

High Volume Computing: Identifying and Characterizing Throughput Oriented Workloads in Data Centers

Jianfeng Zhan, Lixin Zhang, Ninghui Sun, Lei Wang, Zhen Jia, and Chunjie Luo
State Key Laboratory of Computer Architecture
Institute of Computing Technology
Chinese Academy of Sciences
Beijing, China

Email: {zhanjianfeng, zhanglixin, snh}@ict.ac.cn, {wl, jiazhen, luochunjie}@ncic.ac.cn

Abstract—For the first time, this paper systematically identifies three categories of throughput oriented workloads in data centers: services, data processing applications, and interactive real-time applications, whose targets are to increase the volume of throughput in terms of *processed requests or data*, or *supported maximum number of simultaneous subscribers*, respectively, and we coin a new term *high volume computing (in short HVC)* to describe those workloads and data center computer systems designed for them. We characterize and compare HVC with other computing paradigms, e.g., high throughput computing, warehouse-scale computing, and cloud computing, in terms of levels, workloads, metrics, coupling degree, data scales, and number of jobs or service instances. We also preliminarily report our ongoing work on the metrics and benchmarks for HVC systems, which is the foundation of designing innovative data center computer systems for HVC workloads.

Keywords—High volume computing; Throughput-oriented workloads; Data center computer systems; Metrics; Benchmarks;

I. INTRODUCTION

In the past decade, there are three trends in computing domains. First, more and more *services*, involving a large amount of data, are deployed in data centers to serve the masses, e.g., Google search engine and Google Map. Second, massive data are produced, stored, and analyzed in real time or off line. According to the annual survey of the global digital output by IDC, the total amount of global data passes 1.2 zettabytes in 2010. In this paper, we call applications that produce, store, and analyze massive data *data processing applications*,

which is also referred to as big data applications. Third, lots of users tend to use streaming media or VoIP for fun or communications. Different from an ordinary Web server, a VoIP application will maintain a user session of a long period (e.g., more than five minutes) while guaranteeing the *real time* quality of service, which we call *an interactive real-time application*.

The workloads mentioned above consist of a large amount of loosely coupled workloads instead of a big tightly coupled job. The nature of this class of workloads is throughput-oriented, and the target of data center computer systems designed for them is to increase the volume of throughput in terms of processed requests (for services), or processed data (for data processing applications), or the maximum number of simultaneous subscribers (for interactive real-time applications), performed or supported in data centers. So as to pay attention to this class of workloads and computer systems designed for them, in this paper, we coin a new term *high volume computing* (nature: throughput computing; target: high volume, in short HVC) to describe this class of workloads or data center computer systems designed for them.

In this paper, we identify three categories of workloads in HVC: *services*, *data processing applications*, and *interactive real-time applications*, all of which are throughput-oriented workloads. A service is a group of applications that collaborate to receive user requests and return responses to end users. More and more emerging services

are data-intensive, e.g., Google search engine or Google Map. Data processing applications produce, store, and analyze massive data, and we only focus on *loosely coupled* data processing applications, excluding tightly coupled data-intensive computing, e.g., those written in MPI. Typical examples are MapReduce or Dryad based computing. We also include *data stream* applications that process continuous unbounded streams of data in real time into the second category of HVC workloads. Different from an ordinary Web server, an interactive real-time application will maintain a user session of a long period while guaranteeing the *real time* quality of service. Typical interactive real-time applications include streaming media, desktop clouds [33], and Voice over IP (VoIP) applications. The details of three categories of workloads can be found at II-A.

Despite several computing paradigms are not formally or clearly defined, e.g., warehouse-scale computing (WSC) [4], data-intensive scalable computing (DISC) [28], and cloud computing, we compare HVC with several computing paradigms in terms of six dimensions: levels, workloads, metrics, coupling degree, data scales, and number of jobs or service instances as shown in Table I. Our definition of HVC is towards data center computer systems, while many task computing [14] and high throughput computing [26] are defined towards runtime systems. In terms of workloads and respective metrics, both high throughput computing and high performance computing are about scientific computing centering around floating point operations, while most of HVC applications have few floating point operations as uncovered in our preliminary work [3]. Meanwhile, we also notice that many emerging workloads can be included into one or two categories of HVC workloads, e.g., WSC (into the first and second categories), and DISC (into the second category). As for (public) cloud computing, we believe it is basically a business model of renting computing or storage resources, which heavily relies upon virtualization technologies, however HVC is defined in terms of workloads. We think many well-know workloads in cloud [5] can be included into HVC, but HPC

in cloud workloads [1] [38] are excluded, since they are tightly-coupled.

