Portland State University

PDXScholar

Engineering and Technology Management
Faculty Publications and Presentations

Engineering and Technology Management

10-2014

Technological Forecasting of Supercomputer Development: The March to Exascale Computing

Dong-Joon Lim
Portland State University

Timothy R. Anderson
Portland State University, tim.anderson@pdx.edu

Tom Shott
Portland State University, tom.n.pdx@gmail.com

Follow this and additional works at: https://pdxscholar.library.pdx.edu/etm_fac

Part of the [Engineering Commons](#)

Let us know how access to this document benefits you.

Citation Details

Lim, Dong-Joon; Anderson, Timothy R.; and Shott, Tom, "Technological Forecasting of Supercomputer Development: The March to Exascale Computing" (2014). *Engineering and Technology Management Faculty Publications and Presentations*. 46.

https://pdxscholar.library.pdx.edu/etm_fac/46

This Post-Print is brought to you for free and open access. It has been accepted for inclusion in Engineering and Technology Management Faculty Publications and Presentations by an authorized administrator of PDXScholar. Please contact us if we can make this document more accessible: pdxscholar@pdx.edu.

Technological forecasting of supercomputer development: The march to Exascale computing

Dong-Joon Lim*, Timothy R. Anderson, Tom Shott

Dept. of Engineering and Technology Management, Portland State University, USA

Abstract- Advances in supercomputers have come at a steady pace over the past 20 years. The next milestone is to build an Exascale computer however this requires not only speed improvement but also significant enhancements for energy efficiency and massive parallelism. This paper examines technological progress of supercomputer development to identify the innovative potential of three leading technology paths toward Exascale development: hybrid system, multicore system and manycore system. Performance measurement and rate of change calculation were made by Technology Forecasting using Data Envelopment Analysis (TFDEA.) The results indicate that the current level of technology and rate of progress can achieve Exascale performance between early 2021 and late 2022 as either hybrid systems or manycore systems.

1. Introduction

Supercomputers have played a critical role in various fields which require computationally intensive tasks such as pharmaceutical test, genomics research, climate simulation, energy exploration, molecular modeling, astrophysical recreation, etc. The unquenchable need for faster and higher precision analysis in those fields create the demand for even more powerful supercomputers. Furthermore, developing indigenous supercomputer has become a fierce international competition due to its role as a strategic asset for a nationwide scientific research and the prestige of being the maker of the fastest computers [1,2]. While the vast majority of

supercomputers have still been built using processors from Intel, Advanced Micro Devices (AMD), and NVidia, manufacturers are committed to developing their own customized systems, e.g. interconnect, operating system and resource management, as system optimization becomes a crucial factor in today's massively parallel computing paradigm [3].

Advances in supercomputers have come at a steady pace over the past 20 years in terms of speed, which has been enabled by the continual improvement in computer chip manufacturing [4]. The world's fastest supercomputer today (March 2014) is the Tianhe-2 built by China's National University of Defense Technology (NUDT) performing at over 33.86 Petaflops, i.e. 33.86×10^{15} floating point operations per second. This is about 273,000 times faster than the fastest machine 20 years ago, the Fujitsu Numerical Wind Tunnel. On average, progress went from being measured by Gigaflops in 1990 to Teraflops in about 10 years, and then to Petaflops in another 10 years [5]. In line with this, the next milestone to build an Exascale computer, a machine capable of doing a quintillion operations, i.e. 10^{18} , per second had been projected to see light of day by 2018 [6]. However, there are significant industry concerns that this incremental improvement might not continue mainly due to several practical problems.

The biggest challenge to build the Exascale computer is the power consumption [7]. Tianhe-2, which is currently not only the fastest but also the largest supercomputer, uses about 18 megawatts (MW) of power. If the current trend of power use continues, projections for the Exascale computing systems range from 60 to 130 MW which would cost up to \$150 million annually [8]. Therefore, unlike past advancement mainly driven by performance improvement [9], energy efficiency has now gone from being a negligible factor to a fundamental design consideration. Practically, fewer sites in the U.S. will be able to host the Exascale computing systems due to limited availability of facilities with sufficient power and cooling capabilities [10].

To cope with these issues, current efforts are targeting the Exascale machine that draws electrical power of 20 MW using 100 million processors in the 2020 timeframe [7,11].

