
BENCHCOUNCIL'S VIEW ON BENCHMARKING AI AND OTHER EMERGING WORKLOADS

EDITED BY

JIANFENG ZHAN (*BenchCouncil Steering Committee Chair*)

LEI WANG (*BenchCouncil Big Data and CPU Tracks Executive Committee Co-chair*)

WANLING GAO (*BenchCouncil Datacenter AI Track Executive Committee Co-chair*)

RUI REN (*BenchCouncil Testbed Co-chairs*)



BenchCouncil: International Open Benchmark Council

<http://www.benchcouncil.org>

TECHNICAL REPORT NO. BENCHCOUNCIL-BCVIEW-2019
NOV 12, 2019

BenchCouncil’s View on Benchmarking AI and Other Emerging Workloads

Jianfeng Zhan*¹, Lei Wang¹, Wanling Gao¹, and Rui Ren¹

¹BenchCouncil (International Open Benchmark Council)

Nov 12, 2019

Abstract

This paper outlines BenchCouncil’s view on the challenges, rules, and vision of benchmarking modern workloads like Big Data, AI or machine learning, and Internet Services. We conclude the challenges of benchmarking modern workloads as FIDSS (Fragmented, Isolated, Dynamic, Service-based, and Stochastic), and propose the PRDAERS benchmarking rules that the benchmarks should be specified in a paper-and-pencil manner, relevant, diverse, containing different levels of abstractions, specifying the evaluation metrics and methodology, repeatable, and scaleable. We believe proposing simple but elegant abstractions that help achieve both efficiency and general-purpose is the final target of benchmarking in future, which may be not pressing. In the light of this vision, we shortly discuss BenchCouncil’s related projects.

1 Challenges

Our society is increasing relying upon information infrastructure, which consists of massive IoT, edge devices, extreme-scale datacenters, and high-performance computing systems. Those systems collaborate with each other to handle big data and leverage AI technique, and finally provide Internet services for huge end users with guaranteed quality of services. From a perspective of workload characterization, the emerging workloads like Big Data, AI, and Internet Services raise serious FIDSS (Fragmented, Isolated, Dynamic, Service-based, and Stochastic) challenges, which are significantly different from the traditional workloads characterized by SPEC CPU (desktop workloads) [3], TPC-C [4], TPC-Web (Traditional web services) [5], and HPL (high performance computing) [1] benchmarks.

The first challenge is fragmented. There are huge fragmented application scenarios, a marked departure from the past. However, there is a lack of simple but elegant abstractions that help achieve both efficiency and general-purpose. For example, for database, the relation algebra demonstrates its generalized ability, and any complex query can be written using five primitives like select, project, product, union, and difference [13]. However, in a new era of big data, hundreds or even thousands ad-hoc solutions are proposed to handle different application scenarios, most of which are termed with NoSQL or NewSQL. For AI, the same observation holds true. There are tens or even hundreds of organizations who are developing AI training or inference chips to tackle their challenges in different application scenarios, respectively [30]. Though domain-specific software and hardware co-design is promising [23], the lack of simple but unified abstractions has two side effects. On one hand, it is challenging to amortize the cost of building an ad-hoc solution. On the other hand, single-purpose is a structure obstacle to resource sharing. Proposing simple but elegant abstractions that help achieve both efficiency and general-purpose is our final target of workload modeling, benchmarking, or characterization in future, which may be not pressing.

*Jianfeng Zhan is the corresponding author.

The second challenge is isolated. The real-world data sets and workloads or even AI models are treated as first-class confidential issues, and they are hidden with giant Internet service providers’ datacenters, and isolated between academia and industry, or even among different providers [16]. This kind of isolation finally poses a huge obstacle for our communities towards developing an open and mature research field [15].

The third challenge is the dynamic complexity. The dynamic complexity of emerging workloads is in terms with that the common requirements are handled collaboratively by datacenters, edge, and IoT devices. It is also manifested by that different distributions of data sets, workloads, machine learning or AI models may substantially affect the system’s behaviors. Unfortunately, the system architectures are undergoing fast evolutions in terms of the interactions among IoT, edge, and datacenters.

The fourth challenge is the side effect of service-based architecture. On one hand, the software-as-a-service (SaaS) development and deploy model makes the workloads change very fast (so-called workload churn) [10], and it is not scalable or even impossible to create a new benchmark or proxy for every possible workload [18]. On the other hand, modern Internet services adopt a microservice-based architecture, which is often distributed across different datacenters, and consist of diversity of various modules with very long and complex execution paths. As the worst-case performance (tail latency) [14] does matter, the micro-service-based architecture also pose a serious challenge to benchmarking [16].