After widely investigating previous benchmarks, we found there is no systematic work on benchmarking data center computer systems in the context of our identified three categories of throughput-oriented workloads. We present our preliminary work on the metrics and benchmarks for HVC systems.

The remainder of the paper is organized as follows. Section II characterize and compare different computing paradigms. Section III revisits previous benchmarks and report our preliminary work on the HVC metrics and benchmarks. Section IV draws a conclusion.

II. CHARACTERIZING COMPUTING PARADIGMS

In this section, we give out the definition of HVC, and identify its distinguished differences from other computing paradigms.

A. What is HVC?

HVC is a data center based computing paradigm focusing on throughput-oriented workloads. The target of a data center computer system designed for HVC workloads is to increase the volume of throughput in terms of *requests*, or *processed data*, or *the maximum number of simultaneous subscribers*, which are performed or supported in a data center.

In Table I, we characterize HVC from six dimensions: levels, workloads, metrics, coupling degree, data scales, and number of jobs or service instances.

The HVC system is defined on a data center level. We identify three categories of workloads in HVC: *services*, *data processing applications*, and *interactive real-time applications*. Services belong to the first category of HVC workloads. A service is a group of applications that collaborate to receive user requests and return responses to end users. We call a group of applications that independently process requests *a service instance*. For a large Internet service, a large amount of service instances are deployed with requests distribution enabled by load balancers. Since each

Table I: Characterizing different computing paradigms.

<i>Computing paradigm</i>	<i>level</i>	<i>Workloads</i>	<i>Metrics</i>	<i>Coupling degree</i>	<i>Data scale</i>	<i># jobs or service instances</i>
<i>High performance computing</i>	Super computers	Scientific computing: heroic MPI applications	Float point operations per second	Tight	n/a	Low
<i>High performance throughput computing [27]</i>	Processors	Traditional server workloads	Overall work performed over a fixed time period	loose	n/a	Low
<i>High throughput computing [26]</i>	Distributed runtime systems	Scientific computing	Float point operations per month	loose	n/a	Medium
<i>Many task computing [14]</i>	Runtime systems	Scientific computing or data analysis: workflow jobs	Tasks per second	Tight or loose	n/a	Large
<i>Data-intensive scalable computing [28] or data center computing [29]</i>	Runtime systems	Data analysis: MapReduce-like jobs	n/a	Loose	Large	Large
<i>Warehouse-scale computing [4]</i>	Data centers for Internet services, belonging to a single organization	Very large Internet services	n/a	Loose	large	Large
<i>Cloud computing [34] [15]</i>	Hosted data centers	SaaS + utility computing	n/a	Loose	n/a	Large
<i>High volume computing (HVC)</i>	Data centers	Services	Requests per minutes and joule	Loose	Medium	Large
		Data processing applications	Data processed per minute and joule	Loose	Large	Large
		Interactive real-time applications	Maximum number of simultaneous subscribers and subscribers per watt	Loose	From medium to large	Large

request is independent, a service in itself is loosely coupled. For an ordinary Apache Web server, the data scale is lower, while for a search engine provided by Google, the data scale is large. More and more emerging services are data-intensive.

The second category of HVC workloads is *data processing applications*. Please note that we only include loosely coupled data-intensive computing, e.g., MapReduce jobs, into HVC, *excluding data-intensive MPI applications*. In running, MapReduce tasks are independent, significantly different from batch jobs of programming models like MPI in which tasks execute concurrently and communicate during their execution [32]. The data scale of this category of workloads is large, which are also referred to as *big data applications*, and hence will produce large amount of tasks. We also include

data stream applications into this category of HVC workloads. For example, S4 is a platform that allows programmers to easily develop applications for processing continuous unbounded streams of data [37].

