mode	Cache	Cache	N/A
LLC Size	<b>32 MiB</b>	<b>36 MiB</b>	60 MiB
Inst. Set Extension	AVX-512	AVX-512	AVX2
FP32 Peak Perf.	<b>5,324 Gflop/s</b>	<b>13,824 Gflop/s</b>	1,382 Gflop/s
FP64 Peak Perf.	<b>2,662 Gflop/s</b>	<b>1,728 Gflop/s</b>	691 Gflop/s

## II. ARCHITECTURES, ENVIRONMENT, AND APPLICATIONS

Our research objective is to evaluate the impact of migrating from an architecture with (relatively) high amount of double-precision compute to an architecture with less. By high amount of double-precision compute we mean architectures whose Floating-Point Unit (FPU) has most of its silicon dedicated to 64-bit IEEE-754 floating-point operations, and by less double-precision compute we mean architectures that replace those same double-precision FPUs with lower – potentially hybrid – precision units.

To understand and explore the intersection of architectures with high-amount of double-precision and those with hybrid-precision, there is a need to find a processor whose architecture is unchanged with the sole exception of its floating-point unit to silicon distribution. Only one modern processor family allows for such an apples-to-apples comparison: the Xeon Phi family of processors.

### A. Hardware & Software Environment

Intel’s Knights Landing (KNL) and Knights Mill (KNM) are the latest incarnations of a long line of architectures in the Intel’s accelerator family. Both processor consist of a large number of processors cores (64 and 72, respectively), interconnected in a 2-D mesh (prior to KNL: ring interconnection). Each core has a private L1 cache and a slice of the distributed L2 cache. Caches are kept coherent through the directory-based MESIF protocol. Both processors come with two types of external memory: MCDRAM (or, Hybrid Memory Cube) and Double-Data Rate-synchronous (DDR4) memory. Unique to the Xeon Phi processors is that the MCDRAM memory can be configured to one of three modes of operation: it is either (1) directly addressable in the global memory address space (memory-mapped), called `flat` mode, or it (2) acts as last-level cache before the DDR, called `cache` mode. Finally, the third mode (`hybrid` mode [11]) is a combination of the properties from the first two modes.

There are several policies governing where data is homed. A common high-performance configuration [12], which is also the one we used in our study, is the quadrant mode. Quadrant

mode means that the physical cores are divided into four logical parts, where each logical part is assigned two memory controllers; each logical group is treated as a unique Non-Uniform Memory-Access (NUMA) node, allowing the operating system to perform data-locality optimizations. Table I surveys and contrasts the processors against each other, where the main differences are highlighted. The main architectural difference – which is also the difference and its impact we seek to empirically quantify – is the Floating-Point Unit (FPU). In KNL, this unit features two 512-bit wide vector units (AVX), together capable of executing 32 double-precision or 64 single-precision operations per cycle, totaling 2.6 Tflop/s of double- and 5.3 Tflop/s of single-precision peak performance, respectively, across all 64 processing cores. In KNM, however, the FPU is redesigned to replace one 512-bit vector unit with two Virtual Neural Network Instruction (VNNI) units. Those units, although specializing in hybrid-precision FMA, can execute single-precision vector instructions, but have no support for double-precision compute. Thus, in total, the KNM can execute up to 1.7 Tflop/s of double-precision or 13.8 Tflop/s of single-precision computations. In summary, the KNM has 2.59x more single-precision compute, while the KNL have 1.54x more double-precision compute.

While both the KNL and KNM are functionally and architectural similar, there are some note-worthy differences. First, the operating frequency of these processors varies: the KNL operates at a frequency of 1.3 GHz (and up to 1.5 GHz in Turbo mode), while KNM operates at 1.5 GHz (1.6 GHz turbo). Hence, KNM executes 15% more cycles per second over KNL. Furthermore, although the cores of KNM and KNL are similar (except the FPU), the number of cores is different: KNL has 64 cores while KNM has 72 cores. Both processors are manufactured in 14 nm technology. Finally, the amount of on-chip last-level cache between the two processors is different, where KNM has a 4 MiB advantage over KNL.

Additionally, for verification reasons, we include a modern dual-socket Xeon-based compute node in our evaluation. Despite being vastly different from the Xeon Phi systems, our Xeon Broadwell-EP (BDW) general-purpose processor is used to cross-check metrics, such as: execution time and performance (Xeon Phi should perform better), frequency-scaling experiments (BDW has more frequency domains), and performance counters (BDW exposes more performance counters). Aside from those differences mentioned above (and highlighted in Table I), the setup between the Xeon Phi nodes (and BDW node) is *identical*, including the same operating system, software stack, and solid state disk.

For the operating system (OS) and software environment, we use equivalent setups across our three compute nodes. The OS is a fresh installation of CentOS 7 (minimal) with Linux kernel version 3.10.0-862, which by default has the latest versions of the Meltdown and Spectre patches enabled. During our experiments, we limit potential OS noise by disabling all remote storage (Network File System in our case) and allowing only a single user on the system. Most of our applications are compiled with Intel’s Parallel Studio XE

(version 2018; update 3) compilers, and we install the latest versions of Intel TensorFlow and Intel MKL-DNN for the deep learning proxy application, since our assumption is that Intel’s software stack allows for the highest utilization of their hardware. Exceptions to this compiler selection are listed in the subsequent Section II. Furthermore, we use Intel MPI from the Parallel Studio XE suite to execute our measurements.

## B. Benchmark Applications

Over the years, the HPC community developed many benchmarks that represent real workloads in order to test the capabilities of a system – primarily for comparisons across architectures but also for system procurement purposes. The so-called Exascale Computing Project (ECP) proxy applications [9] and RIKEN R-CCS’ (f.k.a. AICS) Fiber Miniapp Suite [10], which we will focus on for this study, are just two examples representing modern HPC workloads. Those benchmarks are designed to evaluate single-node and small-scale test installations, and hence are adequate for our study.

1) *The ECP Proxy-Apps*: The ECP suite (release v1.0) consists of 12 proxy applications primarily written in C (5x), FORTRAN (3x), C++ (3x), and Python (1x), listed hereafter.

a) *Algebraic multi-grid (AMG)*: solver of the *hypr* library is a parallel solver for unstructured grids [13] arising from fluid dynamics problems. We choose *problem 1* for our tests, which applies a 27-point stencil on a 3-D linear system.

b) *CANDLE (CNDL)*: is a deep learning benchmark suite to tackle various problems in cancer research [14]. We select benchmark 1 of pilot 1 (*PIB1*), which builds an autoencoder from a sample of gene expression data to improve the prediction of drug responses.

c) *Co-designed Molecular Dynamics (CoMD)*: serves as the reference implementation for ExMatEx [15] to facilitate co-design for (and evaluation of) classical molecular dynamics algorithms. We are using the included strong-scaling example to calculate the inter-atomic potential for 256,000 atoms.

d) *LAGrangian High-Order Solver – Laghos (LAGO)*: computes compressible gas dynamics through an unstructured high-order finite element method [16]. The input for our study is the simulation of a 2-dimensional Sedov blast wave with default settings as documented for the Laghos proxy-app.

e) *MACSio (MxIO)*: is a synthetic Multi-purpose, Application-Centric, Scalable I/O proxy designed to closely mimic realistic I/O workloads of HPC applications [17]. Our input causes MACSio to write a total of 433.8 MB to disk.