Given the fact that the Exascale computing may require a vast amount of power and massive parallelism, it is crucial to incorporate the power efficiency and multicore characteristics into the measure of technology assessment to have a correct view of current trend [12]. This implies that the extrapolation relying on a single performance measure, i.e. computing speed, may overlook required features of future technology systems and could eventually result in an erroneous forecast. Specifically, the average power efficiency of today's top 10 systems is about 2.2 Petaflops per megawatt. This indicates that it is required to improve power efficiency by a factor of 23 to achieve the Exascale goal. The projection of performance development therefore may have to be adjusted to consider structural and functional challenges involved. This manifestly requires multifaceted approach to investigate the tradeoffs between system attributes, which can tackle the questions such as: how much performance improvement would be restricted by power and/or core reduction? What would be the maximum attainable computing performance with certain levels of power consumption and/or the number of cores?

There are three leading technology paths representing today's supercomputer development: hybrid systems, multicore systems, and manycore systems [13]. The hybrid systems use both central processing units (CPU) and graphics processing units (GPU) to efficiently leverage the performances [14]. The multicore systems maintain a number of complex cores whereas the manycore systems use a large number of less powerful but power efficient cores within the highly-parallel architecture [15]. Manufacturers and researchers are exploring these alternate paths to identify the most promising, namely energy efficient and performance effective, avenue to face challenges of deploying and managing Exascale systems [16–18]. The comparative

analysis on these technology paths can, therefore, give insights into the estimation of the future performance levels as well as the possible disruptive technology changes.

This study employs Technology Forecasting using Data Envelopment Analysis (TFDEA) to measure the technological progress considering tradeoffs among power consumption, multicore processors, and maximum performance so that supercomputers are to be evaluated in terms of both energy efficiency and performance effectiveness. The resulting analysis then provides a forecast of Exascale computer deployment under three different development alternatives in consideration of current business environment as well as emerging technologies.

2. Methodology

Frontier analysis (or best practice) methods that model the frontier of data points rather than model the average possibilities have become popular in modern benchmarking studies [19–22]. As an example, TFDEA has shown its usefulness in a wide range of applications since the first introduction in *PICMET '01* [23–27]. This approach has a strong advantage in capturing technological advancement from the State of the Arts (SOAs) rather than being influenced by the inclusion of mediocre technologies that is frequently observed in the central tendency model [28].

One of the favorable characteristics of TFDEA is its flexibility that can incorporate practical views in the assessment. For example, DEA, which underlies TFDEA, allows dynamic weighting scheme that the model gives freedom to each data point to select its own weights, and as such, the efficiency measure will show it in the best possible light [29,30]. This approach has an advantage to prevent the model from underestimating various types of technologies based on *a priori* fixed weighing scheme. Furthermore, DEA generates a reference set that can be used as reasonable benchmarks for each data point to improve its performance [31]. This allows TFDEA

to capture the technological advancement with consideration of various tradeoffs from different model parameters: orientation, returns to scale, weight restrictions along with variable selections [32].

In addition, TFDEA can deal with multiple variables, i.e. system attributes, thereby tracking efficiency changes over time. Specifically, TFDEA first measures the efficiency of each technology and then analyzes time series efficiency changes to explain the expansion of SOA frontier. This approach has a strong advantage by taking various tradeoffs into account whereas most extrapolation methods can only explain the variance of single dependent variable at a time by directly relating it to independent variables [33,34].

Although time series application of benchmarking practice can shed light on the new product development planning phase, there remains a need to integrate the product positioning with the assessment of technological progress so that analysts can investigate different product segments with the purpose of market research. In particular, the High Performance Computing (HPC) industry has important niches with segmented levels of competition from small-scale multiprocessing computers to the mainframe computers. This necessarily requires an identification of technological advancement suitable for the design target of Exascale computer from corresponding product segments. For this reason, this study presents a new approach which is capable of considering the variable rates of technological advancement from different product segments.

