

Performance Evaluation of Memory-Centric ARMv8 Many-Core Architectures: A Case Study with Phytium 2000+

Jian-Bin Fang, *Member, CCF*, Xiang-Ke Liao, *Fellow, CCF*, Chun Huang, and De-Zun Dong*

College of Computer, National University of Defense Technology, Changsha 410073, China

E-mail: {j.fang, xkliao, chunhuang, dong}@nudt.edu.cn

Received June 24, 2020; accepted December 9, 2020.

Abstract This article presents a comprehensive performance evaluation of Phytium 2000+, an ARMv8-based 64-core architecture. We focus on the cache and memory subsystems, analyzing the characteristics that impact the high-performance computing applications. We provide insights into the memory-relevant performance behaviours of the Phytium 2000+ system through micro-benchmarking. With the help of the well-known roofline model, we analyze the Phytium 2000+ system, taking both memory accesses and computations into account. Based on the knowledge gained from these micro-benchmarks, we evaluate two applications and use them to assess the capabilities of the Phytium 2000+ system. The results show that the ARMv8-based many-core system is capable of delivering high performance for a wide range of scientific kernels.

Keywords many-core architecture, memory-centric design, performance evaluation

1 Introduction

The high-performance computing (HPC) hardware is firmly moving towards the many-core design, and the ARMv8-based processors are emerging as an interesting alternative building block for HPC systems^[1–3]. This can be seen from the fact that the 64-core Phytium 2000+ architecture has been used to build the prototype of China’s new-generation exascale supercomputer^[4], as well as the fact that the Fugaku supercomputer has been built upon a 48-core A64FX architecture^[5]. Thus, it is important to understand such novel microarchitecture designs and their performance impacts on typical HPC applications. Having such knowledge is useful not only for better utilizing the computation resources, but also for justifying a further increase in the processor core provision and driving the innovations in memory system design.

In this article, we present a comprehensive performance evaluation of the Phytium 2000+ many-core architecture for HPC applications. We focus on the characteristics of Phytium 2000+ that have a direct per-

formance impact on these applications — high memory capacity and bandwidth, the unique cache organization, the mesh-based on-chip interconnect, and a large number of available hardware cores.

We first describe and highlight the new design features of the Phytium 2000+ processor (Section 2). Most notably, Phytium 2000+ integrates 64 ARMv8-based hardware cores and uses 128 GB of memory, which can deliver 563.2 GFLOP/s of double-precision performance and 153.6 GB/s of memory throughput when running at 2.2 GHz. This results in a good machine balance between calculations and memory accesses. Besides, Phytium 2000+ leverages a hierarchical heterogeneous on-chip network and cache organization distributed among eight panels, which helps to achieve a good data locality and scalability^[3].

We then use a set of microbenchmarks to characterize the low-level memory-relevant features of the Phytium 2000+ architecture (Section 3). We focus on the memory and cache subsystem in particular, including the core-to-core communication performance and the aggregated memory bandwidth with all hardware

Regular Paper

Special Section on Memory-Centric System Research for High-Performance Computing

This work is partially funded by the National Key Research and Development Program of China under Grant No. 2018YFB0204301, and the National Natural Science Foundation of China under Grant Nos. 61972408 and 61602501.

*Corresponding Author

©Institute of Computing Technology, Chinese Academy of Sciences 2021

cores. We observe that a hierarchy of data locality has to be dealt with so as to achieve high performance. Special care has to be taken as to whether the data is located in the local cache, in the same core group, in the same panel, or across panels. Our results show that the measured `write` bandwidth can reach over 80% of the theoretical peak DRAM bandwidth, which is surprisingly larger than the `read` bandwidth. This is achieved by using all available hardware cores, and pinning each thread to a distinct hardware core. We also demonstrate that the non-contiguous memory access is detrimental to the bandwidth efficiency, with `Phytium 2000+` showing more restrictions on the stanza length of data prefetching than the conventional `x86` processors.

In addition to our low-level evaluation on the memory subsystem, we analyze a `Phytium 2000+` many-core architecture with the well-known roofline model (Section 4). Thanks to its balance between calculations and memory accesses, the `Phytium 2000+` system can deliver promising performance for typical high-performance computing kernels, including sparse matrix-vector multiplication (SpMV), 3D stencil, and general matrix-matrix multiplication (GEMM).

We finally evaluate the `Phytium 2000+` system using two typical high-performance computing kernels (Section 5). They are SpMV and GEMM, which are two essential linear algebra kernels and widely used in high-performance computing applications. We show that for both applications, the `Phytium 2000+` system achieves excellent performance, comparable to

the start-of-the-art results. Our work demonstrates that the `Phytium 2000+` system can be used to efficiently handle much larger datasets, compared with the accelerator-centric design.

2 Overview of the Phytium 2000+ Architecture

Fig.1 gives a high-level view of the `Phytium 2000+` processor. It uses the `Mars II` architecture^①, and features 64 high-performance ARMv8 compatible `xiaomi` cores running at 2.2 GHz. The entire chip offers a peak performance of 563.2 Gflops for double-precision operations, with a maximum power consumption of 96 watts. The 64 hardware cores are organized into eight `panels`, where each panel connects a memory control unit.

The panel architecture of `Phytium 2000+` is shown in Fig.1(b). Each panel has eight `xiaomi` cores, and each core has a private L1 cache of 32 KB for data and instructions, respectively. Every four cores form a `core group` and share a 2 MB L2 cache. The L2 cache of `Phytium 2000+` uses an inclusive policy, i.e., the cachelines stored in L1 are also present in the L2 cache.

Each panel contains two directory control units (DCUs) and one `routing cell`. The DCUs on each panel act as dictionary nodes of the entire on-chip network. `Mars II` uses a hierarchical on-chip network, with a local interconnect on each panel and a global connect for the entire chip. The former couples cores

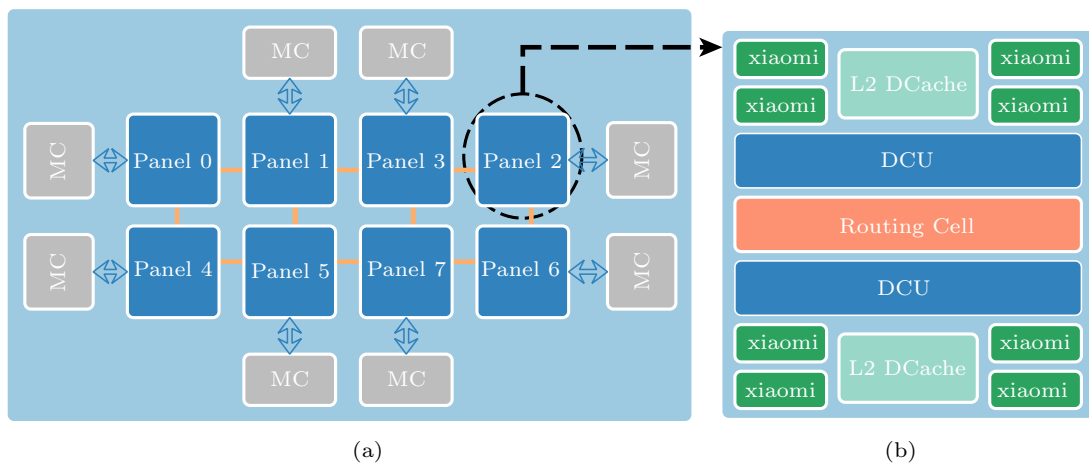


Fig.1. A high-level view of the `Phytium 2000+` architecture. (a) The 64 processor cores are grouped into eight panels, where (b) each panel contains eight ARMv8-based `xiaomi` cores.