f) *MiniAMR (MAMR)*: is an adaptive mesh refinement proxy application of the Mantevo project [18] which applies a stencil computation on a 3-dimensional space, in our case a sphere moving diagonally through a cubic medium.

g) *MiniFE (MiFE)*: is a reference implementation of an implicit finite elements solver [18] for scientific methods resulting in unstructured 3-dimensional grids. For our study, we use  $128 \times 128 \times 128$  input dimensions for the grid.

h) *MiniTri (MTri)*: is able to apply different graph detection algorithms for a given graph, such as community detection or dense subgraph detection [19]. As input for the triangle

detection and approximation of the graph’s largest clique, we download *BCSSTK30* from the MatrixMarket [20].

i) *Nekbone (NekB)*: is a proxy for the Nek5000 application [21], and uses the conjugate gradient method for solving the standard Poisson equation for computational fluid dynamics problems. We enabled the multi-grid preconditioner, and for strong-scaling, see Section III-B, we fixed the elements per process and polynomial order to one number, respectively.

j) *SW4lite (SW4L)*: is a proxy for the computational kernels used in the seismic modelling software, called SW4 [22], and we use the *pointsource* example, which calculates the wave propagation emitted from a single point in a half-space.

k) *SWFFT (FFT)*: represents the compute kernel of the HACC cosmology application [23] for N-body simulations. The 3-D fast Fourier transformation of SWFFT emulates one performance-critical part of HACC’s Poisson solver. In our tests, we perform 32 repetitions on a  $128 \times 128 \times 128$  grid.

l) *XSBench (XSBn)*: is the proxy for the Monte Carlo calculations used by a neutron particle transport simulator for a Hoogenboom-Martin nuclear reactor [24]. We simulate a *large* reactor model represented by a *unionized* grid with  $15 \cdot 10^6$  cross-section lookups per particle.

2) *RIKEN Mini-Apps*: In comparison to the modernized ECP proxy-apps, RIKEN’s eight mini-apps are written in FORTRAN (4x), C (2x), and a mix of FORTRAN/C/C++ (2x).

a) *FrontFlow/blue (FFB)*: uses the finite element method to solve the incompressible Navier-Stokes equation for thermo-fluid analysis [25]. We simulate the 3-D cavity flow in a rectangular space discretized into  $50 \times 50 \times 50$  cubes.

b) *Frontflow/violet Cartesian (FFVC)*: falls into the same problem class as FFB, however the difference is that FFVC uses the finite volume method (FVM) [26]. Here, we calculate the 3-D cavity flow in a  $144 \times 144 \times 144$  cuboid.

c) *MODYLAS (MDYL)*: makes use of the fast multipole method for long-range force evaluations in molecular dynamics simulations [27]. Our input is the *wat222* example which distributes 156,240 atoms over a  $16 \times 16 \times 16$  cell domain.

d) *many-variable Variational Monte Carlo (mVMC) method*: implemented by this mini-app is used to simulate quantum lattice models for studying the physics of condensed matter [28]. We use mVMC’s included strong-scaling test, but downsize it (1/3 lattice dimensions and 1/4 of samples).

e) *Nonhydrostatic ICosahedral Atmospheric Model (NICM)*: is a proxy of NICAM, which computes mesoscale convective cloud systems based on FVM for icosahedral grids [29]. We run Jablonowski’s baroclinic wave test (*gl05r100z40pe10*), but reduce the simulated days from 11 to 1.

f) *Next-Gen Sequencing Analyzer (NGSA)*: is a mini-app of a genome analyzer and a set of alignment tools designed to facilitate cancer research by detecting genetic mutations in human DNA [30]. For our experiments, we rely on pre-generated pseudo-genome data (*ngsa-dummy*).

g) *NTChem (NTCh)*: implements a computational kernel of the NTChem software framework for quantum chemistry calculations of molecular electronic structures, i.e., the solver

TABLE II  
APPLICATION CATEGORIZATION, COMPUTE PATTERNS, AND MAIN PROGRAMMING LANGUAGES USED; MACSIO, HPL, HPCG, AND BABELSTREAM BENCHMARKS OMITTED

ECP	Scientific/Engineering Domain	Compute Pattern	Language
AMG	Physics and Bioscience	Stencil	C
CANDLE	Bioscience	Dense matrix	Python
CoMD	Material Science/Engineering	N-body	C
Laghos	Physics	Irregular	C++
miniAMR	Geoscience/Earthscience	Stencil	C
miniFE	Physics	Irregular	C++
miniTRI	Math/Computer Science	Irregular	C++
Nekbone	Math/Computer Science	Sparse matrix	Fortran
SW4lite	Geoscience/Earthscience	Stencil	C
SWFFT	Physics	FFT	C/Fortran
XSbench	Physics	Irregular	C
RIKEN	Scientific/Engineering Domain	Compute Pattern	Language
FFB	Engineering (Mechanics, CFD)	Stencil	Fortran
FFVC	Engineering (Mechanics, CFD)	Stencil	C++/Fortran
mVMC	Physics	Dense matrix	C
NICAM	Geoscience/Earthscience	Stencil	Fortran
NGSA	Bioscience	Irregular	C
MODYLAS	Physics and Chemistry	N-body	Fortran
NTChem	Chemistry	Dense matrix	Fortran
QCD	Lattice QCD	Stencil	Fortran/C

for the second-order Møller-Plesset perturbation theory [31]. We select the H<sub>2</sub>O test case for our study.

*h) Quantum ChromoDynamics (QCD):* mini-app solves the lattice QCD problem in a 4-D lattice (3-D plus time), represented by a sparse coefficient matrix, to investigate the interaction between quarks [32]. We evaluate QCD with the *Class 2* input for a  $32^3 \times 32$  lattice discretization.

*3) Reference Benchmarks:* In addition to those 20 applications, we use the compute intensive HPL [33] benchmark, and HPCG [34] and stream (both memory intensive) to evaluate the baseline of the investigated architectures.

*a) High Performance Linpack (HPL):* is solving a dense system of linear equations  $Ax = b$  to demonstrate the double-precision compute capabilities of a (HPC) system [35]. Our problem size is 64,512. For both HPL and HPCG (see below), we employ highly tuned versions shipped with Intel’s Parallel Studio XE suite with appropriate parameters for our systems.

*b) High Performance Conjugate Gradients (HPCG):* is applying a conjugate gradient solver to a system of linear equation (sparse matrix  $A$ ), with the intent to demonstrate the system’s memory subsystem and network limits. We choose  $360 \times 360 \times 360$  as global problem dimensions for HPCG.

*c) BabelStream (BABL):* is one of many available “stream” benchmarks supporting evaluations of the memory subsystem for CPUs and accelerators [36]. We will use 2 GiB and 14 GiB input vectors, see Section IV-C for details.

We provide a compressed overview of the ECP and RIKEN’s proxy applications in Table II. In this table, each application is categorized by its scientific domain, as well as the primary workload/kernel classification, for which we use the classifiers employed by Hashimoto et al. [37]. Both, the scientific domain as well as the kernel classification will be important for our subsequent analysis in Sections IV and V.

### III. METHODOLOGY

In this section, we present our rigorous benchmarking approach into investigating the characteristics of each architecture, and extracting the necessary information for our study.