The third category of HVC applications is *interactive real-time applications*. Different from an ordinary Web server, an interactive real-time application will maintain a user session of a long period while guaranteeing the *real time* quality of service. Typical interactive real-time applications include streaming media—multimedia that is constantly delivered to an end-user by a service provider, desktop clouds [33], and Voice over IP (VoIP) applications. For this category of applications, the workload is loosely coupled because of independent requests or desktop applications; the data

scale varies from medium to large, and the number of tasks or service instances is large.

B. Identifying differences of HVC

1) *High performance computing*: There are two-fold differences of HVC from high performance computing (in short, HPC): first, workloads are different. HPC is mainly about scientific computing and usually a large-scale heroic MPI application, which is tightly coupled, while HVC is loosely coupled and commonly composed of a large amount of jobs or service instances. Second, the metric is different. The metric of HPC is *floating point operation per second*. However, in HVC, most of workloads, e.g., web search engines reported in [3], have few float point operations.

2) *High throughput computing*: Livny *et al.* refer to the environments that can be deliver large amounts of processing capacity over very long periods of time as high throughput computing [26]. There are three-fold differences of of HVC: first, high throughput computing is defined on the level of distributed runtime systems, while HVC is defined on the level of data center computer systems; second, the workloads of high throughput computing are towards scientific computing, while our HVC includes three categories of applications; third, the metric of high throughput computing is floating point operations per month or year. However, in HVC, major workloads, e.g., web search engines [3], have few float point operations.

3) *Many task computing*: According to [14], many task computing differs from high throughput computing in the emphasis of using large number of computing resources over short periods of time to accomplish many computational tasks where primary metrics are measured in seconds (e.g. FLOPS, tasks/sec, MB/s I/O rates), as opposed to operations per month. In terms of workloads, many task computing denotes high-performance computations comprising multiple distinct activities, coupled via file system operations [14]. The set of tasks may be loosely coupled or tightly coupled.

The differences of HVC are as follows: first, the concept of many task computing is towards runtime systems, while our HVC is on the level of data center computer systems; second, the workloads and respective metrics are different. Many task computing includes tightly or loosely coupled applications from scientific computing or data analysis domains, while HVC includes three categories of loosely coupled applications.

4) *High performance throughput computing*: Chaudhry *et al.* [27] call the combination of *high-end single-thread performance* and *a high degree of multicore, multithreading* high performance throughput computing. According to [27], systems designed for throughput computing emphasize the the aggregate amount of computation performed by all functional units, threads, cores, chips, coprocessors and network interface cards in a system over a fixed time period as opposed to focusing on a speed metric describing how fast a single core or a thread executes a benchmark.

Different from HVC, high performance throughput computing is defined on the level of processors, targeting traditional server workloads, e.g., database workloads, TPC-C, SPECint2000, SPECfp2000 exemplified in [27], while HVC is defined on the data center level. Different from high performance throughput computing, we have identifies three categories of workloads in HVC. In terms of workloads, high performance throughput computing mainly focuses on the first category of HVC applications (excluding data-intensive services like search engines) in addition to float point operation-intensive benchmarks like SPECfp2000. Moreover, we define different metrics for three different categories of HVC applications in III-B.

5) *Data-intensive scalable computing (in short, DISC) or data center computing*: Bryant *et al.* did not formally define what is DISC [28]. Instead, they characterized DISC through comparing it with cloud computing and high performance computing in [28]: first, cloud computing is towards hosted services, e.g., web-based Email services, while DISC refers to very large, shared data repository enabling complex analysis [28]; sec-

ond, from a perspective of programming model, a HPC program is described at very low level with specifying detailed control of processing communications, while DISC applications are written in terms of high-level operations on data, and the runtime system controls scheduling and load balancing [28]. In general, a DISC application is written in the MapReduce programming model, which splits jobs into small tasks that are run on the clusters compute nodes. No one formally defines what is data center computing. In [29], it refers to computing performed with frameworks such as MapReduce, Hadoop, and Dryad. Through these frameworks, computation can be performed on large datasets in a fault-tolerant way, while hiding the complexities of the distributed nature of the cluster [29]. According to [29], data center computing shares the same view like that of DISC.