^①Phytium Mars II Microarchitectures. <https://en.wikichip.org/wiki/phytium/microarchitectures/mars.ii>, Nov. 2020.

and L2 cache slices as a local cluster, and the latter is implemented with a configurable cell-network to connect panels. Phytium 2000+ uses a home-grown Hawk cache coherency protocol to implement a distributed directory-based global cache coherency across panels. The connected DDR memory modules are working at 2400 MHz, giving a theoretical bandwidth of 153.6 GB/s.

We run a customized Linux OS based on version 4.4 on the Phytium 2000+ system. We use gcc v8.2.0 compiler and the OpenMP/POSIX threading model.

3 MicroBenchmark Results

This section presents the core-to-core communication performance numbers, and the aggregated memory bandwidth for the Phytium 2000+ memory subsystem.

3.1 Core-to-Core Latency Results

3.1.1 Latency Overview

The accessing latency is referred to as the time of moving a cacheline within the local core or between two distinct cores. We use Molka’s approach to measure the core-to-core latency numbers [6]. During the measurement, we use multiple threads to move data between cores. To ensure that the buffer allocated by a thread belongs to a fixed core, we pin each thread to a fixed core, i.e., thread n always runs on core n (C n). Fig.2 shows the latency results when C0 loads the data from its local cache, from C1 sharing an L2 cache slide with C0, from C4 on the same panel, and from C8 on a different panel. Note that the performance numbers are measured when the cachelines are modified initially.

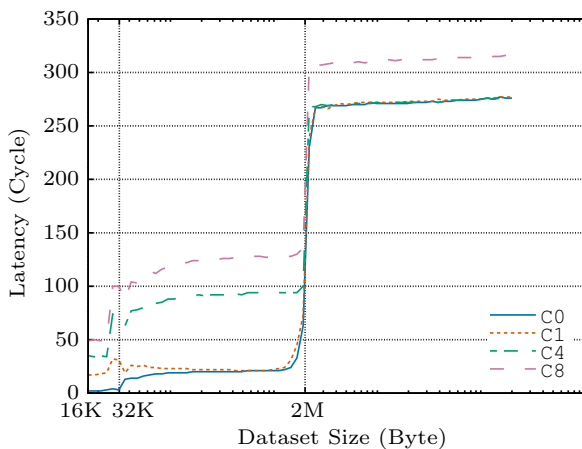


Fig.2. Read latency of C0 accessing the local (C0) cache or the cache of another core (C1, C4 or C8).

We see that, accessing the local L1 and L2 cache takes 3 cycles (1.4 ns) and 21 cycles (9.5 ns), respectively. The specification of the first generation Mars describes that accessing the local L1 and L2 takes 2 ns and 8 ns, respectively, which is in accordance with our measured numbers [3]. When C0 loads data from C1, the latency is the same as that accessing the local L2 cache. Fig.2 shows that, no matter which memory layer the data is suited in, loading cachelines across core groups or panels takes many more cycles than accessing the local cache slices. Thus, loading the data within the local cache slice is the fastest.

3.1.2 Across-Panel Latency Results

We evaluate the performance impact of panel distance on latency when accessing cores fixed to different panels. Fig.3 shows the latency results when C0 is accessing the cores on P1 (panel 1)–P7 (panel 7), respectively.

We see that the latency numbers vary over the panel distance, with a latency variance of up to 105 cycles. Besides, the latency numbers of C0 on P0 accessing C8 on P1 and C32 on P4 are the same. This is because P1 and P4 are at the same distance to P0. This result also agrees with our measured NUMA-aware stream bandwidth (Fig.6) and the theoretical latency results [3].

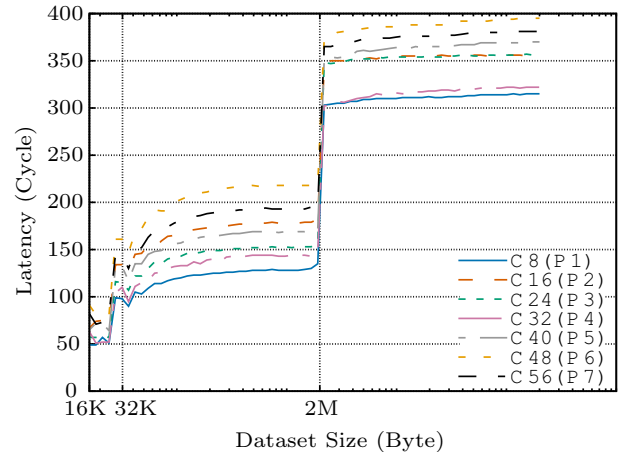


Fig.3. Read latency of C0 accessing the cache on the cores of different panels (P1–P7).

3.2 Core-to-Core Bandwidth

This subsection presents the read bandwidth on the Phytium 2000+ architecture. We measure the core-to-core sustainable bandwidth by continuously accessing a chunk of data elements [6]. Fig.4 shows the bandwidth of C0 loading cachelines which are modified, or shared