#### A. Benchmark Setup and Configuration Selection

Due to the fact that the benchmarks, listed in Section II-B, are firstly realistic proxies of the original applications [38] and secondly are used in the procurement process, we can assume that these benchmarks are well tuned and come with appropriate compiler options for a variety of compilers – a hypothesis we will test in Section IV-D. Hence, we refrain from both manual code optimization and alterations of the compiler options. The only modifications we perform are:

- Enabling interprocedural optimization (`-ipo`) and compilation for the highest instruction set available (`-xHost`)<sup>1</sup>,
- Patching a segmentation fault in MACSio<sup>2</sup>, and
- Injecting our measurement source code, see Section III-B.

With respect to the measurement runs, we follow this five step approach for each benchmark:

- 0) Install, patch, and compile the benchmark, see above,
- 1) Select appropriate inputs/parameters/seeds for execution,
- 2) Determine “best” parallelism: #processes and #threads,
- 3) Execute a *performance*, a *profiling*, and a *frequency* run,
- 4) Analyze the results (go to 0. if anomalies are detected).

and we will further elaborate on those steps hereafter.

For the input selection we have to balance between multiple constraints and choose based on: Which recommended inputs are listed by the benchmark developers?, How long does the benchmark run?<sup>3</sup> Does it occupy a realistic amount of main memory (e.g., avoid cache-only executions)? Are the results repeatable (randomness/seeds)? We optimize for the metrics reported by the benchmark (e.g., select the input with the highest Gflop/s rate).

Furthermore, one of the most important consideration while selecting the right inputs is *strong-scaling*. We require strong-scaling properties of the benchmark for two reasons: the results collected in Step (2) need to be comparable, and even more importantly, the results of Step (3) must be comparable between different architectures, since we may have to use different numbers of MPI processes for KNL and KNL (and our BDW reference architecture) due to their difference in core counts. The only exception is MiniAMR for which we are unable to find a strong-scaling input configuration and instead optimized for the reported Gflop/s of the benchmark. Accordingly, we then choose the same amount of MPI processes on our KNL and KNL compute nodes for MiniAMR.

In Step (2), we evaluate numerous combinations of MPI processes and OpenMP threads for each benchmark, including combinations which over-/undersubscribe the CPU cores, and test each combination with three runs to minimize the

<sup>1</sup> Exceptions: (a) AMG compiled with `-xCORE-AVX2` to avoid arithmetic errors; (b) NGSA’s BWA tool compiled with GNU gcc to avoid segfaults.

<sup>2</sup> After our reporting, the developers patched the upstream version.

<sup>3</sup> Our aim is 1 sec–10 min due to the large sample size we have to cover.

potential for outliers due to system noise. For all subsequent measurements, we select the number of processes and threads based on the “best” (w.r.t time-to-solution of the solver) combination among these tested versions, see Table IV at the end of this paper for details. We are not applying specific tuning options to Intel’s MPI library, except for using Intel’s recommended settings for HPCG with respect to thread affinity and MPI\_allreduce. The reason is that our pretests (with a subset of the benchmarks) with non-default parameters for Intel MPI consistently resulted in longer time-to-solution.

For Step (3), we run each benchmark ten times to identify the fastest time-to-solution for the (compute) kernel of the benchmark. Additionally, for the profiling runs, we execute the benchmark once for each of the profiling tools and/or metrics (in case the tool is used for multiple metrics), see Section III-B for details. Finally, we perform frequency scaling experiments for each benchmark, where we throttle the CPU frequency to all of the available lower CPU states below the maximum CPU frequency, which we use for the performance runs, and record the lowest kernel time-to-solution among ten trials per frequency. The reason and results of the frequency scaling test will be further explained in Section IV-E. One may argue for more than ten runs per benchmark to find the optimal time-to-solution, however, given the prediction interval theory and our deterministic benchmarks executed on a single node, it is unlikely to obtain a much faster run and we confirmed that the fastest 50% of executions per benchmark only vary by 3.9% on average. The collected metrics, see the following section, will be analyzed in Section IV in detail.

### B. Metrics and Measurement Tools

To study and analyze the floating point requirements by applications, it is not only important to evaluate an established metric (floating point operations per second), but also other metrics, such as memory throughput, cache utilization, or speedup with increased CPU frequency. The detailed list of metrics (and derived metrics) and the methodology and tools we use to collect these metrics will be explained hereafter.

One observation is that the amount of time spent on initializing and post processing within each proxy application can be relatively high (e.g., HPCG spends only 11% and 30% of its time in the solver part on BDW and Phi, respectively) and is usually not consistent with the real workloads, e.g., one can reduce the epochs for performance evaluation purposes in CANDLE but not the input data pre-processing to execute those epochs. These mismatches in kernel-to-[pre|post]processing ratio requires us to extract all metrics only for the (computational) kernel of the benchmark. Hence, we identify and inject profiling instructions around the kernels to start or pause the collection of raw metric data by the analysis tools. This code injection is exemplified in PseudoCode 1. Therefore, unless otherwise stated in this Section or subsequent sections, all presented data will be based exclusively on the kernel portion of each benchmark.

For tool stability reason, attention to detail/accuracy, and overlap with our needs, we settle on the use of the MPI API

### PseudoCode 1: Injecting analysis instructions

```
#define START_ASSAY {measure time; toggle on [PCM | SDE | VTune]}
#define STOP_ASSAY {measure time; toggle off [PCM | SDE | VTune]}
Function main is
  STOP_ASSAY
  Initialize benchmark
  foreach solver loop do
    START_ASSAY
    Call benchmark solver/kernel
    STOP_ASSAY
    Post-processing
  Verify benchmark result
  START_ASSAY
```

TABLE III  
SUMMARY OF METRICS AND METHOD/TOOL TO COLLECT THESE METRICS

Raw Metric	Method/Tools
Runtime [s]	MPI_Wtime()
{FP / integer operations}	SDE
{Branches operations}	SDE
Memory throughput [B/s]	PCM (pcm-memory.x)
{L2/LLC cache hits/misses}	PCM (pcm.x)
Consumed Power [Watt]	PCM (pcm-power.x)
SIMD instructions per cycle	perf + VTune ('hpc-performance')
Memory/Back-end boundedness	perf + VTune ('memory-access')

for runtime measurements, alongside with Intel’s Processor Counter Monitor (PCM) [39], Intel’s Software Development Emulator (SDE) [40], and Intel’s VTune Amplifier [41]<sup>4</sup>. Furthermore, as auxiliary tools we rely on RRZE’s Likwid [42] for frequency scaling<sup>5</sup> and LLNL’s msr-safe [43] for allowing access to CPU model-specific registers. An overview of (raw) metrics which we extract with these tools for the benchmarks, listed in Section II-B, is shown in Table III. Furthermore, derived metrics, such as Gflop/s, will be explained on-demand in Section IV.

## IV. EVALUATION

The following subsections will primarily focus on visualizing and analyzing the key metrics we collect for each proxy- and mini-app, such as Gflop/s. The significance of our findings with respect to future software, CPU, and HPC system design will then be discussed in the next Section V. While we will argue for less flop/s-centric performance reporting of HPC benchmarks in Section V, we have to adhere the current standards. By analyzing the performed FP operation/s instead of concealing them, not only do we gain insight into realistic flop/s of HPC applications, we also have the capability to evaluate FP unit requirements. Furthermore, this analysis will strengthen our argument that flop/s should not be the only reported performance metric – especially if the majority of benchmarks does not even achieve 10% of theoretical peak.