The final but not least challenge is the side effect of the stochastic nature of AI. The AI techniques are widely used to augment modern products or Internet services. The nature of AI is stochastic, allowing multiple different but equally valid solutions [29]. The uncertainness of AI is manifested by adverse effect of lower-precision optimization on the quality of the final model, effect of scaling training on time-to-quality, and run-to-run variation in terms of epochs-to-quality [29]. However, the benchmarks mandates being repeatable, reliable, and reproducible. This conflict raises serious challenges.

2 BenchCouncil Benchmarking Rules

After revisiting the previous successful benchmarks, we propose the PRDAERS benchmarking rules as follows.

First, the common requirements should be specified only algorithmically in a paper-and pencil approach. This rule is firstly proposed by the NAS parallel benchmarks [8], and well-practiced in the database community. Interestingly, this rule is often overlooked by the architecture and system communities. Following this rule, the benchmark specification should be proposed firstly and reasonably divorced from individual implementations. In general, the benchmark specification should define a problem domain in a high-level language.

Second, the benchmark should be relevant [20]. On one hand, the benchmark should be domain-specific [27] and distinguish between different contexts like IoT, edge, Datacenter and HPC. Under each context, each benchmark should provide application scenarios abstracted from details, semantic-preserving data sets, or even quality targets (for AI tasks) that represent real-world deployments [30]. On the other hand, the benchmark should be simplified. It is not a copy of a real-world application. Instead, it is a distillation of the essential attributes of a workload [27]. Generally, a real-world application is not portable across different systems and architectures.

Third, the diversity and representativeness of a widely accepted benchmark suite are of paramount importance. On one hand, this is a tradition witnessed by the past. For example, SPEC CPU 2017 contains 43 benchmarks. The other examples include PARSEC3.0 (30), TPC-DS (99). On the other hand, modern workloads manifest much higher complexity. For example, Google’s datacenter workloads show significant diversity in workload behavior with no single “silver-bullet” application to optimize for [26]. Modern AI models vary wildly, and a small accuracy change (e.g., a few percent) can drastically change the computational requirements (e.g., 5-10x) [30, 11]. For modern deep learning workloads, someone may argue running an entire training session is costly, so few benchmarks should be included for reducing the cost. However, we believe the cost of execution time cannot justify including only a few benchmarks. Actually, the cost of execution time for other benchmarks (like HPC, SPEC CPU on simulators) is also

prohibitively costly. So, for workload characterization, diverse workloads should be included to exhibit the range of behavior of the target applications, or else it will oversimplify the typical environment [27]. On the other hand, for performance ranking (benchmarking), it may be reasonable for us to choose a few representative benchmarks for reducing the cost just like that the HPC Top500 ranking only reports HPL, HPCG, and Graph500 (three benchmarks out of 20+ representative HPC benchmarks like HPCC, NPB).

Fourth, the benchmarks should contain different levels of abstractions, and usually a combination of micro, component and end-to-end application benchmarks is preferred to. From an architectural perspective, porting a full-scale application to a new architecture at an earlier stage is difficult and even impossible [9], while using micro or component benchmarks alone are insufficient to discover the time breakdown of different modules and locate the bottleneck within a realistic application scenario at a later stage [9]. Hence, a realistic benchmark suite should have the ability to run not only collectively as a whole end-to-end application to discover the time breakdown of different modules but also individually as a micro or component benchmark for fine tuning hot spot functions or kernels [16].

Fifth, it should specify the evaluation metrics and methodology. The performance number should be simple, linear, orthogonal, and monotonic [27]. Meanwhile, it should be domain relevant. For example, the time-to-quality (i.e., state-of-the-art accuracy) metric is relevant to the AI domain, because some optimizations immediately improve throughput while adversely affect the quality of the final model, which can only be observed by running an entire training session [29].

Sixth, the benchmark should be repeatable, reliable, and reproducible [7]. Since many modern deep learning workloads are intrinsically approximate and stochastic, allowing multiple different but equally valid solutions [29], it will raise serious challenges for AI benchmarking.

Finally, but not least, the benchmark should be scaleable [20], and the benchmark users can scale up the problem size, so the benchmark is applicable for running on both small and large systems. However, this is not trivial for modern deep learning workloads, as accommodating system scale even requires changing hyperparameters, which can affect the amount of computation to reach a particular quality target [29].