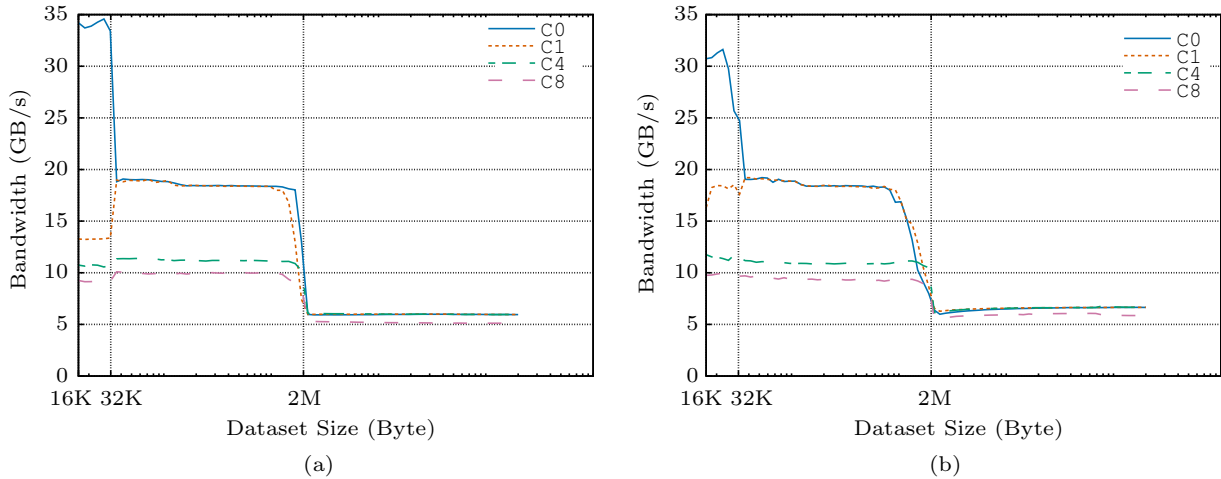


Fig.4. Read bandwidth of C0 accessing the local or another core (C1, C4 or C8). (a) modified. (b) shared.

in different cores and different cache levels. Note that, the **exclusive** and **modified** states are the same for the directory-based caching architectures. We measure the bandwidth of C0 loading data from its local cache, from C1 sharing an L2 cache with C0, from C4 on the same panel, and from C8 on a different panel.

In Fig.4, we find that the read bandwidth results show a clear phase change as the size of the dataset increases. Moreover, the size of the dataset when the staged change occurs is basically consistent with the size of various levels of the cache. Compared with the first change point occurring exactly at 32 KB (the size of L1 cache), the second change occurs earlier than 2 M (the size of L2 cache). This is because the L1 cache is a pure data cache, while the L2 cache is a hybrid cache for both data and instructions.

3.2.1 Local Cache Accesses

Whatever the state of the cachelines is, the data can be loaded from C0's local caches. The obtained bandwidth has nothing to do with the coherency state of the accessed data. The read bandwidth to its local L1 cache can reach 33.6 GB/s, while reading data from the local L2 cache can reach a bandwidth of 18.5 GB/s. Given that the L1 read port of Phytium 2000+ is 128 bits in width and runs at 2.2 GHz, we calculate the theoretical L1 read bandwidth as $2.2 \times 128 \div 8 = 35.2$ GB/s. We see that the measured bandwidth is quite close to its theoretical counterpart (33.6 GB/s vs 35.2 GB/s). The measured write bandwidth stays about 17.4 GB/s for L1. We note that the write bandwidth is around a half of the read bandwidth. This is because storing data into L1 occurs at 64 bits per cycle, while loading data from L1 occurs at 128 bits per cycle.

3.2.2 Within a Core Group

Differing from Intel's MESIF cache coherence protocol, Phytium 2000+ uses a home-grown Hawk cache coherence protocol to implement a distributed directory-based global cache coherency across panels. Note that the directory-based protocol cannot distinguish a data block cached in an **exclusive** or **modified** state.

Given that C1 and C0 share the same L2 cache slice, data can be loaded from the local L2 cache when the cacheline is **shared**. And the memory bandwidth of accessing the local L2 can reach 18.5 GB/s. But the bandwidth is reduced to be around 13.3 GB/s when C0 loading **exclusive** or **modified** cachelines is suited in C1's L1 cache. This is a notable difference from the x86 processor that the **exclusive** and **modified** states are treated the same on Phytium 2000+.

3.2.3 Within a Panel

When C0 loads the data from C4 of the same panel, where the two cores share no common cache slices, the bandwidth will be limited by the cross-group links. As can be seen from Fig.4, the bandwidth is significantly smaller (by around 40%) than the case when sharing the same L2 cache slice.

Similar to C1, when performing cross-group access to C4 for **exclusive** or **modified** cachelines, the bandwidth for reading the remote L1 cache is always smaller than that for accessing the remote L2. This is also because the data can be obtained directly from the L2 cache only when its state is **shared** initially.

3.2.4 Across Panel Accesses

C8 does not share a common L2 cache slice with C0,

and the two cores have to communicate via the cross-panel routing cells. The read bandwidth of C0 accessing C8 ranges from 9.2 GB/s to 9.7 GB/s, which is smaller than the bandwidth of accessing C1 or C4 within the same panel with C0.

3.2.5 NUMA Memory Accesses

Since C0, C1, C4 are within the same panel, they are connected directly to the same MCU and memory module. When accessing the data in the local memory module for C1 and C4, the bandwidth can reach around 6 GB/s. On the other hand, C8 is connected directly to another memory module. The bandwidth of C0 loading data from C8's memory module is around 5.1 GB/s.

To summarize, there is another difference between the Phytium 2000+ processor and the x86 processor when accessing the shared cachelines. The x86 processor uses an extension of the MESIF protocol, which requires the data to be fetched from the core with the latest copy (*forward*). Meanwhile, the Phytium 2000+ processor uses a MOSEI-like coherency protocol. There is no need to find the *forward* copy, but it can directly obtain the data with an arbitrary *shared* copy.

3.3 Overall Bandwidth

3.3.1 Aggregated Stream Bandwidth

We use the `stream v5.9` benchmark suite (`copy`, `scale`, `add`, `triad`) to measure the aggregated memory bandwidth^[7], where the array size is 200 000 000 (i.e., 4577.6 MB memory required) and each test is run 20 times. Also, we use separate benchmarks to measure the memory bandwidth for pure `read` and `write` operations. The `read` benchmark reads the data from an array A ($b = b + A[k]$). The `write` benchmark writes a constant value into an array A ($A[k] = C$). Note that A needs to be large enough (e.g., 1 GB) such that it cannot fit in the on-chip memory. To avoid the impact of “cold” TLBs, we start with two “warm-up” iterations of the benchmarks, before we measure a third one. We use different numbers of running threads — from 1 to 64. Note that each thread is pinned to a hardware core in a sequential order with `GOMP_CPU_AFFINITY`.

The aggregated memory bandwidth of `read`, `write`, `copy`, `scale`, `add` and `triad` is shown in Fig.5, when using different numbers of cores. We see that the overall memory bandwidth increases over the number of used cores. It happens because when using more threads, we can generate more requests to memory controllers, thus making the memory channels busier. Thus, if aiming

to achieve high memory bandwidth, programmers need to launch enough threads to saturate the interconnect and the memory channels. We also see that the maximum bandwidth for the six benchmarks is far below the theoretical peak of 153.6 GB/s. Specifically, the `read` bandwidth peaks at 100.42 GB/s and the `write` bandwidth peaks at 123.43 GB/s, achieved with 64 threads by pinning each thread to a fixed core.