Furthermore, analyzing other metrics such as the instruction mix, time-to-solution, or memory throughput, see Section IV-A, IV-B, and IV-C, in an isolated fashion does also not give good indications about the system’s bottlenecks, and

<sup>4</sup> To avoid persistent compute node crashes (likely due to incompatibilities with the Spectre/Meltdown patches), we had to disable VTune’s build-in sampling driver and instead rely on Linux’ perf tool.

<sup>5</sup> Our Linux kernel version required us to disable the default Intel P-State driver to have full access to the fine-grained frequency scaling.




Fig. 1. Ratio of integer vs. single-precision FP vs. double-precision FP per proxy-app as counted by Intel’s SDE; Per application: **Left bar = BDW**, **middle bar = KNL**, **right bar = KNM**; Missing bars for CANDLE due to SDE crashes on Xeon Phi; Proxy-app abbreviations acc. to Section II-B

hence, especially when reasoning about FPU requirements, we have to understand the applications’ compute-boundedness, which we evaluate in Section IV-E. Only when analyzing all these metrics in the same context, we attain the needed understanding. Table III summarizes the primary metrics and method/tool used to collect these metrics, while Table IV includes additionally collected metrics.

### A. Integer vs. Single-Precision FP vs. Double-Precision FP

The breakdown of total number of integer and single/double-precision floating point (FP) operations, as depicted in Figure 1, shows two rather unexpected trends. First, the number of proxy-apps relying on 32-bit FP instructions is four out of 22, which is surprisingly low, and furthermore, only one of them utilizes both 32-bit and 64-bit FP instructions. Minor variances in integer to FP ratio between the architectures can likely be explained by the difference in AVX vector length, quality of compiler optimization for each CPU, and execution/parallelization approach. The second unexpected trend is the imbalance of integer to FP operations, i.e., 16 of 22 applications issue at least 50% integer operations. However, one has to keep in mind that the Intel SDE output includes AVX vector instructions for integers, where the granularity can be as low as 1-bit per operand (cf. 4 or 8 byte per FP operand). Hence, the total integer operations count might be slightly inflated. Lastly, the results for HPCG show a big discrepancy between BDW and KNL/KNM. While the total FP operations count is similar, Intel’s optimized binary for KNL/KNM issues far more integer operations, see Table IV for details, and we are unaware of the reason.

### B. Floating-Point Operation/s and Time-to-Solution

Figure 2 shows the relative performance improvement of KNL/KNM over the dual-socket BDW node and the absolute achieved Gflop/s on each processor. It is important to note that all proxy-/mini-apps, with the exception of HPL, have less than 21.5% (BDW), 10.5% (KNL), and 15.1% (KNM) FP efficiency. Given that these applications are presumably optimized, and still achieve this low FP efficiency, implies a limited relevance of FP unit’s availability. The figure shows that the majority of codes have comparable performance on KNM versus KNL. Notable mentions are: a) CANDLE which benefits from VNNI units in mixed precision, b) MiFE, NekB, and XSBn which improve probably due to increased core count and KNM’s higher CPU frequency, and c) some




Fig. 2. **Relative floating-point performance** (FP32 and FP64 Gflop/s accumulated) of KNL and KNM in comparison to dual-socket Broadwell-EP (see **top plot**) and **Absolute achieved Gflop/s** w.r.t dominant FP operations (cf. Fig. 1) in comparison to theoretical peak performance listed in Tab. I (see **bottom plot**); Due to missing SDE data for CANDLE, we assume the total number of FP operations is equivalent to BDW and divide by CANDLE’s time-to-solution; Filtered proxy-apps with negligible FP operations: MxIO, MTri, and NGSA; Filtered out MiniAMR because of the strong-scaling issue described in Section III-A; Proxy-app abbreviations acc. to Section II-B




Fig. 3. Runtime speedup of KNL/KNM in comparison to dual-socket Broadwell-EP; MiniAMR included, but only KNL-to-KNM comparison valid due to differences in  $\#\{\text{MPI processes}\}$  and aforementioned strong-scaling issues (see Section III-A); Proxy-app abbreviations acc. to Section II-B

memory-bound applications (i.e., AMG, HPCG, and MTri) which get slower supposedly due to the difference in peak throughput demonstrated in Figure 4 in addition to the increased core count causing higher competition for bandwidth.

While we filtered out applications which do not perform a significant amount of FP operations in Figure 2, we added these applications in the time-to-solution comparison, shown in Figure 3, to gain a more comprehensive view. Overall, the speedups of KNL/KNM over our reference system match the expectation we reached from Figure 2 (top plot). However, one noticeable outlier is Laghos, which is caused by the application executing  $\approx 2x$  more FP64 operations on KNL/KNM, but also running about two times longer, and hence flop/s are roughly the same, while the time-to-solution differs from BDW.

### C. Memory Throughput of (MC-)DRAM

For the memory throughput measurements, shown in Figure 4, we use Intel’s PCM tool to analyze DRAM and MCDRAM throughput. Our measurements with BabelStream are included as well to demonstrate the maximum achievable






Fig. 4. Memory throughput (only DRAM for BDW, DRAM+MCDRAM for Phi) per proxy-app; Dotted lines indicate Triad stream bandwidth (flat mode, cf. Tab. I); BabelStream for 2 GiB (BABL2) and 14 GiB (BABL14) vector length added (measured in cache mode); Proxy-app labels acc. to Section II-B




Fig. 5. Roofline plot (w.r.t dominant FP operations and DRAM bandwidth) for Broadwell-EP reference system; Filtered proxy-apps with negligible FP operations: MxIO, MTri, and NGSa; Proxy-app labels acc. to Section II-B

bandwidth, see horizontal lines for MCDRAM (in *flat* mode), which is lower when the MCDRAM is used in *cache* mode. We still achieve 86% on KNL and 75% on KNM when the vectors fit into MCDRAM, but drop to slightly higher than DRAM throughput (due to minor prefetching benefits) when the vectors do not fit (see BABL14 for 14 GiB vectors). This throughput advantage of the MCDRAM translates into a performance boost for six proxy-apps (AMG, MAMR, MiFE, NekB, XSBn, and QCD; in comparison to BDW) which heavily utilize the available bandwidth, see Figure 4, and which are memory-bound on our reference system. This can easily be verified when comparing the time-to-solution for these kernels as show in Figure 3 and broken down into numbers in Table IV. Only HPCG cannot benefit from the higher bandwidth and, despite showing  $\approx 2x$  throughput, the runtime drops by more than 10%, indicating a memory-latency issue of HPCG on KNL/KNM, which is one of the design goals for the benchmark [34].

#### D. Roofline Model Analysis for Broadwell-EP System

Based on the collected data of Sections IV-B and IV-C, we calculate the location of each of our FP-intensive proxies (including BabelStream Triad) within the roofline graph for our x86-based reference system, as visualized in Figure 5.

These results match largely our expectations about the applications’ optimizations for x86-based architectures, as well as similar analysis performed for other HPC workloads and benchmarks [44], [45], [46]. The only noticeable outlier is Laghos, which leaves room for performance tuning, and hence challenges our initial assumption of Section III-A.

Given that almost all proxy-apps are in fact optimized for x86, and that both KNL and KNM are x86 ISA, likely all levels of (threads-, vector-, instruction-) parallelism are already exposed. The remaining question is how well the runtime and the compiler utilizes this parallelism. We covered both aspects by our approach (see Section III-A) of determining the best combination of MPI processes and OpenMP threads and instructing the compiler to optimize for the host architecture. Consequently, our roofline plots for KNL/KNM reveal similar information, and are therefore omitted.