Basically, DISC or data center computing can be included into the second category of HVC workloads. In addition, our HVC is defined toward data center computer systems instead of MapReduce-like runtime systems mentioned in [28] [29]. Second, Bryant *et al.* do not formally define the metrics, while we define three different metrics in evaluating data center computer systems for three categories of applications in III-B.

6) *Warehouse-scale computing*: According to [4], the trend toward server-side computing and the exploding popularity of Internet services has created a new class of computing systems that Hoelzle *et al.* have named warehouse-scale computers, or WSC. WSC is meant to call attention to the most distinguishing feature of these machines: the massive scale of their software infrastructure, data repositories, and hardware platform [4]. WSC demonstrates the following characteristics [4]: belonging to a single organization, using a relatively homogeneous hardware and system software platform, and sharing a common systems management layer [4]. Most importantly, WSCs run a smaller number of very large Internet services [4].

There are two differences of HVC from WSC: first, WSC is towards data centers for Internet services, belonging to a single organization, while in HVC, many small-to-medium scale services

from different organization may be hosted in the same data center, the aggregate number of which is large. Second, in terms of workloads, WSC can be included in HVC, meanwhile HVC covers more categories of workloads, e.g., interactive real-time applications. Moreover, we have given out the specific metrics for different categories of workloads, while in WSC, no metric is defined in [4].

7) *Cloud computing*: Vaquero *et al.* [34] defined cloud as a large pool of easily usable and accessible virtualized resources (such as hardware, development platforms and/or services). According to the widely cited Berkeley's technology report [15], when a cloud is made available in a pay-as-you-go manner to the general public, it is called a public cloud; the service being sold is utility computing [15]. The term private cloud is used to refer to internal data centers of a business or other organization, not made available to the general public [15], which is also characterized as warehouse-scale computing in [4] or data center computing in [29].

Basically, we believe that (public) cloud computing is a fine grain pay-per-use business model of renting computing or storage resources, which heavily relies upon virtualization technologies, while our HVC is defined in terms of workloads. Though virtualization technologies indeed bring changes to workloads, which is worth further investigation, however, well-known workloads in cloud studied in [5], e.g., *NoSQL* data serving, MapReduce, Web search, can be included into HVC. A new trend is that practitioners in HPC communities advocate an innovative computing paradigm—HPC in cloud [1] [38], which suggests renting virtual machines from a public cloud for running parallel applications. Since HPC in cloud workloads are tightly coupled, we exclude them from HVC. Moreover, in the context of a business model, no one formally defines the metrics for evaluating cloud from perspectives of both hardware and software systems.

C. Discussion

Despite several computing paradigms are not formally or clearly defined, e.g., DISC, WSC, and cloud computing, we characterize several new computing paradigms and compare the definition of our HVC from them.

We draw three conclusions in this Section: first, the definition of HVC is towards data center computer systems, while many task computing, high throughput computing, or DISC are defined towards runtime systems. Second, in terms of workloads and respective metrics, both high throughput computing and high performance computing are about scientific computing centering around floating point operations, while most of HVC applications have few floating point operations. Third, many emerging workloads can be included into one or two categories of HVC workloads e.g., WSC (into the first and second categories), DISC or data center computing (into the second category); With the exception of HPC in cloud workloads[1], well-know workloads in cloud studied in [5] can be included into HVC. Moreover, (public) cloud computing is basically a business model of renting computing or storage resources, while our HVC is defined in terms of workloads.

III. BENCHMARKING HVC SYSTEMS

In this section, we revisit previous benchmarks, and present our current work on HVC metrics and benchmarks.

A. Revisiting previous benchmarks and metrics

Table II summarizes different benchmarks and their respective levels, workloads, and metrics.

The LINPACK benchmark report [7] describes the performance for solving a general dense matrix problem $Ax = b$. Performance is often measured in terms of floating point operations per second (flop/s). Focusing on evaluating the ability of traversing the whole memory of the machine [20], Murphy *et al.* have put together a benchmark they call Graph 500, and the metric is the traversed edges per second (TEPS). Since energy efficiency becomes more and more important, the Green500

List ranks supercomputers used primarily for scientific production codes according to the amount of power needed to complete a fixed amount of work [17]. Sun Microsystems also proposed the SWaP (space, watts, and performance) metric [24] to evaluate enterprise systems from perspectives of both data center space efficiency and power consumption.