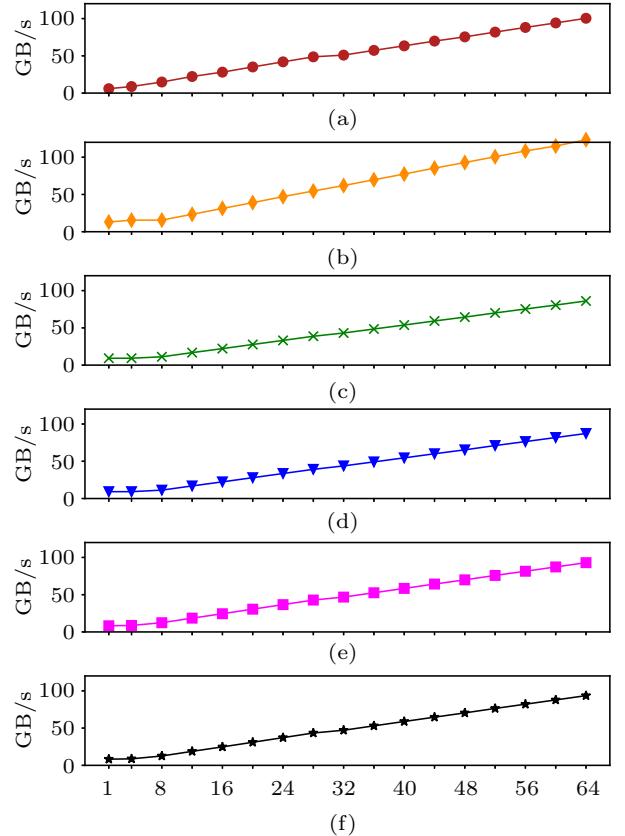


Fig.5. Memory bandwidth with the `stream` benchmark. (a) `read`. (b) `write`. (c) `copy`. (d) `scale`. (e) `add`. (f) `triad`.

Fig.5 shows that the achieved memory bandwidth increases very slightly, when using 1, 4, or 8 cores. This is because using less than 8 cores will load the data from the same memory channel and module. In such a case, the maximum `triad` bandwidth stays around 12.57 GB/s, which is 65.47% of the theoretical peak bandwidth. When using 64 cores, the achieved memory bandwidth reaches their maximum, with a bandwidth of 86.13 GB/s for `copy`, 87.27 GB/s for `scale`, 93.00 GB/s for `add`, and 93.57 GB/s for `triad`.

3.3.2 NUMA-Aware Stream Bandwidth

We measure the NUMA-aware `stream` bandwidth on the Phytium 2000+ processor with the `libnuma` li-

brary. This is achieved by running a Pthread version of the `stream` v5.9 benchmark^[7] and pinning 64 threads to eight NUMA nodes. Fig.6 shows the triad bandwidth of all pairs of NUMA nodes. We see that the bandwidth results obtained on the local NUMA node are the largest. The maximum bandwidth within the same nodes is 14.46 GB/s. With the help of the `libnuma` library, the aggregated triad memory bandwidth reaches around 103.96 GB/s (with 64 threads), which is around 10% larger than the measured number in Subsection 3.3.1. On the other hand, the across-node bandwidth numbers are noticeably much smaller. The cross-border accessing bandwidth can be reduced to only 10.16 GB/s on the Phytium 2000+ processor.

0	14.46	12.16	12.22	12.37	12.26	11.99	11.12	11.14
1	12.97	14.42	11.88	13.28	11.35	12.61	11.47	11.41
2	11.55	11.60	13.70	12.83	10.57	10.85	12.52	11.88
3	12.09	11.96	12.63	14.00	11.18	12.55	11.30	12.57
4	12.73	11.35	11.44	11.18	13.84	12.44	11.93	11.78
5	10.77	11.53	11.49	11.50	13.06	13.55	12.78	12.61
6	10.16	11.27	12.45	11.54	11.13	12.04	13.87	12.78
7	11.50	12.09	12.10	11.91	13.14	12.53	12.77	13.88
	0	1	2	3	4	5	6	7

Fig.6. NUMA-aware memory bandwidth with the `stream` benchmark (the ID of x -axis and y -axis denotes the ID of NUMA nodes, and the numbers on the heatmap represent the achieved memory bandwidth in GB/s).

3.3.3 Prefetching Effect

To evaluate the prefetching efficiency of Phytium 2000+, we use the Stanza Triad (STriad)^[8] benchmark with a single thread. STriad works by performing a DAXPY (Triad) inner loop for a length L stanza, then jumps over k elements, and continues with the next L elements, until reaching the end of the array. We set the total array size to 1 GB, and set k (the jump length) to 2048 which is large enough to ensure no prefetch between stanzas, but small enough to avoid penalties from TLB misses and DDR precharge^[8]. For each stanza, we run the experiment 10 times, with the L2 cache flushed each time, and calculate the median value of the 10 runs to get the memory bandwidth for each stanza length.

Fig.7 shows the results of the STriad experiments on both Phytium 2000+ and a regular Xeon processor (Intel Xeon Gold 6130). We see an increase in memory bandwidth over stanza length L , and it eventually

approaches a peak of 8.4 GB/s on Phytium 2000+ and 12.8 GB/s on Xeon. We conclude that the prefetching mechanism on both system does offer a significant benefit. Further, we see the transition point from the bandwidth-increasing state to the bandwidth-stable state appears at the same point when $L = 2^{10}$, but the Xeon processor can reach a much larger bandwidth. Therefore, we conclude that the non-contiguous access to memory is detrimental to memory bandwidth efficiency, with Phytium 2000+ showing more restrictions on the stanza length of prefetching than the regular Xeons. To comply with this restriction, programmers have to create the longest possible stanzas of contiguous memory accesses, improving memory bandwidth.

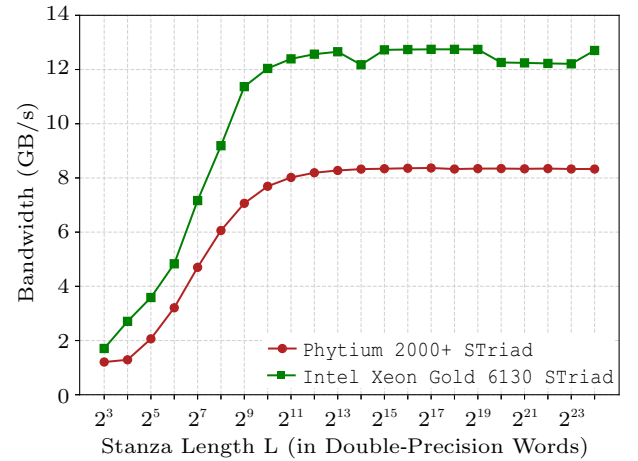


Fig.7. Performance of STriad on Phytium 2000+ (the x -axis is in log scale and the results on Intel Xeon Gold 6130 are normalized to those on Phytium 2000+).