#### E. Frequency Scaling to Identify Compute-Boundedness

For this test, we disable turbo boost and throttle the core frequency, but keep the uncore at maximum frequency which would otherwise negatively affect the memory subsystem, to identify each application’s dependency on ALU/FPU performance. The shown speedup (w.r.t time-to-solution) in Figure 6 for each proxy-app is relative to the lowest CPU frequency on each architecture, and we include our *performance* results (cf. Section III-A) with maximum frequency plus enabled turbo boost (labeled with “+TB”). It should be noted, that Intel abandoned single-frequency turbo boost long ago, and the real TB frequency band depends on multiple factors [47], such as #cores, utilized units, etc. Hence, we choose an universal, but pessimistic +100 Mhz (cf. Table I) for the TB plot in this figure, and therefore an application may exceed our pessimistic peak, as it is evident for Knights Landing in the top plot.

While a benefit from enabled turbo boost on BDW is near invisible (except for MTri), the proxy-apps clearly reduce their time-to-solution on KNL and KNM when these CPUs are allowed to turbo. Overall, the benchmarks seem to be less memory-bound and more compute-bound, especially salient for AMG and MiniFE, when moving to Xeon Phi, indicating a clear benefit from the much bigger/faster MCDRAM used as last-level cache and indicating a more balanced (w.r.t bandwidth to flop/s ratio) architecture. However, the limited speedup for HPL on KNL clearly shows the CPU’s abundance of FP64 units. Here, the successor, Knights Mill, shows a better balance. Another interesting observation is the inverse behavior of AMG and HPCG on our tested architecture. Both benchmarks are supposed to be memory-bound, but the absence of signs of any scalability with frequency on Xeon Phi strengthens our hypothesis from Section IV-C that HPCG is primarily bound by memory latency.

For I/O portions of an application, the Figure 6 reveals another observation, i.e., MACSio’s write speed scales with increased frequency. Since, MACSio performs only single figure Glop/s and negligible flop/s, increasing the CPU’s compute capabilities cannot explain the shown speedup. Hence, our theory is: MACSio (and I/O in general) is bound by the Linux kernel,





## D. Compute-bound Applications

Investing more in data-centric architectures to accommodate memory-bound applications can have a negative impact on the remaining minority of applications: compute-bound applications. Considering the market trends that are already pushing away from dedicating the majority of chip area to FP64 units, it is likely that libraries with compute-bound code (e.g., BLAS) would support mixed precision or emulation by lower precision FPUs. The remaining applications that do not rely on external libraries might suffer a performance hit.

## VI. RELATED WORK

Apart from RIKEN's mini-apps and the ECP proxy-apps, which we use for our study, there are numerous benchmark suites based on proxy applications from other HPC centers and institutes available [51], [52], [53], [54], [55], [56]. Overall those lists show a partial overlap, either directly (i.e., same benchmark) or indirectly (same scientific domain), between all these suites, which, for example, were used to analyze message passing characteristic [57] or to assess how predictable full application performance is based on proxy-app measurements [58]. Hence, our systematic approach and published framework <https://gitlab.com/domke/PAstudy> can be transferred to these alternative benchmarks for complementary studies, and our included raw data can be investigated further w.r.t metrics which were outside the scope of our study.

Furthermore, the HPC community has already started to analyze relevant workloads with respect to arithmetic intensity or memory and other potential bottlenecks for some proxy-apps [38], [44], [59] and individual applications [60], [61], [62], revealing similar results to ours that most realistic HPC codes are not compute-bound and achieve very low computational efficiency, which in demonstrated cases affected procurement decisions [50]. However, to the best of our knowledge, we are the first to present a broad study across a wide spectrum of HPC workloads which aims at characterizing bottlenecks and aims specifically at identifying floating-point unit/precision requirements for modern architectures.

## VII. CONCLUSION

We compared two architectural similar processors that have different double-precision silicon budget. By studying a large number of HPC proxy application, we found no significant performance difference between these two processors, despite one having more double-precision compute than the other. Our study points toward a growing need to re-iterate and re-think architecture design decisions in high-performance computing, especially with respect to precision. Do we really need the amount of double-precision compute that modern processors offer? Our results on the Intel Xeon Phi twins points towards a 'No', and we hope that this work inspires other researchers to also challenge the floating-point to silicon distribution for the available and future general-purpose processors, graphical processors, or accelerators in HPC systems.

## ACKNOWLEDGMENT & AUTHOR CONTRIBUTIONS

This work was supported by MEXT, JST special appointed survey 30593/2018, JST-CREST Grant Number JPMJCR1303, JSPS KAKENHI Grant Number JP16F16764, the New Energy and Industrial Technology Development Organization (NEDO), and the AIST/TokyoTech Real-world Big-Data Computation Open Innovation Laboratory (RWBC-OIL). Moreover, we thank Intel for their technical support. The authors K.M., J.D., H.Z., K.Y., T.T. and Y.T. performed the experiments and data collection. J.D., M.W., A.P. designed the study, analyzed the data, and supervised its execution together with S.M., while all authors contributed to writing and editing.

## REFERENCES

- [1] R. H. Dennard *et al.*, "Design of ion-implanted MOSFET's with very small physical dimensions," *IEEE Journal of Solid-State Circuits*, vol. 9, no. 5, pp. 256–268, 1974.
- [2] G. E. Moore, "Lithography and the Future of Moore's Law," in *Integrated Circuit Metrology, Inspection, and Process Control IX*, vol. 2439. International Society for Optics and Photonics, 1995, pp. 2–18.
- [3] T. Yoshida, "Fujitsu High Performance CPU for the Post-K Computer," 2018. URL: <http://www.fujitsu.com/jp/Images/20180821hotchips30.pdf>
- [4] J. Choquette *et al.*, "Volta: Performance and Programmability," *IEEE Micro*, vol. 38, no. 2, pp. 42–52, Mar. 2018.
- [5] J. Pu *et al.*, "FPMaX: a 106gflops/W at 217gflops/mm<sup>2</sup> Single-Precision FPU, and a 43.7 GFLOPS/W at 74.6 GFLOPS/mm<sup>2</sup> Double-Precision FPU, in 28nm UTBB FDSOI," 2016. URL: <http://arxiv.org/abs/1606.07852>
- [6] A. Haidar *et al.*, "Harnessing GPU's Tensor Cores Fast FP16 Arithmetic to Speedup Mixed-Precision Iterative Refinement Solvers," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '18, Dallas, Texas, Nov. 2018, accepted at SC '18.
- [7] A. Sodani *et al.*, "Knights Landing: Second-Generation Intel Xeon Phi Product," *IEEE Micro*, vol. 36, no. 2, pp. 34–46, Mar. 2016.
- [8] D. Bradford *et al.*, "KNIGHTS MILL: New Intel Processor for Machine Learning," 2017. URL: [https://www.hotchips.org/wp-content/uploads/hc\\_archives/hc29/Hc29.21-Monday-Pub/Hc29.21.40-Processors-Pub/Hc29.21.421-Knights-Mill-Bradford-Intel-APPROVED.pdf](https://www.hotchips.org/wp-content/uploads/hc_archives/hc29/Hc29.21-Monday-Pub/Hc29.21.40-Processors-Pub/Hc29.21.421-Knights-Mill-Bradford-Intel-APPROVED.pdf)
- [9] "ECP Proxy Apps Suite," 2018. URL: <https://proxyapps.exascaleproject.org/ecp-proxy-apps-suite/>
- [10] RIKEN AICS, "Fiber Miniapp Suite," 2015. URL: <https://fiber-miniapp.github.io/>
- [11] A. Heinecke *et al.*, "High Order Seismic Simulations on the Intel Xeon Phi Processor (Knights Landing)," in *International Conference on High Performance Computing*, ser. ISC '16. Springer, 2016, pp. 343–362.
- [12] N. A. Gawande *et al.*, "Scaling Deep Learning Workloads: NVIDIA DGX-1/Pascal and Intel Knights Landing," in *2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, May 2017, pp. 399–408.
- [13] J. Park *et al.*, "High-performance Algebraic Multigrid Solver Optimized for Multi-core Based Distributed Parallel Systems," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '15. Austin, TX, USA: ACM, 2015, pp. 54:1–54:12.
- [14] J. Wozniak *et al.*, "CANDLE/Supervisor: A Workflow Framework for Machine Learning Applied to Cancer Research," *BMC Bioinformatics*, 2018.
- [15] J. Mohd-Yusof *et al.*, "Co-design for molecular dynamics: An exascale proxy application," Los Alamos National Laboratory, Tech. Rep. LA-UR 13-20839, 2013. URL: [http://www.lanl.gov/orgs/adts/publications/science\\_highlights\\_2013/docs/Pg88\\_89.pdf](http://www.lanl.gov/orgs/adts/publications/science_highlights_2013/docs/Pg88_89.pdf)
- [16] V. Dobrev *et al.*, "High-Order Curvilinear Finite Element Methods for Lagrangian Hydrodynamics," *SIAM Journal on Scientific Computing*, vol. 34, no. 5, pp. B606–B641, 2012.
- [17] J. Dickson *et al.*, "Replicating HPC I/O Workloads with Proxy Applications," in *Proceedings of the 1st Joint International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems*, ser. PDSW-DISCS '16. Piscataway, NJ, USA: IEEE Press, 2016, pp. 13–18.