Rivoire *et al.* [18] proposed the JouleSort benchmark. The external sort from the sort benchmark specification (<http://research.microsoft.com/research/barc/SortBenchmark/default.htm>) is chosen for the benchmarks workload. The metrics is records sorted per Joule. SPECpower_ssj2008 is the first industry-standard SPEC benchmark that evaluates the power and performance characteristics of volume server class and multi-node class computers [36]. The initial benchmark addresses only the performance of server side Java—SPECjbb2005. The metric is performance per watt metric in terms of `ssj_ops/watt`.

TPC defined transaction processing and database benchmarks, the goal of which is to define a set of functional requirements that can be run on any transaction processing system, regardless of hardware or operating system [35]. Most of the TPC benchmarks are obsolete, and only three are still used: TPC-C, TPC-E, and TPC-H. TPC-C centers around the principal activities (transactions) of an order-entry environment [35]. TPC-E models a brokerage firm with customers who generate transactions related to trades, account inquiries, and market research [35]. Different from TPC-C and TPC-E, TPC-H models the analysis end of the business environment where trends are computed and refined data are produced to support the making of sound business decisions [35]. For the TPC benchmarks, the metrics are application-specific. For example, the metric of TPC-C is the number of New-Order transactions executed per minute.

In the context of data center computing, Hi-Bench [13], GridMix2 or GridMix 3 [39], WL Suite [9] is proposed to evaluate MapReduce runtime systems, respectively. The workloads are data analysis applications. The metrics are throughput

in terms of the number of tasks per minute [13], and job running time, widely used in batch queuing systems [16]. YCSB [19] and an extension benchmark—YCSB++ [10] are proposed to evaluate NoSQL systems for scale-out data services. The metrics are throughput—total operations per second, including reads and writes, and average latency per requests.

PARSEC is a benchmark suite for studies of Chip-Multiprocessors (CMPs) [11]. PARSEC includes emerging applications in recognition, mining and synthesis (RMS) as well as systems applications which mimic large-scale multi-threaded commercial programs. SPEC CPU2006 provides a snapshot of current scientific and engineering applications, including a suite of serial programs that is not intended for studies of parallel machines [11]. SPEC CPU2006 component suite include both CINT2006—the Integer Benchmarks and CFP2006—the Floating Point Benchmarks. After the benchmarks are run on the system under test (SUT), a ratio for each of them is calculated using the run time on the system under test and a SPEC-determined reference time [36].

SPEC also proposed a series of benchmarks for Java applications [36]. Among them, SPECjvm2008 is a client JVM benchmark; SPECjbb2005 is a server JVM benchmark, while SPECjEnterprise2010 is a Java enterprise edition application server benchmark. SPECjms2007 is the first industry-standard benchmark for evaluating the performance of enterprise message-oriented middleware servers based on JMS (Java Message Service). SPECweb2009 emulates users sending browser requests over broadband Internet connections to a web server using both HTTP and HTTPS. It provides banking, e-commerce, and support workloads, along with a new power workload based on the e-commerce workload [36]. SPECsip_Infrastructure2011 is designed to evaluate a system’s ability to act as a SIP server supporting a particular SIP application [36]. The application is modeled after a VoIP deployment [36]. The metric is the simultaneous number of supported subscribers [36].

1) discussion: We draw four conclusions in this subsection: first, there is no systematic work on benchmarking throughput-oriented workloads in data centers. In Section I, we have identified three categories of HVC workloads. Few previous benchmarks pay attentions to all three categories of applications.

Second, some benchmarking efforts [13] [9] have focused on MapReduce-based data analysis applications (belong to our second category of HVC workloads), and the metric is toward evaluating MapReduce runtime systems. They failed to notice that there are diverse programming models in this field, e.g., Dryad, AllPair [2], since MapReduce is not a one-size-fits-all solution [2]. For example, MapReduce or Dryad is not appropriate for applications such as iterative jobs, nested parallelism, and irregular parallelism [2].