4 Roofline Model Analysis

In this section, we use the well-known roofline model to estimate the Phytium 2000+'s performance for various scientific kernels. The roofline model is a visually-intuitive and throughput-oriented approach of characterizing a system's performance for various arithmetic intensities (floating-point operations per byte of DRAM traffic) of algorithms^[9]. This visualization and mapping of performance to algorithms helps to identify and quantify the primary factors that limit the performance of a given application. Such type of "bound and bottleneck analysis" has the advantage of being more intuitive and user-friendly than traditional analytical performance models.

Fig.8 shows the roofline model specialized for the Phytium 2000+ system. The system's double-precision performance is 563.2 GFLOP/s and its peak DDR4

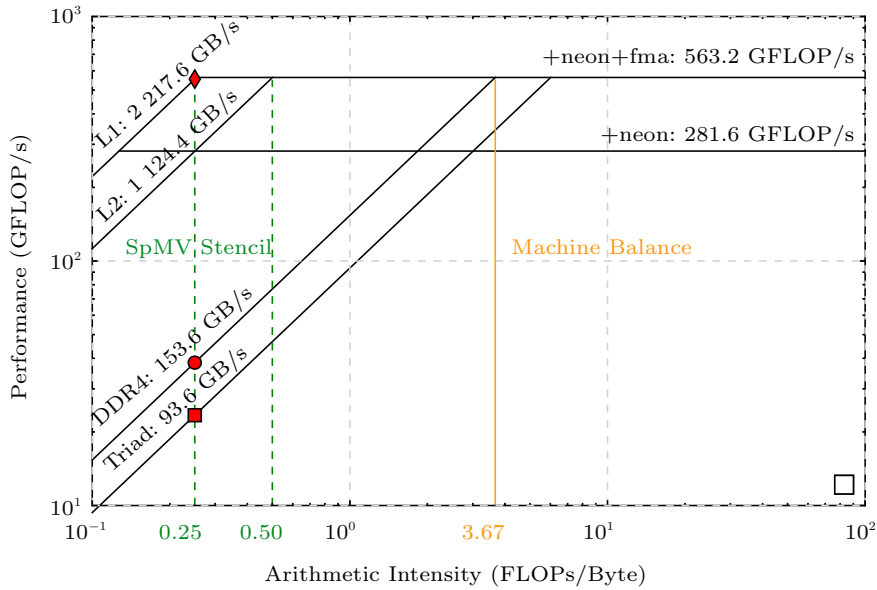


Fig.8. Roofline for the Phytium 2000+ processor.

memory bandwidth is 153.6 GB/s. This gives a machine balance of 3.67 (denoted by the orange vertical line), which is calculated as the ratio of peak compute throughput to memory throughput. This ridge point, where the two lines meet, is the arithmetic intensity at which an algorithm goes from being memory bound to being compute bound on that system. To derive more practical performance numbers, we also use the sustainable triad memory bandwidth (93.6 GB/s).

Phytium 2000+ can be represented as a set of cores with two-level on-chip caches connected to a DRAM memory. And using the original roofline modeling approach is not sufficient to fully capture the performance of Phytium 2000+ which relies on the on-chip memory hierarchy. Ilic *et al.* further proposed a cache-aware roofline model^[10], based on a core-centric concept where FP operations, data traffic and memory bandwidth at different levels are perceived. Based on the microbenchmarking results on the L1 and L2 caches (Section 3), we add two more cache-relevant rooflines to capture the memory accessing capabilities.

Using the roofline model, we can conduct a performance analysis for a variety of scientific kernels. As shown in Fig.8, for three kernels that are frequently used in scientific applications—sparse matrix-vector multiplication (SpMV), 3D stencil (Stencil), and general matrix-matrix multiplication (GEMM)—we can estimate a parallel performance upper-bound on the Phytium 2000+ system. Note that, the arithmetic intensity of GEMM depends on the matrix size, and we

do not show it in the figure. For example, for the SpMV kernel, whose operational intensity is around 0.25, we can expect a peak performance of 38.4 Gflop/s (red circle) with 64 threads or cores. The above analysis, however, assumes the optimal memory bandwidth use of Phytium 2000+. Instead, if we use the sustainable triad memory bandwidth, then the expected performance would be 23.4 Gflop/s (red square) instead. If the SpMV dataset is sufficiently small to be within the L1 cache (i.e., smaller than 64×32 KB), then the expected performance would be 554.4 Gflop/s (red diamond) on the Phytium 2000+ system.

Our results show that the Phytium 2000+ system is optimized for data intensive applications with low operational intensities. Phytium 2000+ has a system balance of 3.67, which suggests that it is well balanced. This is different from many other systems whose system balance typically ranges from 6 to 7 and which prefer increasing the number of processing units rather than the memory bandwidth^[11]. By taking the difficult, but effective, strategy of keeping a “balance” between compute and memory throughputs, the Phytium 2000+ system is capable of delivering excellent performance for a wide range of scientific kernels.

5 Applications

5.1 GEMM

General matrix-matrix multiplication (GEMM) is a common algorithm in linear algebra, machine learn-

ing, and many other domains. Given that large-scale GEMM and small-scale GEMM are used in different scenarios, this subsection demonstrates how well they perform on the Phytium 2000+ architecture. We use four open-source GEMM libraries (OpenBLAS, BLIS [12], BLASFEO [13], and Eigen) for the evaluation.

5.1.1 Small-Scale GEMM

Fig.9(a) shows the performance results of small-scale SGEMM on the Phytium 2000+ system. In our experiments, square matrices are used with their sizes ranging from 5 to 200, with a step of 5. Since the size of the matrix is small, we use a single thread to evaluate the performance of each BLAS library. The execution time is calculated as the average of 10 runs.

BLASFEO stands out as the best performer in nearly all the input sizes tested, reaching a maximum thread efficiency of 94%. This is because the matrix represented in BLASFEO is stored in a panel-major format, which has no data packing overhead, while the other libraries are stored in either row-major or column-major format. When the matrix dimension in BLASFEO is a multiple of 4, there occurs a ridge point for the GEMM performance, which is related to the index computation of elements in the panel-major format [13]. When the dimension of the matrix is not a multiple of 4, BLASFEO has to pad the data structure with zeros, resulting in an overhead.