3 Our Viewpoint

Domain-specific software-hardware co-design is promising as a single-purpose solution will achieve high energy-efficiency than that of a general-purpose one. However, we believe the final target in future is to propose simple but elegant abstractions that help achieve both efficiency and general-purpose for big data, AI, and Internet services.

To meet with the fragmented application scenarios, currently, BenchCouncil sets up (datacenter) Big Data, datacenter AI, HPC AI, AIoT, and Edge AI benchmarking tracks, and released Big-DataBench [19] (<http://www.benchcouncil.org/BigDataBench/index.html>) and AIBench [16, 15](<http://www.benchcouncil.org/AIBench/index.html>) benchmark suites for datacenter big data and AI, Edge AIBench [22](<http://www.benchcouncil.org/EdgeAIBench/index.html>) for edge AI, AIoT Bench [28](<http://www.benchcouncil.org/AIoTBench/index.html>)for IoT AI, and HPC AI500 [25] (<http://www.benchcouncil.org/HPCAI500/index.html>) for HPC AI. Those benchmarks are in fast evolution. We release the source code, pre-trained model, and container-based deployment on the BenchCouncil web site (<http://benchcouncil.org/testbed/index.html>).

On the other side, we propose an innovative approach to modeling and characterizing the emerging workloads. We consider each big data, AI and Internet service workload as a pipeline of one or more classes of unit of computation performed on different initial or intermediate data inputs, which we call a data motif [18]. After thoroughly analyzing a majority of workloads in five typical big data application domains (search engine, social network, e-commerce, multimedia and bio-informatics), we identify eight data motifs that take up most of run time, including Matrix, Sampling, Logic, Transform, Set, Graph, Sort and Statistic [18]. We found the combinations of one or more data motifs with different weights in terms of runtime can describe most of big data and AI workloads we investigated [18, 17]. In conclusion, the data motifs are promising as simple but elegant abstractions that achieve both efficiency

and general-purpose.

On the basis of the data motif methodology, we are proposing a new benchmark suite, named BENCHCPU [6], to characterize emerging workloads, including Big Data, AI, and Internet Services. The goal of BENCHCPU is to abstract ISA (Instruction Set Architecture) independent workload Characterizations from emerging workloads. BENCHCPU will be portable across edge, IoT, and datacenter processor architectures. Furthermore, we are working on an open-source chip project, named EChip. On the basis of data dwarf approaches, the goal of EChip is to design an open source general-purpose ISA for emerging workloads, which is a marked departure from single-purpose accelerators. The ISA of EChip is composed of the general-purpose instruction set and the domain-specific instruction set. The general-purpose instruction set is modular-based basic instruction set, which always remains unchanged (Minimum implementation subset) and is compatibility with Linux ecosystems. The domain-specific instruction set is an extension of domain-specific customization, proposed for emerging workloads like big data, AI, internet service.

4 Revisiting the Benchmarking Principles and Methodology

Benchmarking principle defines the rules and guidelines on what the important criteria need to be considered for a good benchmark [7, 31]. The benchmarking methodology specifies systematic strategies and processes on how to construct a benchmark. Jim Gray summarizes four key criteria to define a good benchmark—relevant, portable, scalable, and simple [21]. Also, other benchmark consortiums propose their own principles and methodologies.

The TPC Benchmarks are a series of domain-specific benchmarks targeting measuring transaction processing (TP) and database (DB) performance. They believe that a domain-specific benchmark should satisfy three criteria [27, 24]: (1) No single metric can measure the application performance of all domains, (2) The more general the benchmark, the less useful it is for anything in particular, (3) A benchmark is a distillation of the essential attributes of a workload. They adopt a benchmark methodology with the concept of "functions of abstraction" and "functional workload model", which abstract the compute units that frequently appeared with repetitions or similarities.

SPEC benchmarks are a set of benchmarks for the newest generation of computing systems [2], guided by six implicit design principles: application-oriented; portable; repeatable and reliable of benchmarking results; consistent and fair across different users or different systems; diverse workloads that can run independently, and reporting unit of measurement, e.g., throughput.

PARSEC (Princeton Application Repository for Shared-Memory Computers) is a benchmark suite for chip multiprocessors. The PARSEC benchmarks are constructed following five requirements [12]. (1) The benchmarks should be multi-threaded. (2) Emerging workloads should be considered. (3) The benchmarks should cover diverse applications and a variety of platforms, and accommodate different usage models. (4) The benchmarks should use state-of-the-art algorithms and data structures. (5) The benchmarks should be designed to support research.