The whole process of the proposed model can be divided into three computational stages. In the first stage, (1)-(7), the notation x_{ij} represents the i th input and y_{rj} represents the r th output of technology j and $j=k$ identifies the technology being evaluated. The variables for the linear program underlying DEA are $\lambda_{jk}^{h \in \{R,C\}}$ and $\phi_k^{h \in \{R,C\}}$. The variable $\phi_k^{h \in \{R,C\}}$ represents the

proportion of output that should be generated to become an SOA at time period R (release time) or at time period C (current frontier) to actual output that technology k produced. Since each reference set, λ_{jk}^R , only includes technologies that had been released up to t_k by constraint (4), ϕ_k^R indicates how superior the technology k is at the time of release. Similarly, constraint (5) restrict each reference set, λ_{jk}^C , to include technologies that had been released up to current frontier time T , therefore, ϕ_k^C measures the amount by which technology k is surpassed by the current SOA frontier. The objective function (1) also incorporates effective date minimization to ensure reproducible results from possible alternate optimal solutions [35]. The variable returns to scale (VRS) are enforced by constraint (6).

$$\max \sum_{k=1}^n \left[\phi_k^h - \varepsilon \left(\frac{\sum_{j=1}^n \lambda_{jk}^h \cdot t_j}{\sum_{j=1}^n \lambda_{jk}^h} \right) \right] \quad (1)$$

$$s. t. \sum_{j=1}^n \lambda_{jk}^h \cdot y_{rj} \geq \phi_k^h \cdot y_{rk}, \quad r = 1, \dots, s \quad (2)$$

$$s. t. \sum_{j=1}^n \lambda_{jk}^h \cdot x_{ij} \leq x_{ik}, \quad i = 1, \dots, m \quad (3)$$

$$s. t. \lambda_{jk}^R = 0, \quad \forall (j, k) | t_j > t_k \quad (4)$$

$$s. t. \lambda_{jk}^C = 0, \quad \forall (j, k) | t_j > T \quad (5)$$

$$s. t. \sum_{j=1}^k \lambda_{jk}^h = 1, \quad \forall k \quad (6)$$

$$s. t. \lambda_{jk}^h \geq 0, \quad \forall j, k, h \in \{R, C\} \quad (7)$$

Once efficiency measurements both at time of release (R) and at time of current frontier (C) are completed, the rate of change (RoC), γ_k^C , may then be calculated in (8) by taking all

technologies that were efficient at time of release, $\phi_k^{R*} = 1$, but were superseded by technology at current frontier, $\phi_k^{C*} > 1$. The local RoCs, δ_j^C , can be obtained in (9) by taking the weighted average of RoCs for each technology on the current frontier. Each local RoC therefore represents a growth potential of adjacent frontier facets based on the technological advancement observed from related past products. Consequently, the local RoC enables the model to identify an individualized RoC under which each forecasting target is expected to arrive [36,37]. Note that the traditional TFDEA model makes a forecast based on a single aggregated RoC, i.e. average of γ_k^C , without consideration of the unique growth patterns of different product segments.

$$\gamma_k^C = (\phi_k^{C*})^{\frac{1}{\sum_{j=1}^n \lambda_{j,k}^{C*} \cdot t_j - t_k}}, \quad \forall k \mid \phi_k^{R*} = 1, \phi_k^{C*} > 1 \quad (8)$$

$$\delta_j^C = \frac{\sum_{k=1}^n \lambda_{j,k}^{C*} \cdot \gamma_k^C}{\sum_{k=1, \gamma_k^C > 0}^n \lambda_{j,k}^{C*}}, \quad \forall j \mid \phi_k^{C*} = 1 \quad (9)$$

The forecasting process is denoted in (10) where ϕ_k^{C*} indicates super-efficiency measuring how much the future technology k outperforms SOA technologies on the current frontier. The individualized RoC for each forecasting target k can be computed by combining the local RoCs of SOA technology j that constitutes the frontier facet onto which technology k is being projected. The forecasted time $t_k^{forecast}$ is, therefore, obtained by the sum of estimated elapsed time and the effective date for the projection.

$$t_k^{forecast} = \frac{\ln\left(\frac{1}{\phi_k^{C*}}\right)}{\ln\left(\frac{\sum_{j=1}^n \lambda_{j,k}^{C*} \cdot \delta_j^C}{\sum_{j=1}^n \lambda_{j,k}^{C*}}\right)} + \frac{\sum_{j=1}^n \lambda_{j,k}^{C*} \cdot t_j}{\sum_{j=1}^n \lambda_{j,k}^{C*}}, \quad \forall k \mid t_k > T \quad (10)$$

3. Analysis

3.1. Dataset

The TOP500 list was first created in 1993 to assemble and maintain a list of the 500 most powerful computer systems [38]. Since the list has been compiled twice a year, datasets from 1993 to 2013 have been combined and cleaned up so that each machine appears once in the final dataset. The purpose of this study is to consider both energy efficiency and performance effectiveness, therefore lists up to 2007 were excluded due to the lack of information on the power consumption (see table 1.) Variables selected for this study are as follows:

- Name (text): name of machine
- Year (year): year of installation/last major update
- Total Cores (number): number of processors
- Rmax (Gigaflops): maximal LINPACK performance achieved
- Power (Kilowatts): power consumption
- Interconnect family (text): interconnect being used
- Processor technology/family (text): processor architecture being used

In the final dataset, there were total 1,199 machines, with number of cores ranging from 960 to 3.12 million, power ranging from 19KW to 17.81MW, Rmax ranging from 9 Teraflops to 33.86 Petaflops from 2002 to 2013. Note that logarithmic transformation was applied to three variables prior to the analysis due to their exponentially increasing trends.

Table 1 TOP500 dataset from 1993 to 2013

Data column	1993 ~2007	2008 ~2009	2010-1	2010-2	2011-1	2011-2 ~2013
Rank	O	O	O	O	O	O
Site	O	O	O	O	O	O
Manufacturer	O	O	O	O	O	O
Name	x	x	x	x	x	O
Computer	O	O	O	O	O	O
Country	O	O	O	O	O	O
Year	O	O	O	O	O	O
Total Cores	O	O	O	O	O	O
Accelerator/Co-Processor Cores	x	x	x	x	x	O
Rmax	O	O	O	O	O	O
Rpeak	O	O	O	O	O	O
Efficiency (%)	x	x	x	x	x	O
Nmax	O	O	O	O	O	O
Nhalf	O	O	O	O	O	O
Power	x	O	O	O	O	O
Mflops/Watt	x	x	x	x	x	O
Measured Size	x	x	O	x	x	x
Processor Technology	O	O	O	O	O	O
Processor Generation	x	x	x	x	x	O
Processor	O	O	O	O	O	O
Proc. Frequency	O	O	O	O	O	x
Processor Cores	x	x	O	O	O	x
Processor Speed (MHz)	x	x	x	x	x	O
System Family	O	O	O	O	O	O
System Model	x	O	O	O	O	O
Operating System	O	O	O	O	O	O
OS Family	x	x	x	x	x	O
Cores per Socket	x	x	x	x	x	O
Architecture	O	O	O	O	O	O
Accelerator/Co-Processor	x	x	x	x	O	O
Segment	O	O	O	O	O	O
Application Area	O	O	O	O	O	x
Interconnect Family	O	O	O	O	O	O
Interconnect	O	O	O	O	O	O
Region	O	O	O	O	O	O
Continent	O	O	O	O	O	O

O: Available, x: Unavailable

3.2. Model building

The TFDEA model has been implemented using the software developed by Lim and Anderson [39]. As discussed earlier, power consumption and the number of cores were used as input variables and the maximum LINPACK performance (Rmax) was used as the output

variable. This allows the model to identify ‘the better performing’ supercomputer which has lower power, fewer cores, and higher performance. Orientation can be either input-oriented or output-oriented and can be best thought of as whether the technological progress is better characterized as “input reduction” or “output augmentation [40].” While power consumption will be a key concern in the Exascale computing, the advancement of this industry has been driven primarily by computing performance, i.e. flops, improvement. Besides, the Exascale computing is a clearly defined development goal therefore an output orientation was selected for this application. It should be noted here that either orientation can deal with tradeoffs among input and output variables. As with many DEA applications, variable returns to scale (VRS) was selected for appropriate returns to scale assumption since doubling the input(s) doesn’t correspond to doubling the output(s) here as well. The main purpose of this study is to make a forecast of the Exascale computer deployment by examining past rate of progress, thus the frontier year of 2013 was used so as to cover the whole dataset. Lastly, minimizing the sum of effective dates was added as a secondary goal into the model to handle the potential issue of multiple optima from the dynamic frontier year [41]. Table 2 summarizes the model parameters used in this study.

Table 2 TFDEA model parameters

Inputs	Output	Orientation	RTS	Frontier year	Frontier type	Second goal
Power, Cores	Rmax	Output	VRS	2013	Dynamic	Min

Figure 1 shows thirteen supercomputers identified as SOAs from the analysis. Intel provided the processors for the largest share (62%) and, *inter alia*, GPU/Accelerator based systems dominated both energy and core efficient systems, while IBM’s Blue Gene, NNSA/SC and Blue

Gene/Q, showed comparable energy efficiency as manycore based systems. As also seen from specifications of these supercomputers in Table 3, supercomputers are characterized as being competitive in consideration of tradeoffs within different scale sizes. This enables the model to construct technology frontiers from which various production possibilities can be identified. This characteristic, in fact, differentiates the TFDEA process from a single dimensional measure such as the TOP500 list in which technological efforts to become energy efficient and/or core efficient are not taken into account.