OpenBLAS, BLIS and Eigen all suffer from the packing overhead. We see that OpenBLAS achieves better performance than the other two. This is due to the fact that it employs hand-crafted assembly-coded

GEMM kernels. Eigen employs C++-coded GEMM kernels, which fails to schedule instructions, resulting in poor performance. Although both OpenBLAS and BLIS employ hand-crafted assembly-coded GEMM kernels, there exist differences in loop organization and register blocking strategies.

5.1.2 Large-Scale GEMM

Fig.9(b) shows the performance results of large-scale SGEMM on the Phytium 2000+ system. Because BLASFEO does not have a multi-threaded version, we only measure the performance of OpenBLAS, BLIS, and Eigen. The matrix size ranges from 5120 to 10240, with a step of 512, and we use 64 threads.

The performance of OpenBLAS, BLIS and Eigen does not change significantly over the matrix size. BLIS runs slightly faster than OpenBLAS, reaching over 70% of Phytium 2000+'s theoretical peak. Eigen employs C++-coded GEMM kernel, leading to bad performance among the three BLAS libraries, achieving only 38% of the theoretical peak. According to the prior experimental results [14], the single-thread GEMM performance on Phytium 2000+ can reach 91% of the theoretical peak. In terms of per thread GEMM performance, using 64 threads is 20%–25% slower than using a single thread. This slowdown comes from the overhead of thread management and the shared L2 cache by every four cores on Phytium 2000+.

5.2 SpMV

Sparse matrix-vector multiplication (SpMV) is an essential kernel in linear algebra and is widely used in

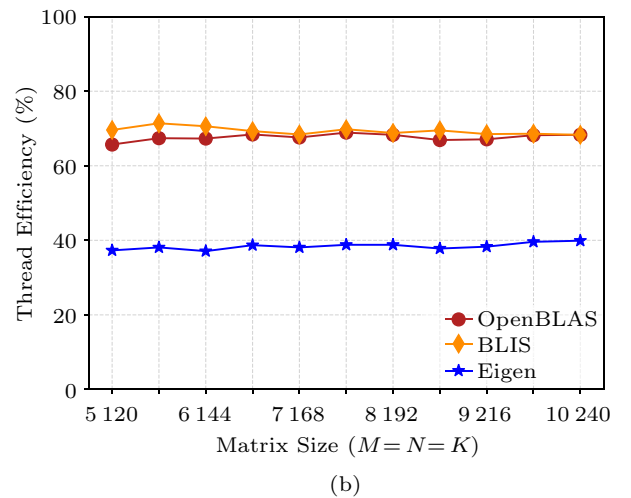
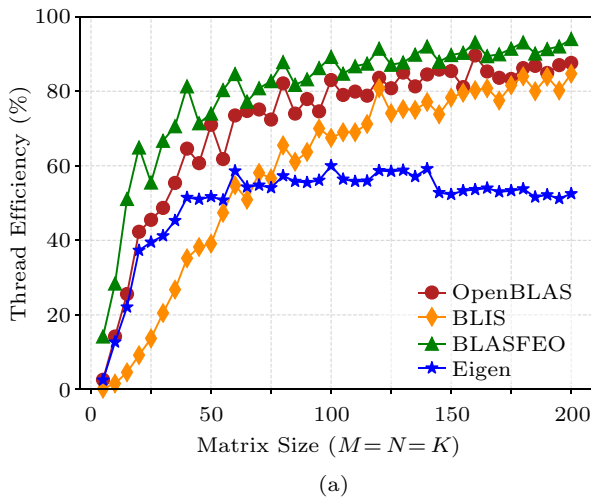


Fig.9. Performance of small-scale and large-scale SGEMM. (a) Small SGEMM with a single thread. (b) Large SGEMM with 64 threads.

scientific and engineering applications. This subsection evaluates the performance of the SpMV kernel on the Phytium 2000+ systems.

5.2.1 Impact of NUMA Bindings

As shown in Fig.8, SpMV is a memory-bound kernel with an arithmetic intensity of 0.25. Thus, memory accesses on Phytium 2000+ have a direct impact on SpMV’s performance. Phytium 2000+ exposes eight NUMA nodes where a group of eight cores are directly connected to a local memory module. Indirect access to remote memory modules is possible but 1.5x slower than accessing the local module (Subsection 3.2.5). We use the Linux NUMA utility, `numactl`, to allocate the required data buffers from the local memory module for a thread that performs SpMV computation [15].

As can be seen from Fig.10, the NUMA-aware memory allocation significantly outperforms the non-NUMA-aware counterpart, giving an average speedup ranging from 1.5x to 6x across five storage formats. As such, we enable static NUMA bindings on the Phytium 2000+ processor. We also observe that the ELL format consumes the largest memory buffers among the five storage formats, and thus can benefit the most from using manual NUMA bindings.

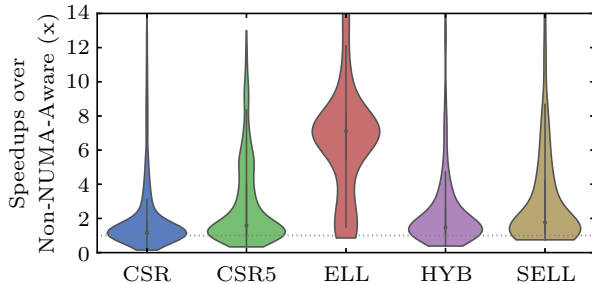


Fig.10. Speedup distribution of NUMA-aware memory allocation on Phytium 2000+ (64 threads). The thick black line shows where 50% of the data locates.

5.2.2 Overall SpMV Performance

Fig.11 shows the overall SpMV performance on Phytium 2000+ with five storage formats (CSR, CSR5 [16], ELL [17], SELL [18], and HYB [19]). Our estimation with the roofline model (Fig.8) indicates that the maximum SpMV performance with 64 threads is 23.4 Gflop/s. But the achieved performance is much larger than this estimated performance. This occurs when the input sparse matrix can be hold within the on-chip cache of Phytium 2000+.

We also see that there is no “one-size-fits-all” format across inputs. On the Phytium 2000+ platform, SELL

is the optimal format for around 50% of the sparse matrices and ELL gives the worst performance on most of the cases. Therefore we need an adaptive scheme to help developers to choose the optimal sparse matrix format [20–22].