Third, a lot of previous benchmarks, e.g., TPC or SPEC efforts, pay attention to services. Unfortunately, emerging data-intensive services, such as Web search engines are ignored.

Last, in addition to SPECsip_Infrastructure2011, little previous work focuses on interactive real-time applications, which are important emerging workloads, since more and more users tend to use streaming media or VoIP for fun or communications.

B. Our ongoing work on metrics and benchmarks.

Benchmarking is the foundation of evaluating HVC systems. However, to be relevant, an HVC benchmark suite needs to satisfy a number of properties as follows:

First, the applications in the suite should consider a target class of machines [11], that is a data center system designed for throughput-oriented workloads, not a processor or server in this paper.

Second, the HVC benchmark suite should represent three categories of important applications, including services, data processing applications, and interactive real-time applications, which is the right target of our ongoing DCBenchmarks project (<http://prof.ncic.ac.cn/DCBenchmarks>). In this project, we have released a SEARCH benchmark [3]. We will release a benchmark for shared

data center computer systems running data processing applications soon.

Third, the workloads in the HVC benchmark suite should be diverse enough to exhibit the range of behavior of the target applications [11]. Meanwhile, since service providers may deploy different applications, it is important for a service provider to customize their chosen benchmarks relevant to their applications.

Fourth, no single metric can measure the performance of computer systems on all applications [8]. Since different categories of workloads in HVC have different focuses, we propose different metrics. For each category of workloads, we propose an aggregate metric which measures a data center system on the whole, and an auxiliary metric to evaluate energy efficiency, which can be measured on the level of not only a server but also a data center system. As shown in Table I, for *services*, we propose *requests per minute* as an aggregate metric, and *requests per joule* as an auxiliary metric evaluating energy efficiency. For *data processing applications*, we propose *data processed per minute* as an aggregate metric and *data processed per joule* as an auxiliary metric. For interactive real-time applications, we propose *maximum number of simultaneous subscribers* as an aggregate metrics and *subscribers per watt*—*a ratio of the maximum number of simultaneous subscribers to the power consumption in a unit time* as an auxiliary metric. A subscribe here can be a user or device. Due to the space limitation, we will report the experiment results in detail in another paper.

We also developed several innovative performance analysis tools [40] [41] [42], aiding with understanding HVC workloads.

IV. CONCLUSION

There are four-fold contributions in this paper. For the first time, we systematically characterized HVC systems and identified three categories of HVC workloads: services, data processing or interactive real-time applications; we compared HVC with other computing paradigms in terms of six dimensions: levels, workloads, metrics, coupling

degree, data scales, and number of jobs or service instances; we widely investigated previous benchmarks, and found there is no systematic work on benchmarking HVC systems; we presented our preliminary work on HVC metrics and benchmarks.

ACKNOWLEDGMENT

We are very grateful to anonymous reviewers. This work is supported by the Chinese 973 project (Grant No.2011CB302500) and the NSFC project (Grant No.60933003).

REFERENCES

- [1] L. Wang *et al.* 2012. In Cloud, Can Scientific Communities Benefit from the Economies of Scale?. IEEE TPDS. vol.23, no.2, pp.296-303, Feb. 2012.
- [2] P. Wang *et al.* 2010. Transformer: A New Paradigm for Building Data-Parallel Programming Models. IEEE Micro 30, 4 (Jul. 2010), 55-64.
- [3] H. Xi *et al.* 2011. Characterization of Real Workloads of Web Search Engines. In Proc. IISWC 2011.
- [4] U. Hoelzle *et al.* 2009. The Datacenter as a Computer: an Introduction to the Design of Warehouse-Scale Machines. 1st. Morgan and Claypool Publishers.
- [5] M. Ferdman *et al.* 2012. Clearing the Clouds: A Study of Emerging Workloads on Modern Hardware. In Proc. ASPLOS 2012.
- [6] D. J. DeWitt. 1993. The Wisconsin Benchmark: Past, present and future. In J. Gray, editor, The Benchmark Handbook. Morgan Kaufmann, 1993.
- [7] J. J. Dongarra *et al.* 2003. The LINPACK Benchmark: past, present and future. Concurrency and Computation: Practice and Experience, 15: 803-820.
- [8] J. Gray, editor. 1993. The Benchmark Handbook For Database and Transaction Processing Systems. Morgan Kaufmann, 1993.
- [9] Y. Chen *et al.* 2011. The Case for Evaluating MapReduce Performance Using Workload Suites. In Proc. MASCOTS, 2011.
- [10] S. Patil *et al.* 2011. YCSB++: benchmarking and performance debugging advanced features in scalable table stores. In Proc. SoCC, 2011.