5 Conclusion

This paper outlines BenchCouncil's view on the challenges, rules, and vision of benchmarking modern workloads. We conclude the challenges of benchmarking modern workloads as Fragmented, Isolated, Dynamic, Service-based, and Stochastic. After revisiting previous successful benchmarks, we propose the PRDAERS benchmarking rules that the benchmarks should be specified in a paper-and-pencil manner, relevant, diverse, containing different levels of abstractions, specifying the evaluation metrics and methodology, repeatable, and scaleable. We believe proposing simple but elegant abstractions that help achieve both efficiency and general-purpose is the final target of benchmarking in future, which may be not pressing. In the light of this vision, we shortly discuss BenchCouncil's related projects, including BigDataBench, AIBench, HPC AI500, Edge AIBench, AIoT Bench, and BENCHCPU, and EChip projects.

References

- [1] “Hpl-a portable implementation of the high-performance linpack benchmark for distributed-memory computers,” <http://www.netlib.org/benchmark/hpl/>.
- [2] “Spec,” <https://www.spec.org/>.
- [3] “Spec cpu 2017 benchmark,” <https://www.spec.org/cpu2017/>.
- [4] “Tpc-c benchmark,” <http://www.tpc.org/tpcc/>.
- [5] “Tpc-w benchmarks,” <http://www.tpc.org/tpcw/>.
- [6] “Benchcpu,” <http://http://www.benchcouncil.org/BenchCPU/index.html>, 2019.
- [7] R. Angles, “Benchmark principles and methods,” in *Linked Data Benchmark Council (LDBC). Project No 317548, European Community’s Seventh Framework Programme FP7*, 2012-2014.
- [8] D. H. Bailey, “Nas parallel benchmarks,” *Encyclopedia of Parallel Computing*, pp. 1254–1259, 2011.
- [9] D. H. Bailey, E. Barszcz, J. T. Barton, D. S. Browning, R. L. Carter, L. Dagum, R. A. Fatoohi, P. O. Frederickson, T. A. Lasinski, R. S. Schreiber, H. Simon, V. Venkatakrisnan, and S. Weeratunga, “The nas parallel benchmarks,” *The International Journal of Supercomputing Applications*, vol. 5, no. 3, pp. 63–73, 1991.
- [10] L. A. Barroso and U. Hölzl, “The datacenter as a computer: An introduction to the design of warehouse-scale machines,” *Synthesis Lectures on Computer Architecture*, vol. 8, no. 3, 2009.
- [11] S. Bianco, R. Cadene, L. Celona, and P. Napoletano, “Benchmark analysis of representative deep neural network architectures,” *IEEE Access*, vol. 6, pp. 64 270–64 277, 2018.
- [12] C. Bienia and K. Li, *Benchmarking modern multiprocessors*. Princeton University USA, 2011.
- [13] E. F. Codd, “A relational model of data for large shared data banks,” *Communications of the ACM*, vol. 13, no. 6, pp. 377–387, 1970.
- [14] J. Dean and L. A. Barroso, “The tail at scale,” *Communications of the ACM*, vol. 56, no. 2, pp. 74–80, 2013.
- [15] W. Gao, C. Luo, L. Wang, X. Xiong, J. Chen, T. Hao, Z. Jiang, F. Fan, M. Du, Y. Huang, F. Zhang, X. Wen, C. Zheng, X. He, and J. Dai, “AIBench: Towards scalable and comprehensive datacenter AI benchmarking,” in *International Symposium on Benchmarking, Measuring and Optimization(Bench)*, 2018.
- [16] W. Gao, F. Tang, L. Wang, J. Zhan, C. Lan, C. Luo, Y. Huang, C. Zheng, J. Dai, Z. Cao, H. Tang, K. Zhan, B. Wang, D. Kong, T. Wu, M. Yu, C. Tan, H. Li, X. Tian, Y. Li, G. Lu, J. Shao, Z. Wang, X. Wang, and H. Ye, “AIBench: An Industry Standard Internet Service AI Benchmark Suite,” <https://arxiv.org/abs/1908.08998>, 2019.
- [17] W. Gao, J. Zhan, L. Wang, C. Luo, Z. Jia, D. Zheng, C. Zheng, H. Ye, H. Wang, and R. Ren, “Data motif-based proxy benchmarks for big data and AI workloads,” *IISWC 2018*, 2018.
- [18] W. Gao, J. Zhan, L. Wang, C. Luo, D. Zheng, F. Tang, B. Xie, C. Zheng, X. Wen, X. He, H. Ye, and R. Ren, “Data motifs: A lens towards fully understanding big data and ai workloads,” in *Parallel Architectures and Compilation Techniques (PACT), the 27th International Conference*, 2018.