- [18] M. A. Heroux *et al.*, “Improving Performance via Mini-applications,” Sandia National Laboratories, Tech. Rep. SAND2009-5574, 2009.
- [19] M. M. Wolf *et al.*, “A task-based linear algebra Building Blocks approach for scalable graph analytics,” in *2015 IEEE High Performance Extreme Computing Conference (HPEC)*, Sep. 2015, pp. 1–6.
- [20] R. F. Boisvert *et al.*, “Matrix Market: A Web Resource for Test Matrix Collections,” in *Proceedings of the IFIP TC2/WG2.5 Working Conference on Quality of Numerical Software: Assessment and Enhancement*. London, UK, UK: Chapman & Hall, Ltd., 1997, pp. 125–137. URL: <http://dl.acm.org/citation.cfm?id=265834.265854>
- [21] Argonne National Laboratory, “NEK5000.” URL: <http://nek5000.mcs.anl.gov>
- [22] N. A. Petersson and B. Sjögreen, “User’s guide to SW4, version 2.0,” Lawrence Livermore National Laboratory, Tech. Rep. LLNL-SM-741439, 2017, (Source code available from [\tt geodynamics.org/cig](http://tt.geodynamics.org/cig)).
- [23] S. Habib *et al.*, “HACC: Extreme Scaling and Performance Across Diverse Architectures,” *Commun. ACM*, vol. 60, no. 1, pp. 97–104, Dec. 2016.
- [24] J. R. Tramm *et al.*, “XSbench - The Development and Verification of a Performance Abstraction for Monte Carlo Reactor Analysis,” in *PHYSOR 2014 - The Role of Reactor Physics toward a Sustainable Future*, Kyoto, 2014.
- [25] Y. GUO *et al.*, “Basic Features of the Fluid Dynamics Simulation Software “FrontFlow/Blue”,” *SEISAN KENKYU*, vol. 58, no. 1, pp. 11–15, 2006.
- [26] K. Ono *et al.*, “FFV-C package.” URL: [http://avr-aics-riken.github.io/ffvc\\_package/](http://avr-aics-riken.github.io/ffvc_package/)
- [27] Y. Andoh *et al.*, “MODYLAS: A Highly Parallelized General-Purpose Molecular Dynamics Simulation Program for Large-Scale Systems with Long-Range Forces Calculated by Fast Multipole Method (FMM) and Highly Scalable Fine-Grained New Parallel Processing Algorithms,” *Journal of Chemical Theory and Computation*, vol. 9, no. 7, pp. 3201–3209, 2013.
- [28] T. Misawa *et al.*, “mVMC—Open-source software for many-variable variational Monte Carlo method,” *Computer Physics Communications*, 2018. URL: <http://www.sciencedirect.com/science/article/pii/S0010465518303102>
- [29] H. Tomita and M. Satoh, “A new dynamical framework of nonhydrostatic global model using the icosahedral grid,” *Fluid Dynamics Research*, vol. 34, no. 6, pp. 357–400, 2004. URL: <http://stacks.iop.org/1873-7005/34/i=6/a=A03>
- [30] RIKEN CSR, “Grand Challenge Application Project for Life Science,” 2013. URL: [http://www.csrp.riken.jp/application\\_d\\_e.html#D2](http://www.csrp.riken.jp/application_d_e.html#D2)
- [31] T. Nakajima *et al.*, “NTChem: A High-Performance Software Package for Quantum Molecular Simulation,” *International Journal of Quantum Chemistry*, vol. 115, no. 5, pp. 349–359, Dec. 2014.
- [32] T. Boku *et al.*, “Multi-block/multi-core SSOR preconditioner for the QCD quark solver for K computer,” *Proceedings, 30th International Symposium on Lattice Field Theory (Lattice 2012): Cairns, Australia, June 24-29, 2012*, vol. LATTICE2012, p. 188, 2012.
- [33] J. Dongarra, “The LINPACK Benchmark: An Explanation,” in *Proceedings of the 1st International Conference on Supercomputing*. London, UK, UK: Springer-Verlag, 1988, pp. 456–474. URL: <http://dl.acm.org/citation.cfm?id=647970.742568>
- [34] J. Dongarra *et al.*, “A new metric for ranking high-performance computing systems,” *National Science Review*, vol. 3, no. 1, pp. 30–35, 2016.
- [35] E. Strohmaier *et al.*, “TOP500,” Jun. 2018. URL: <http://www.top500.org/>
- [36] T. Deakin *et al.*, “GPU-STREAM v2.0: Benchmarking the Achievable Memory Bandwidth of Many-Core Processors Across Diverse Parallel Programming Models,” in *High Performance Computing*, M. Tauber *et al.*, Eds. Cham: Springer International Publishing, 2016, pp. 489–507.
- [37] M. Hashimoto *et al.*, “An Empirical Study of Computation-Intensive Loops for Identifying and Classifying Loop Kernels: Full Research Paper,” in *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering*, ser. ICPE ’17. New York, NY, USA: ACM, 2017, pp. 361–372.
- [38] O. Aaziz *et al.*, “A Methodology for Characterizing the Correspondence Between Real and Proxy Applications,” in *2018 IEEE International Conference on Cluster Computing (CLUSTER)*, Belfast, UK, Sep. 2018.
- [39] T. Willhalm *et al.*, “Intel Performance Counter Monitor - A better way to measure CPU utilization,” Jan. 2017. URL: <https://software.intel.com/en-us/articles/intel-performance-counter-monitor>
- [40] K. Raman, “Calculating “FLOP” using Intel Software Development Emulator (Intel SDE),” Mar. 2015. URL: <https://software.intel.com/en-us/articles/calculating-flop-using-intel-software-development-emulator-intel-sde>
- [41] S. Sobhee, “Intel VTune Amplifier Release Notes and New Features,” Sep. 2018. URL: <https://software.intel.com/en-us/articles/intel-vtune-amplifier-release-notes>
- [42] J. Treibig *et al.*, “LIKWID: A lightweight performance-oriented tool suite for x86 multicore environments,” in *Proceedings of PSTI2010, the First International Workshop on Parallel Software Tools and Tool Infrastructures*, San Diego, CA, 2010.
- [43] S. Walker and M. McFadden, “Best Practices for Scalable Power Measurement and Control,” in *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, May 2016, pp. 1122–1131.
- [44] K. Asifuzzaman *et al.*, “Report on the HPC application bottlenecks,” ExaNoDe, Tech Report ExaNoDe Deliverable D2.5, 2017. URL: <http://exanode.eu/wp-content/uploads/2017/04/D2.5.pdf>
- [45] N. P. Jouppi *et al.*, “In-Datcenter Performance Analysis of a Tensor Processing Unit,” in *Proceedings of the 44th Annual International Symposium on Computer Architecture*, ser. ISCA ’17. New York, NY, USA: ACM, 2017, pp. 1–12.
- [46] G. Ofenbeck *et al.*, “Applying the Roofline Model,” in *2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Mar. 2014, pp. 76–85.
- [47] G. Lento, “Whitepaper: Optimizing Performance with Intel® Advanced Vector Extensions,” Intel Corporation, Tech. Rep., 2014. URL: <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/performance-xeon-e5-v3-advanced-vector-extensions-paper.pdf>
- [48] T. Guérout *et al.*, “Energy-aware simulation with DVFS,” *Simulation Modelling Practice and Theory*, vol. 39, pp. 76–91, 2013.
- [49] Japan Meteorological Agency (JMA), “JMA begins operation of its 10th-generation supercomputer system,” Jun. 2018. URL: [https://www.jma.go.jp/jma/en/News/JMA\\_Super\\_Computer\\_upgrade2018.html](https://www.jma.go.jp/jma/en/News/JMA_Super_Computer_upgrade2018.html)
- [50] S. Saini *et al.*, “Performance Evaluation of an Intel Haswell and Ivy Bridge-Based Supercomputer Using Scientific and Engineering Applications,” in *2016 IEEE 18th International Conference on High Performance Computing and Communications; IEEE 14th International Conference on Smart City; IEEE 2nd International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, 2016, pp. 1196–1203.
- [51] PRACE, “Unified European Applications Benchmark Suite,” Oct. 2016. URL: <http://www.prace-ri.eu/ueabs/>
- [52] “Mantevo Suite.” URL: <https://mantevo.org/packages/>
- [53] NERSC, “Characterization of the DOE Mini-apps.” URL: <https://portal.nersc.gov/project/CAL/designforward.htm>
- [54] LLNL, “LLNL ASC Proxy Apps.” URL: <https://computation.llnl.gov/projects/co-design/proxy-apps>
- [55] —, “CORAL Benchmark Codes.” URL: <https://asc.llnl.gov/CORAL-benchmarks/>
- [56] SPEC, “SPEC HPG: HPG Benchmark Suites.” URL: <https://www.spec.org/hpg/>
- [57] B. Klenk and H. Fröning, “An Overview of MPI Characteristics of Exascale Proxy Applications,” in *High Performance Computing: 32nd International Conference, ISC High Performance 2017*, ser. ISC ’17, Frankfurt, Germany, Jun. 2017, pp. 217–236.
- [58] R. F. Barrett *et al.*, “Assessing the role of mini-applications in predicting key performance characteristics of scientific and engineering applications,” *Journal of Parallel and Distributed Computing*, vol. 75, pp. 107–122, 2015.
- [59] T. Koskela *et al.*, “A Novel Multi-level Integrated Roofline Model Approach for Performance Characterization,” in *High Performance Computing: 33rd International Conference, ISC High Performance 2018*, ser. ISC ’18, Frankfurt, Germany, Jun. 2018, pp. 226–245.
- [60] M. Culpo, “Current Bottlenecks in the Scalability of OpenFOAM on Massively Parallel Clusters,” PRACE, Tech Report, Aug. 2012. URL: <https://doi.org/10.5281/zenodo.807482>
- [61] J. R. Tramm and A. R. Siegel, “Memory Bottlenecks and Memory Contention in Multi-Core Monte Carlo Transport Codes,” *Annals of Nuclear Energy*, vol. 82, pp. 195–202, 2015.
- [62] K. Kumahata *et al.*, “Kernel Performance Improvement for the FEM-based Fluid Analysis Code on the K Computer,” *Procedia Computer Science*, vol. 18, pp. 2496–2499, 2013.