Table II: Comparison of different benchmarks and metrics.

<i>Benchmark</i>	<i>Domains</i>	<i>Level</i>	<i>Workloads</i>	<i>Metrics</i>
<i>Linpack</i> [7]	High performance computing	Super computers	scientific computing code	Float point operations per second
<i>SWap</i> [24]	Enterprise	Systems	Undefined	Performance/(space * watts)
<i>Green 500</i> [17]	High performance computing	Super computers	Scientific computing code	Flops per watt
<i>Graph 500</i> [20]	High performance computing	Super computers	Computing in the field of graph theory	Traversed edges per second (TEPS)
<i>JouleSort</i> [18]	Mobile, desktop, enterprise	Systems	External sort	Records sorted per Joule
<i>SPECpower _ssj2008</i> [36]	Enterprise	Systems	SPECjbb2005	Ssj_ops/watt
[22]	Storage I/O	Storage systems	Transaction processing or scientific applications	I/O or data rates
<i>SPECsfs2008</i>	Network file systms	File servers	N/a	Operations per second and overall latency of the operations
<i>HiBench</i> [13]	Data-intensive scalable computing	MapReduce run-time systems	Data analysis	Job running time and number of tasks completed per minute
<i>GridMix2 or Grid-Mix3</i> [39]	Data-intensive scalable computing	MapReduce run-time systems	Data analysis	Number of completed jobs and running time
<i>WL Suite</i> [9]	Data-intensive scalable computing	MapReduce run-time systems	Data Analysis	n/a
<i>YCSB</i> [19] or <i>YCSB++</i> [10]	Warehouse-scale computing	NoSQL systems	Scale-out data services	Total operations per second and average latency per requests
<i>PARSEC</i> [11]	n/a	Chip-Multiprocessors	Recognition, mining, synthesis, and mimic large-scale multithreaded commercial programs	n/a
<i>TPC C/E/H</i> [23]	Throughput-oriented workloads	Server systems	Transaction processing and decision support	Application-specific
<i>SPEC CPU2006</i> [36]	scientific and engineering applications	Processors	Serial programs	A ratio is calculated using the run time on the system under test and a SPEC-determined reference time
<i>SPECjvm2008, SPECjbb2005, SPECjEnterprise2010, SPECjms2007</i> [36]	Throughput-oriented workloads	Both hardware and software	Java applications	Throughput (application-specific)
<i>SPECsip _Infrastructure2011</i> [36]	Throughput-oriented workloads	Systems	A VoIP deployment	Simultaneous number of supported subscribers
<i>SPECvirt _sc2010</i> [36]	Throughput-oriented workloads	Systems	SPECweb2005, SPECjAppServer2004, and SPECmail2008	Performance only and performance per watt
<i>SPECweb2009</i> [36]	Throughput-oriented workloads	Systems	Banking, Ecommerce, and Support	Maximum number of simultaneous user sessions, ratio of the sum of simultaneous user sessions to the sum of watts used