Figure 1 13 State of the art supercomputers considering system tradeoffs

Table 3 Specifications of 13 state of the art supercomputers considering system tradeoffs

Name	Date	Cores	Power	Rmax	Interconnect	Processor Family
Eurora Eurotech Aurora HPC 10-20	2013	2,688	46.00	100,900	InfiniBand	Intel
Tianhe-2 TH-IVB-FEP	2013	3,120,000	17,808.00	33,862,700	Custom	Intel
HPCC	2013	10,920	237.00	531,600	InfiniBand	Intel
Titan Cray XK7	2012	560,640	8,209.00	17,590,000	Cray	AMD
Beacon Appro GreenBlade GB824M	2012	9,216	45.11	110,500	InfiniBand	Intel
BlueGene/Q, Power BQC 16C 1.60GHz	2012	8,192	41.09	86,346	Custom	IBM Power
iDataPlex DX360M3	2011	3,072	160.00	142,700	InfiniBand	Intel
NNSA/SC Blue Gene/Q Prototype 2	2011	8,192	40.95	85,880	Custom	IBM Power
DEGIMA Cluster	2011	7,920	31.13	42,830	InfiniBand	Intel
BladeCenter QS22 Cluster	2008	1,260	18.97	9,259	InfiniBand	IBM Power
Cluster Platform 3000 BL2x220	2008	1,024	42.60	9,669	InfiniBand	Intel
Power 575, p6 4.7 GHz, Infiniband	2008	960	153.43	14,669	InfiniBand	IBM Power
BladeCenter HS21 Cluster	2007	960	91.55	9,058	InfiniBand	Intel

Figure 2 illustrates performance trajectories based on 1,199 supercomputers from three dominant processor families: AMD X86, IBM Power, and Intel (IA-32/64, Core, Nehalem, Westmere, and Sandy Bridge.) Since the Japanese supercomputer, Earth-Simulator, in 2002 was built using a Nippon Electric Company (NEC) chip which was not adopted by other supercomputer manufacturers thereafter, Fig. 2 is drawn from 2005 to focus on main vendors of processor for today's systems. The ordinate is the overall performance score from the DEA model. As such, each line indicates performance trajectory of the top performing supercomputers from each year against the frontier year of 2013. That is, a performance score of 100% indicates that the supercomputer has a superior performance enough to be on the SOA frontier in 2013. A

performance score higher than 100% denotes super-efficiency from the DEA model which can show how much the supercomputer is outperforming other SOA supercomputers.

Figure 2 Performance trajectories of different processor families

The trajectory of Many/Multicore systems shows that IBM Power (PC) processor based machines are outperforming AMD X86 processor based machines. AMD X86 processor based machines however showed surpassing performances over IBM Power (PC) based machines when they were adopted by Cray to build hybrid systems in 2011 and 2012. In fact, the successful development of Titan Cray XK7 using AMD Opteron CPUs coupled with NVidia coprocessors has made Cray Inc. one of the leading supercomputer vendors to date. Interestingly, this is also consistent with the fact that Cray Inc. was awarded the \$188M U.S. Blue Waters

contract, which is a project funded by National Science Foundation (NSF), replacing IBM which had pulled out of the project prior to completion in 2011 [42].

It is also interesting to point out that the performance gap between Many/Multicore based machines and GPU/Accelerator based machines is larger in supercomputers using AMD X86 processors than Intel processors. This can be attributed to the partnership between Cray and AMD. In fact, Cray has been a staunch supporter of AMD processors since 2007 and their collaboration has delivered continued advancement in HPC [43]. In particular, Cray's recent interconnect technology, Gemini, was customized for the AMD Opteron CPUs Hyper-Transport links to optimize internal bandwidth [44]. Since modern supercomputers are deployed as massively centralized parallel systems, the speed and flexibility of interconnect becomes important for the overall performance of supercomputer. Given that hybrid machines using AMD X86 processors all use Cray's interconnect system, one may notice that AMD X86 based supercomputers had a significant performance contribution from Cray interconnect as well as NVidia coprocessors.