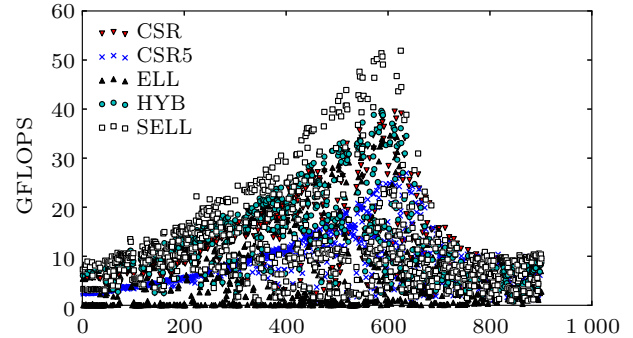


Fig.11. Overall SpMV performance on Phytium 2000+ (64 threads). The x -axis labels different sparse matrices ordered by the number of nonzeros, and the y -axis denotes the achieved SpMV performance in Gflops.

6 Related Work

For the effective use of the memory systems on modern architectures, researchers have obtained their performance results and disclosed implementation details through measurements.

Babka and Tuma [23] proposed experiments that investigate detailed parameters of the x86 processors. The experiment is built on a general benchmark framework and obtains the required memory parameters by performing one or a combination of multiple open-source benchmarks. It focuses on detailed parameters including the address translation miss penalties, the parameters of the additional translation caches, the cacheline size, and the cache miss penalties.

McCalpin [7] presented four benchmark kernels (`copy`, `scale`, `add`, and `triad`), `STREAM`, to assess the memory bandwidth for a large variety of current computers, including uniprocessors, vector processors, shared-memory systems, and distributed-memory systems. `STREAM` is one of the most commonly used memory bandwidth measurement tools in Fortran and C. But it focuses on throughput measurement without considering the latency metric.

Molka et al. [6] proposed a set of benchmarks, including to study the performance details of the Nehalem architecture. Based on these benchmarks, they obtained undocumented performance data and architectural properties. This is the first work to measure the core-to-core communication overhead, but it is only

applicable to the x86 architectures. Fang *et al.* extended the microkernels to Intel Xeon Phi [24]. Ramos and Hoefler [25] proposed a state-based modelling approach for memory communication, allowing algorithm designers to abstract away from the architecture and the detailed cache coherency protocols. The model is built based on the measurement numbers of the cache-coherent memory hierarchy.

As for the Phytium 2000+ architecture, there are a few related studies on performance evaluation and optimization. You *et al.* [4] evaluated three linear algebra kernels including matrix-matrix multiplication, matrix-vector multiplication, and triangular solver with both sparse and dense datasets, aiming to provide performance indicators for the prototype Tianhe-3 cluster built from the Phytium 2000+ architecture. Su *et al.* [14] presented a shared cache partitioning (SCP) method to eliminate inter-thread cache conflicts in the GEMM routines on the Phytium 2000+ architecture. This is achieved by partitioning a shared cache into physically disjoint sets and assigning different sets to different threads. Chen *et al.* evaluated and optimized the SpMV kernel from various angles [20–22].

Different from the previous work, we provide a systematic evaluation on the Phytium 2000+ many-core architecture with a memory-centric perspective. This evaluation work is performed at both the microbenchmark level and the kernel level, guided by a Phytium 2000+ specified roofline model.

7 Conclusions

This article presents a comprehensive performance evaluation of a 64-core ARMv8-based architecture. We focused on cache and memory subsystems, analyzing the characteristics that have an impact on high-performance computing applications. We provided insight into the relevant characteristics of the Phytium 2000+ processor using a set of micro-benchmarks. We then analyzed the Phytium 2000+ processor at the system level using the well-known roofline model. Using the knowledge gained from these micro-benchmarks, we optimized two applications and used them to assess the capabilities of the Phytium 2000+ system. The results showed that the ARMv8-based many-core system is capable of delivering high performance for a wide range of scientific kernels. Our evaluation results also indicated that the shared L2 cache by four cores can be a performance bottleneck for GEMM and leveraging a private L2 per core

is highly recommended for building future generations of Phytium processors.

Acknowledgement(s) We would like to thank the anonymous reviewers for their valuable and constructive comments. We thank Wei-Ling Yang and Wan-Rong Gao from National University of Defense Technology for the experiment support.

References

- [1] Laurenzano M A, Tiwari A, Cauble-Chantrenne A *et al.* Characterization and bottleneck analysis of a 64-bit ARMv8 platform. In *Proc. the 2016 IEEE International Symposium on Performance Analysis of Systems and Software*, April 2016, pp.36-45. DOI: [10.1109/ISPASS.2016.7482072](https://doi.org/10.1109/ISPASS.2016.7482072).
- [2] Stephens N. ARMv8-a next-generation vector architecture for HPC. In *Proc. the 2016 IEEE Hot Chips 28 Symposium*, August 2016. DOI: [10.1109/HOTCHIPS.2016.7936203](https://doi.org/10.1109/HOTCHIPS.2016.7936203).
- [3] Zhang C. Mars: A 64-core ARMv8 processor. In *Proc. the 2015 IEEE Hot Chips 27 Symposium*, Aug. 2015. DOI: [10.1109/HOTCHIPS.2015.7477454](https://doi.org/10.1109/HOTCHIPS.2015.7477454).
- [4] You X, Yang H, Luan Z, Liu Y, Qian D. Performance evaluation and analysis of linear algebra kernels in the prototype Tianhe-3 cluster. In *Proc. the 5th Asian Conference on Supercomputing Frontiers*, March 2019, pp.86-105. DOI: [10.1007/978-3-030-18645-6_6](https://doi.org/10.1007/978-3-030-18645-6_6).
- [5] Dongarra J. Report on the Fujitsu Fugaku system. Technical Report, University of Tennessee, 2020. <https://www.icl.utk.edu/files/publications/2020/icl-utk-1379-2020.pdf>, Nov. 2020.
- [6] Molka D, Hackenberg D, Schöne R, Müller M S. Memory performance and cache coherency effects on an Intel Nehalem multiprocessor system. In *Proc. the 18th International Conference on Parallel Architectures and Compilation Techniques*, September 2009, pp.261-270. DOI: [10.1109/PACT.2009.22](https://doi.org/10.1109/PACT.2009.22).
- [7] McCalpin J. Memory bandwidth and machine balance in current high performance computers. <https://www.cs.virginia.edu/stream/analyses.html>, Dec. 2020.
- [8] Kamil S, Husbands P, Oliker L, Shalf J, Yelick K A. Impact of modern memory subsystems on cache optimizations for stencil computations. In *Proc. the 2005 Workshop on Memory System Performance*, June 2005, pp.36-43. DOI: [10.1145/1111583.1111589](https://doi.org/10.1145/1111583.1111589).
- [9] Williams S, Waterman A, Patterson D A. Roofline: An insightful visual performance model for multicore architectures. *Commun. ACM*, 2009, 52(4): 65-76. DOI: [10.1145/1498765.1498785](https://doi.org/10.1145/1498765.1498785)
- [10] Ilic A, Pratas F, Sousa L. Cache-aware roofline model: Upgrading the loft. *IEEE Comput. Archit. Lett.*, 2014, 13(1): 21-24. DOI: [10.1109/L-CA.2013.6](https://doi.org/10.1109/L-CA.2013.6).
- [11] Liu X, Buono D, Checconi F, Choi J W, Que X, Petrini F, Gunnels J A, Stuecheli J. An early performance study of large-scale POWER8 SMP systems. In *Proc. the 2016 IEEE International Parallel and Distributed Processing Symposium*, May 2016, pp.263-272. DOI: [10.1109/IPDPS.2016.14](https://doi.org/10.1109/IPDPS.2016.14).