TABLE IV

APPLICATION CONFIGURATION AND MEASURED METRICS; MISSING DATA FOR CANDLE DUE TO SDE CRASHES ON PHI; MEASUREMENTS INDICATE CANDLE/MKL-DNN IGNORES OPENMP SETTINGS AND TRIES TO UTILIZE FULL CHIP → LISTED IN ITALIC; LABEL EXPLANATION: t2sol = TIME-TO-SOLUTION (KERNEL), Gop (D | S | I) = GIGA OPERATIONS (FP64 | FP32 | INTEGER), SIMD1/cyc = SIMD INSTRUCTIONS PER CYCLE, FPAIP|R | W] = FP ARITHMETIC INSTRUCTIONS PER MEMORY [READ | WRITE], [B | M]Bd = [BACK-END | MEMORY] BOUND (SEE [41] FOR DETAILS), L2h = L2 CACHE HIT RATE, LLh = LAST LEVEL CACHE HIT RATE (L3 FOR BDW, MCDRAM FOR KNL/KNM), Gbra/s = GIGA BRANCHES/S; NOTE: SIMD1/cyc AND FPAIP\* AS WELL AS BBd AND MBd OCCUPY THE SAME COLUMNS DUE TO THEIR SIMILARITY AND SPACE CONSTRAINTS