One may notice that top supercomputers based on Intel processors have switched to hybrid systems since 2010. This is because combining CPUs and GPUs is advantageous in data parallelism which makes it possible to balance the workload distribution as efficient use of computing resources becomes more important in today's HPC structure [45]. Hybrid machines using Intel processors have all adopted InfiniBand interconnect for their cluster architectures regardless of GPUs/Accelerators; NVidia, ATI Radeon, Xeon Phi, PowerXCell, etc. InfiniBand, manufactured by Mellanox and Intel, enables low processing overhead and is ideal to carry multiple traffic types such as clustering, communications, and storage over a single connection [46]. Especially, its GPU-Direct technology facilitates faster communication and lower latency of GPU/Accelerator based systems that can maximize computing and accelerator resources, as

well as improves productivity and scalable performance [47]. Intel acquired the InfiniBand business from Qlogic in 2012 to support innovating on fabric architectures not only for the HPC but data centers, cloud, and Web 2.0 market [48].

Recent attention is focusing on Intel's next generation supercomputer which will adopt Cray's Aries interconnect with Intel Xeon Phi accelerator as their first non-InfiniBand based hybrid system after their acquisition of interconnect business of Cray [49]. This transition reflects the strategic decision of Cray to use an independent interconnect architecture rather than a processor specific one as AMD's performance and supply stability fell behind competitors' [44,50].

Unlike AMD X86 or Intel processor based systems, the top performing supercomputers using IBM Power (PC) processor were Many/Multicore systems. IBM initially developed the multicore architecture, later evolved to manycore systems, known as 'Blue Gene' technology. The Blue Gene approach is to use a large number of simple processing cores and to connect them via a low latency, highly-scalable custom interconnect [51]. This has the advantage of achieving a high aggregate memory bandwidth, whereas GPU clusters require messages to be copied from the GPU to the main memory and then from main memory to the remote node, whilst maintaining low power consumption as well as cost and floor space efficiency [52]. Currently, GPU/Accelerator based systems suggest smaller cluster solutions for the next generation HPC with its promising performance potential, however the Blue Gene architecture demonstrates an alternate direction of massively parallel quantities of independently operating cores with fewer programming challenges [53].

3.3. Model validation

To validate a predictive performance of the constructed model, we conducted hold-out sample tests. Specifically, a rolling origin was used to determine the forecast accuracy by collecting deviations from multiple forecasting origins so that the performance of the model can be tested both in near-term and far-term. This thus provides an objective measure of accuracy without being affected by occurrences unique to a certain fixed origin [54]. The comparative results with planar model and random walk are summarized in Table 4.

Since the first hybrid system, Blade Center QS22, appeared in 2008 in our dataset, the hold-out sample test was conducted from the origin of 2009 for hybrid systems. That is, the mean absolute deviation of 1.58 years was obtained from TFDEA when the model made a forecast on arrivals of post-2009 hybrid systems based on the rate of technological progress observed from 2008 to 2009. The overall forecasting error across the forecasting origins was found to be 1.32 years which is more accurate than planar model and random walk.

Although multicore systems showed successive introductions from 2007 to 2012, technological progress, i.e. expansion of SOA frontier surface, hasn't been observed until 2010. This rendered the model able to make a forecast only in 2011. The resulting forecast error of TFDEA was found to be about a year which is slightly bigger than that of planar model however care must be taken to interpret this error since this was obtained only from a year ahead forecast in 2011.

Consecutive introductions of manycore systems with a steady technological progress made it possible to conduct hold-out sample tests from the origin of 2007 to 2012. Notwithstanding a bigger average forecasting error of 1.49 years due to the inclusion of errors from longer forecasting windows than other two systems, TFDEA showed outperforming forecast results compared to the planar model and random walk.

Overall, it is shown that TFDEA model provides a reasonable forecast for three types of supercomputer systems with the maximum possible deviation of 18 months. In addition, it is interesting to note that forecasts from TFDEA tended to be less sensitive to the forecasting window than the planar model or random walk. This implies that the current technological progress of supercomputer technologies exhibits multifaceted characteristics that can be better explained by various tradeoffs derived from the frontier analysis. In contrast, a single design tradeoff identified from the central tendency model was shown to be vulnerable especially to the long-term forecast.