- [12] Goto K, van de Geijn R A. Anatomy of high performance matrix multiplication. *ACM Trans. Math. Softw.*, 2008, 34(3): Article No. 12. DOI: [10.1145/1356052.1356053](https://doi.org/10.1145/1356052.1356053).
- [13] Frison G, Kouzoupis D, Sartor T, Zanelli A, Diehl M. BLASFEO: Basic linear algebra subroutines for embedded optimization. *ACM Trans. Math. Softw.*, 2018, 44(4): Article No. 42. DOI: [10.1145/3210754](https://doi.org/10.1145/3210754).
- [14] Su X, Liao X, Jiang H, Yang C, Xue J. SCP: Shared cache partitioning for high-performance GEMM. *ACM Transactions on Architecture and Code Optimization*, 2019, 15(4): Article No. 43. DOI: [10.1145/3274654](https://doi.org/10.1145/3274654).
- [15] Hollowell C, Caramarcu C, Strecker-Kellogg W, Wong A, Zaytsev A. The effect of NUMA tunings on CPU performance. *Journal of Physics: Conference Series*, 2015, 664(9): Article No. 092010. DOI: [10.1088/1742-6596/664/9/092010](https://doi.org/10.1088/1742-6596/664/9/092010).
- [16] Liu W, Vinter B. CSR5: An efficient storage format for cross-platform sparse matrix-vector multiplication. In *Proc. the 29th ACM on International Conference on Supercomputing*, June 2015, pp.339-350. DOI: [10.1145/2751205.2751209](https://doi.org/10.1145/2751205.2751209).
- [17] Grimes R, Kincaid D, Young D. ITPACK 2.0 user's guide. Technical Report, Center for Numerical Analysis, University of Texas, 1979.
- [18] Kreutzer M, Hager G, Wellein G, Fehske H, Bishop A R. A unified sparse matrix data format for efficient general sparse matrix-vector multiplication on modern processors with wide SIMD units. *SIAM J. Sci. Comput.*, 2014, 36(5): 401-423. DOI: [10.1137/130930352](https://doi.org/10.1137/130930352).
- [19] Bell N, Garland M. Implementing sparse matrix-vector multiplication on throughput-oriented processors. In *Proc. the ACM/IEEE Conference on High Performance Computing*, November 2009. DOI: [10.1145/1654059.1654078](https://doi.org/10.1145/1654059.1654078).
- [20] Chen D, Fang J, Xu C, Chen S, Wang Z. Characterizing scalability of sparse matrix-vector multiplications on Phytium FT-2000+. *Int. J. Parallel Program.*, 2020, 48(1): 80-97. DOI: [10.1007/s10766-019-00646-x](https://doi.org/10.1007/s10766-019-00646-x).
- [21] Chen D, Fang J, Chen S, Xu C, Wang Z. Optimizing sparse matrix-vector multiplications on an ARMv8-based many-core architecture. *Int. J. Parallel Program.*, 2019, 47(3): 418-432. DOI: [10.1007/s10766-018-00625-8](https://doi.org/10.1007/s10766-018-00625-8).
- [22] Chen S, Fang J, Chen D, Xu C, Wang Z. Adaptive optimization of sparse matrix-vector multiplication on emerging many-core architectures. In *Proc. the 20th IEEE International Conference on High Performance Computing*, June 2018, pp.649-658. DOI: [10.1109/HPCC/SmartCity/DSS.2018.00116](https://doi.org/10.1109/HPCC/SmartCity/DSS.2018.00116).
- [23] Babka V, Tuma P. Investigating cache parameters of x86 family processors. In *Proc. the 2009 SPEC Benchmark Workshop*, January 2009, pp.77-96. DOI: [10.1007/978-3-540-93799-9_5](https://doi.org/10.1007/978-3-540-93799-9_5).
- [24] Fang J, Sips H J, Zhang L, Xu C, Che Y, Varbanescu A L. Test-driving Intel Xeon Phi. In *Proc. the 5th ACM/SPEC International Conference on Performance Engineering*, March 2014, pp.137-148. DOI: [10.1145/2568088.2576799](https://doi.org/10.1145/2568088.2576799).
- [25] Ramos S, Hoer T. Modeling communication in cache-coherent SMP systems: A case-study with Xeon Phi. In *Proc. the 22nd International Symposium on High-Performance Parallel and Distributed Computing*, June 2013, pp.97-108. DOI: [10.1145/2462902.2462916](https://doi.org/10.1145/2462902.2462916).




Jian-Bin Fang is an assistant professor in computer science at National University of Defense Technology (NUDT), Changsha. He obtained his Ph.D. degree in computer science from Delft University of Technology in 2014. His research interests include parallel programming for many-cores, parallel compilers, performance modeling, and scalable algorithms. He is a member of CCF.



Xiang-Ke Liao received his B.S. degree from Tsinghua University, Beijing, in 1985, and M.S. degree from National University of Defense Technology (NUDT), Changsha, in 1988, both in computer science. Currently he is a full professor of College of Computer at NUDT, Changsha. His research interests include high-performance computing systems, operating systems, and parallel and distributed computing. Prof. Liao is a fellow of CCF and an academician of Chinese Academy of Engineering.



Chun Huang is a full professor in computer science at National University of Defense Technology (NUDT), Changsha. Her research interests are high-performance computing, system software, parallel compilers, parallel programming, performance optimization, and high-performance math libraries.



De-Zun Dong received his B.S., M.S., and Ph.D. degrees in computer science from the National University of Defense Technology (NUDT), Changsha, in 2002, 2004, and 2010, respectively. He is a professor in the College of Computer, NUDT, Changsha, where he leads the research group of High-Performance Network And Architecture (HiNA) and serves as the deputy director designer of Tianhe supercomputer. His research interests focus on high-performance network and architecture for supercomputer, datacenter, and deep learning systems. He has published over 60 peer-reviewed papers in reputed international journals and conferences.