- [11] C. Bienia. 2011. Benchmarking Modern Multiprocessors. PhD thesis, Princeton University, January 2011.
- [12] J. Laudon. 2005. Performance/Watt: the new server focus. SIGARCH Comput. Archit. News 33, 4 (November 2005), 5-13.
- [13] S. Huang *et al.* 2010. The HiBench benchmark suite: Characterization of the MapReduce-based data analysis. In Proc. ICDEW 2010.
- [14] I. Raicu *et al.* 2008. Many-task computing for grids and supercomputers, In Proc. of MTAGS 2008.
- [15] M. Armbrust *et al.* 2009. Above the clouds: A Berkeley view of cloud computing, Tech. Rep. UCB/EECS-2009-28.
- [16] K. Gaj *et al.* 2002. Performance evaluation of selected job management systems, In Proc. IPDPS 2002.
- [17] W. Feng *et al.* 2007. The Green500 List: Encouraging Sustainable Supercomputing. Computer, vol.40, no.12, pp.50-55, Dec. 2007.
- [18] S. Rivoire *et al.* 2007. Models and Metrics to Enable Energy-Efficiency Optimizations. Computer, vol.40, no.12, pp.39-48, Dec. 2007.
- [19] B. F. Cooper *et al.* 2010. Benchmarking cloud serving systems with YCSB. In Proc. SoCC '10, pp.143-154.
- [20] M. Anderson. 2011. Better benchmarking for supercomputers. Spectrum, IEEE , vol.48, no.1, pp.12-14, Jan. 2011.
- [21] M. Michael *et al.* 2007. Scale-up x scale-out: A case study using nutch/lucene. In Proc. IPDPS 2007.
- [22] P.M. Chen *et al.* 1993. Storage performance-metrics and benchmarks. In Proc. of the IEEE , vol.81, no.8, pp.1151-1165, Aug 1993.
- [23] M. Poess *et al.* 2000. New TPC benchmarks for decision support and web commerce. SIGMOD Rec. 29, 4 (December 2000), 64-71.
- [24] Sun Microsystems, SWaP (Space, Watts, and Performance) Metric; www.sun.com/servers/coolthreads/swap.
- [25] D. Thain, T. Tannenbaum, M. Livny, Distributed Computing in Practice: The Condor Experience, Concurrency and Computation: Practice and Experience 17(2-4), pp. 323-356, 2005
- [26] M. Livny *et al.* 1997. Mechanisms for High Throughput Computing, SPEEDUP Journal 1(1), 1997
- [27] S. Chaudhry *et al.* High-performance throughput computing. Micro, IEEE , vol.25, no.3, pp. 32-45, May-June 2005
- [28] R. E Bryant. 2008. Data intensive scalable computing. Retrieved January 5, 2010, from: <http://www.cs.cmu.edu/bryant/presentations/DISC-concept.ppt>.
- [29] G. Ananthanarayanan *et al.* 2011. Disk-locality in datacenter computing considered irrelevant. In Proc. HotOS 2011.
- [30] J. Hamilton. 2009. Internet-scale service infrastructure efficiency. SIGARCH Comput. Archit. News 37, 3 (Jun. 2009), 232-232.
- [31] D. Gannon. 2011. Challenges for large scale scientific data analytics. Keynote given at PDAC-11, in conjunction with SC'11.
- [32] M. Isard *et al.* 2009. Quincy: fair scheduling for distributed computing clusters. In Proc. of SOSP 09. ACM, New York, NY, 261-276.
- [33] K. Beaty *et al.* 2009. Desktop to cloud transformation planning. In Proc. IPDPS 2009.
- [34] L. M. Vaquero *et al.* 2008. A break in the clouds: towards a cloud definition. SIGCOMM Comput. Commun. Rev. 39, 1 (Dec. 2008), 50-55.
- [35] TPC BENCHMARK : <http://www.tpc.org/>
- [36] Standard Performance Evaluation Corporation: <http://www.spec.org/>
- [37] Distributed stream computing platform: <http://incubator.apache.org/s4/>
- [38] E. Deelman *et al.* 2008. The cost of doing science on the cloud: the montage example, In Proc. of SC 08.
- [39] GridMix: <http://hadoop.apache.org/mapreduce/docs/current/gridmix.html>
- [40] X. Liu *et al.* 2011. Automatic performance debugging of SPMD-style parallel programs. J. Parallel Distrib. Comput. 71, 7 (July 2011), 925-937.

- [41] Z. Zhang *et al.* 2009. Precise request tracing and performance debugging for multi-tier services of black boxes. In Proc. DSN 2009, 337-346.
- [42] B. Sang *et al.* 2012. Precise, scalable, and online request tracing for multi-tier services of black boxes. Accepted by TPDS.