Table 4 Model validation using a rolling origin hold-out sample tests

Forecast Origin	Mean absolute deviation (unit: year)								
	Hybrid systems			Multicore systems			Manycore systems		
	TFDEA	Planar model	Random walk	TFDEA	Planar model	Random walk	TFDEA	Planar model	Random walk
2007	N/A	N/A	N/A	N/A	N/A	N/A	1.8075	2.8166	2.9127
2008	N/A	N/A	N/A	N/A	N/A	N/A	1.4470	2.5171	2.4949
2009	1.5814	2.7531	2.1852	N/A	N/A	N/A	2.0060	2.3593	2.0509
2010	1.1185	1.9956	1.5610	N/A	N/A	N/A	1.4996	2.0863	1.6016
2011	1.8304	1.5411	1.2778	0.9899	0.7498	1.0000	1.2739	1.8687	1.3720
2012	0.7564	1.2012	1.0000	N/A	N/A	N/A	0.8866	2.2269	1.0000
Average	1.3217	1.8728	1.5060	0.9899	0.7498	1.0000	1.4867	2.3125	1.9053

N/A: insufficient data

3.4. Forecasting

We now turn to the forecasting of the Exascale systems. As previously noted, the design goal of the Exascale supercomputer is expected to have the Exaflops (10^{18} flop / second) with

20MW power consumption and 100 million total cores (see Table 5) [7,11]. These specifications were set as a forecasting target to estimate when this level of systems could be reached considering the RoCs identified from the past advancements in a relevant segment.

Table 5 Exascale computer as a forecasting target

Cores	Power	Rmax
100 million	20 MW	1 Exaflops

Table 6 summarizes the forecasting results from three development possibilities. Exascale performance was forecasted to be achieved earliest by hybrid systems in 2021.13. Hybrid systems are expected to accomplish this with a relatively high individualized RoC of 2.22% and having the best current level of performance represented by Tianhe-2. Considering the possible deviations identified in the previous section, one could expect the arrival of hybrid Exascale system within the 2020 timeframe. In fact, many industry experts claim that GPU/Accelerator based systems will be more popular in TOP500 list for their outstanding energy efficiency, which may spur the Exascale development [13,16].

The forecasted arrival time of the first multicore based Exascale system is far beyond 2020 due to the slow rate of technological advancement: 1.19% as well as relatively underperforming performances. Note that projection from the planar model also estimated the arrival of multicore based Exascale system farther beyond 2020 timeframe. This implies that innovative engineering efforts are required for multicore based architecture to be scaled up to the Exaflop performance. Even though the RIKEN embarked on the project to develop the Exascale system continuing the preceding success of K-computer, IBM's cancellation of Blue Water contract and recent movement toward design house raise questions on the prospect of multicore based HPCs [55,56].

The first manycore based system is expected to reach the Exascale target by 2022.28. This technology path has been mostly led by the progress of Blue Gene architecture and shown the individualized RoC of 2.34%. Nonetheless, this fast advancement couldn't overcome the current performance gap with hybrid systems in the Exascale race.

Table 6 Forecast results of Exascale supercomputer

	Hybrid system	Multicore system	Manycore system
Individualized Rate of change (RoC)	1.022183	1.011872	1.023437
Forecasted arrival of Exascale supercomputer	2021.13	2031.74	2022.28

5. Discussion

The analysis of technologies' RoCs makes it possible to forecast a date for achieving Exascale performance from three different approaches; however it is worthwhile to examine these forecasts with consideration for the business environment and emerging technologies to anticipate the actual deployment possibilities of the Exascale systems.

The optimistic forecast is that, as seen from the high performing Tianhe-2 and Titan Cray XK7 system, there would be an Intel or AMD based system with a Xeon Phi or NVidia coprocessor and a custom Cray interconnect system. However, given business realities it's unlikely that the first Exascale system will use AMD processors. Intel purchased the Cray interconnect division and is expected to design the next generation Cray interconnect optimized for Intel processors and Xeon Phi coprocessors [57]. Existing technology trends and the

changing business environment would make a forecast a hybrid Exascale system with a Cray interconnect, Intel Processors and Xeon Phi coprocessors.