- [19] W. Gao, J. Zhan, L. Wang, C. Luo, D. Zheng, X. Wen, R. Ren, C. Zheng, X. He, H. Ye, H. Tang, Z. Cao, S. Zhang, and J. Dai, “BigDataBench: A scalable and unified big data and ai benchmark suite,” *arXiv preprint arXiv:1802.08254*, 2018.
- [20] J. Gray, *Benchmark handbook: for database and transaction processing systems*. Morgan Kaufmann Publishers Inc., 1992.
- [21] J. Gray, “Database and transaction processing performance handbook.” 1993.
- [22] T. Hao, Y. Huang, X. Wen, W. Gao, F. Zhang, C. Zheng, L. Wang, H. Ye, K. Hwang, Z. Ren, and J. Zhan, “Edge AIBench: Towards comprehensive end-to-end edge computing benchmarking,” *2018 BenchCouncil International Symposium on Benchmarking, Measuring and Optimizing (Bench18)*, 2018.
- [23] J. Hennessy and D. Patterson, “A new golden age for computer architecture: Domain-specific hardware/software co-design, enhanced security, open instruction sets, and agile chip development.” 2018.
- [24] K. Huppler, “The art of building a good benchmark,” in *Technology Conference on Performance Evaluation and Benchmarking*. Springer, 2009, pp. 18–30.
- [25] Z. Jiang, W. Gao, L. Wang, X. Xiong, Y. Zhang, X. Wen, C. Luo, H. Ye, X. Lu, Y. Zhang, S. Feng, K. Li, W. Xu, and J. Zhan, “HPC AI500: A benchmark suite for hpc ai systems,” *2018 BenchCouncil International Symposium on Benchmarking, Measuring and Optimizing (Bench18)*, 2018.
- [26] S. Kanev, J. P. Darago, K. Hazelwood, P. Ranganathan, T. Moseley, G.-Y. Wei, and D. Brooks, “Profiling a warehouse-scale computer,” in *Computer Architecture (ISCA), 2015 ACM/IEEE 42nd Annual International Symposium on*. IEEE, 2015, pp. 158–169.
- [27] C. Levine, “Tpc benchmarks,” in *SIGMOD International Conference on Management of Data - Industrial Session*, 1997.
- [28] C. Luo, F. Zhang, C. Huang, X. Xiong, J. Chen, L. Wang, W. Gao, H. Ye, T. Wu, R. Zhou, and J. Zhan, “AIoT Bench: Towards comprehensive benchmarking mobile and embedded device intelligence,” *BenchCouncil International Symposium on Benchmarking, Measuring and Optimizing (Bench18)*, 2018.
- [29] P. Mattson, C. Cheng, C. Coleman, G. Diamos, P. Micikevicius, D. Patterson, H. Tang, G.-Y. Wei, P. Bailis, V. Bittorf *et al.*, “MLPerf training benchmark,” *arXiv preprint arXiv:1910.01500*, 2019.
- [30] V. J. Reddi, C. Cheng, D. Kanter, P. Mattson, G. Schmuelling, Carole-Jean Wu, B. Anderson, M. Breughe, M. Charlebois, W. Chou, R. Chukka, C. Coleman, S. Davis, P. Deng, G. Diamos, J. Duke, D. Fick, J. S. Gardner, I. Hubara, S. Idgunji, T. B. Jablin, J. Jiao, T. S. John, P. Kanwar, D. Lee, J. Liao, A. Lokhmotov, F. Massa, P. Meng, P. Micikevicius, C. Osborne, G. Pekhimenko, A. T. R. Rajan, D. Sequeira, A. Sirasao, F. Sun, H. Tang, M. Thomson, F. Wei, Ephrem Wu, L. Xu, K. Yamada, B. Yu, G. Yuan, A. Zhong, P. Zhang, and Y. Zhou, “MLPerf inference benchmark,” <https://arxiv.org/abs/1911.02549>, 2019.
- [31] L. Zhao, W. Gao, and Y. Jin, “Revisiting benchmarking principles and methodologies for big data benchmarking,” in *BPOE*. Springer, 2015, pp. 3–9.