KNL	#MPI	#OMP	t2sol [s]	#Gop (D)	#Gop (S)	#Gop (I)	Power [W]	#SIMD1/cyc	BBd [%]	L2h [%]	LLh [%]	Gbra/s
AMG	1	128	6.057	110.271	0	352.640	202.16	0.063	77.3	93	74.6	7.310
CANDLE	1	32	59.796	N/A	N/A	N/A	143.79	0.105	67.4	86	89.7	N/A
CoMD	32	8	3.199	161.691	14.842	3.476	189.24	0.077	81.0	85	99.4	11.551
Laghos	64	4	13.508	85.547	0.422	1055.977	143.84	0.021	23.7	98	99.7	7.839
MACSio	64	1	35.110	0.613	0.007	77.884	140.02	0.002	52.8	98	98.6	18.115
miniAMR	128	1	47.150	291.536	0.014	3358.569	153.44	0.009	75.9	71	97.6	5.023
miniFE	1	256	0.694	28.961	0	177.704	221.69	0.022	81.8	93	93.6	10.930
miniTri	1	128	8.630	0	0	118.261	131.49	0.001	81.9	66	99.5	4.531
Nekbone	128	1	3.290	410.361	0	23.371	221.48	0.050	76.5	87	97.6	6.125
SW4lite	64	4	1.686	145.938	0	0.761	214.57	0.096	80.6	95	98.4	2.218
SWFFT	128	1	1.235	12.688	0.005	42.509	174.12	0.029	76.7	83	98.6	20.905
XSBench	1	256	1.290	27.283	0.417	16.441	192.46	0.041	93.7	22	99.5	2.629
FFB	64	2	8.244	2.300	258.561	1785.716	179.55	0.159	38.4	89	99.7	2.717
FFVC	1	64	13.009	134.589	1579.917	20174.483	180.58	0.169	36.0	95	99.7	4.415
mVMC	32	6	20.679	1141.865	1.345	1746.001	180.98	0.036	81.9	91	98.9	6.073
MODYLAS	64	4	22.514	6287.279	2.063	23104.728	206.98	0.072	80.4	97	95.7	7.742
NGSA	4	32	829.675	0.826	0.023	69.117	97.91	0.002	51.9	71	95.9	1.050
NICAM	10	15	37.802	422.504	0.066	925.228	119.46	0.193	67.8	92	99.2	0.231
NTChem	16	8	18.985	1629.210	0.627	2303.804	167.13	0.060	64.4	91	99.2	5.429
QCD	1	128	8.437	631.522	0	3823.335	215.67	0.220	69.4	88	95.4	1.151
HPCG	96	1	44.612	612.799	0	17530.136	181.69	0.023	86.1	91	45.7	1.446
HPL	64	1	145.400	184191.774	0.015	20226.567	221.13	0.374	52.3	93	87.9	1.232
KNM	#MPI	#OMP	t2sol [s]	#Gop (D)	#Gop (S)	#Gop (I)	Power [W]	#SIMD1/cyc	BBd [%]	L2h [%]	LLh [%]	Gbra/s
AMG	1	128	7.434	110.271	0	352.639	202.52	0.062	75.4	94	73.3	6.392
CANDLE	1	144	50.527	N/A	N/A	N/A	153.69	0.040	82.4	92	90.9	N/A
CoMD	72	2	3.194	161.842	14.880	3.479	196.64	0.177	67.5	86	99.1	11.546
Laghos	64	4	12.725	85.383	0.422	1056.141	139.33	0.023	25.1	98	99.8	8.345
MACSio	64	1	33.236	0.613	0.007	77.884	135.48	0.002	53.8	98	98.2	19.206
miniAMR	128	1	44.653	291.536	0.014	3358.570	177.31	0.009	75.3	71	97.3	5.337
miniFE	72	1	0.669	32.892	0	669.371	210.18	0.097	55.6	60	98.3	7.393
miniTri	1	128	9.545	0	0	118.262	122.02	0	80.9	68	99.6	4.102
Nekbone	144	1	2.984	410.381	0	23.470	233.46	0.040	76.1	87	96.6	6.494
SW4lite	72	4	1.569	146.048	0	0.764	228.01	0.090	81.3	96	97.8	2.753
SWFFT	128	1	1.189	12.555	0.005	41.732	172.66	0.026	77.2	83	98.5	21.990
XSBench	1	288	1.220	30.603	0.417	16.440	197.16	0.038	91.5	22	98.5	2.783
FFB	64	2	7.750	2.300	258.565	1785.712	178.72	0.171	38.6	89	99.7	2.886
FFVC	1	72	13.497	134.589	1579.917	20174.587	182.05	0.162	55.2	94	99.9	5.055
mVMC	72	4	19.659	1140.670	1.347	1802.663	197.64	0.012	76.0	91	98.5	8.869
MODYLAS	64	4	24.026	6287.279	2.063	23104.728	217.47	0.062	80.0	97	95.6	7.153
NGSA	4	18	724.546	0.826	0.023	69.300	88.67	0.002	39.5	68	94.9	1.138
NICAM	10	7	34.380	422.504	0.066	925.229	113.88	0.208	68.2	92	99.1	0.248
NTChem	72	2	14.606	1575.310	0.623	1985.255	176.51	0.066	59.0	90	98.4	7.038
QCD	1	144	9.662	631.522	0	3823.337	200.86	0.175	72.6	88	95.9	2.121
HPCG	64	1	42.865	612.605	0	17532.326	174.58	0.041	86.5	95	42.9	2.878
HPL	72	1	146.562	184893.073	0.016	20414.548	263.59	0.351	57.0	92	87.0	1.885
BDW	#MPI	#OMP	t2sol [s]	#Gop (D)	#Gop (S)	#Gop (I)	Power [W]	FPAIP R : W]	MBd [%]	L2h [%]	LLh [%]	Gbra/s
AMG	8	6	10.780	110.810	0	362.209	152.21	0.361 : 5.516	44.8	21	17	4.354
CANDLE	1	12	78.240	0.012	6918.340	2783.532	132.38	1.078 : 2.800	26.7	23	11	1.242
CoMD	48	1	2.921	152.022	0	0.205	133.17	0.845 : 6.615	1.5	15	15	11.391
Laghos	24	1	5.5472	44.534	0	421.465	126.51	0.184 : 0.476	13.2	81	56	16.808
MACSio	4	1	10.498	0.070	0	72.582	89.3	0 : 0	0.8	48	59	3.274
miniAMR	96	1	55.386	40.816	0	172.317	133.29	0.059 : 0.311	55.1	24	23	4.013
miniFE	24	1	1.475	30.693	0	120.715	152.77	0.311 : 5.454	55.2	15	12	4.699
miniTri	1	48	5.478	0	0	118.178	112.61	0 : 0	34.0	47	90	7.106
Nekbone	96	1	5.671	301.559	0	10.139	154.74	0.593 : 2.431	36.9	36	24	3.915
SW4lite	24	2	2.056	136.835	0	1.585	146.65	1.044 : 4.580	9.1	75	18	1.112
SWFFT	32	1	1.088	12.239	0	38.782	134.55	0.117 : 0.675	28.3	23	32	20.932
XSBench	1	96	2.022	19.921	0	20.280	132.25	0.807 : 3.847	71.7	5	18	1.653
FFB	24	1	5.327	1.300	233.640	2116.421	144.35	0.635 : 2.200	21.3	79	33	3.723
FFVC	12	4	12.691	127.322	1573.782	27857.376	151.85	0.481 : 2.844	3.3	84	57	9.045
mVMC	24	2	13.489	1092.394	0	2224.092	152.28	0.601 : 2.456	12.0	36	24	10.170
MODYLAS	16	3	36.101	5363.366	0	10888.745	135.75	0.875 : 8.736	8.1	60	31	5.385
NGSA	12	4	105.879	0.826	0.023	64.249	107.15	0.002 : 0.006	6.5	21	36	8.566
NICAM	10	6	28.449	428.282	0.003	687.852	118.32	0.540 : 3.732	49.6	27	19	0.585
NTChem	24	1	8.963	1315.509	0	778.829	141.3	0.867 : 4.931	9.4	56	39	10.173
QCD	1	24	13.102	612.303	0	3817.944	153.2	1.152 : 4.542	45.2	27	24	0.368
HPCG	2	24	38.595	559.046	0	90.171	166.18	0.143 : 0.628	11.3	34	23	10.928
HPL	24	1	271.794	181484.240	0	31919.479	189.37	2.280 : 122.693	3.9	10	3	2.147