The 2.22% year improvement for hybrid systems has come mostly from a combination of advances in Cray systems, such as their transverse cooling system, Cray interconnects, AMD processors and NVidia coprocessors. It is difficult to determine the contribution of each component however it is worth noting that only Cray systems using AMD processors were SOAs. This implies that Cray's improvements are the highest contributor to the RoC for AMD based systems. Intel moving production of Cray interconnect chips from TSMC to Intel's more advanced processes will likely result in additional performance improvement. Thus, we can expect that Cray / Intel collaboration will result in RoC greater than the 2.22% and might reach the Exascale goal earlier.

As another possibility of achieving Exascale systems, IBM's Blue Gene architecture using IBM Power (PC) processor with custom interconnects has shown a 2.34% yearly improvement building on the 3rd highest rated Sequoia system. The Blue Gene architecture, with a high bandwidth, low latency interconnects and no coprocessors to consume bandwidth or complicate programming, is an alternative to the coprocessor architectures being driven by Intel and AMD. Given their stable business environment they may be more effective moving forward while Intel / Cray work out their new relationship.

Who has the system experience to build an Exascale system? Cray, IBM and Appro have built the largest SOA OEM systems. In 2012, Cray purchased Appro leaving two major supercomputer manufactures [58]. Based upon the captured RoCs and the business changes we can expect that the first Exascale system will be built by either Cray or IBM.

Data driven forecasting techniques, such as TFDEA, make a forecast on technical capabilities based upon released products, so emerging technologies that are not yet being

integrated into products are not considered. In the supercomputer academic literature, there is an ongoing debate about when the currently dominating large core processors (Intel, AMD) will be displaced by larger numbers power-efficient lower performance small cores such as ARM; much as what happened when microprocessors displaced vector machines in the 1990's and ARM based mobile computing platforms are affecting both Intel and AMD X86 desktop and laptop sales [13,18]. Although there is no ARM based supercomputer in the TOP500 yet, the European Mont-Blanc project is targeting getting one on the list by 2017 and NVidia is developing an ARM based supercomputer processor for use with its coprocessor chips [59]. Small cores are a potentially disruptive technology as long as power efficiency is concerned, therefore the further analysis is needed to investigate when it will overcome the challenges of building interconnects to handle a larger number of smaller cores or when software developers will overcome the synchronization challenges of effectively using more cores.

6. Conclusion

The HPC industry is experiencing a radical transition which requires improvement of power efficiency by a factor of 23 to deploy and/or manage the Exascale systems. This has created an industry concern that the naïve forecast based on the past performance curve may have to be adjusted. TFDEA is highly relevant in this context especially to deal with multiple tradeoffs between systems attributes. This study examined comparative prospects of three competing technology alternatives with various design possibilities considering business environment to achieve the Exascale computing so that researchers and manufacturers can have an accurate view on their development targets. In sum, the results showed that current development target of 2020 might entail technical risks considering the rate of change toward the energy efficiency observed

in the past. It is anticipated that either a Cray built hybrid systems using Intel processors or an IBM built Blue Gene architecture system using PowerPC processors will likely achieve the goal between early 2021 and late 2022.

In addition, the results provided a systematical measure of technological change which can guide a decision on the new product target setting practice. Specifically, the rate of change contains information not only about how much performance improvement is expected to be competitive but also about how much technical capability should be relinquished to achieve a specific level of technical capabilities in other attributes. One can also utilize this information to anticipate the possible disruptions. As shown in the HPC industry, the rate of change of manycore system was found to be faster than that of hybrid system. Although the arrival of hybrid Exascale system is forecasted earlier because of its current surpassing level of performance, fast rate of change of manycore system implies that the performance gap could be overcome and Blue Gene architecture might accomplish the Exascale goal earlier if hybrid system development couldn't keep up with the expected progress. Furthermore, the benchmark results can provide information about market segments relevant to the new technology alternatives. This makes it possible to identify competitors as well as dominant designs in a certain segment under which new product developers may be searching for market opportunities.

A new approach presented in this study can take segmented rate of change into account. Unlike traditional TFDEA relying on a constant rate of change, presented model enables to obtain variable rates of change for each product segment thereby either estimating technical capabilities of products at a certain point in time or forecasting the time by which desired levels of products will be operational.

Lastly, as an extension of this study, future research can consider:

- imposing weight restrictions to assess technologies in line with practical views,

- elaborating disruption possibilities from small core systems,
- incorporating external factors that can either stimulate or constrain the technological progress.

Acknowledgement

Authors would like to thank Dr.Wilfred Pinfold for his insightful comments on an earlier draft, the Associate Editor and three anonymous reviewers for their thoughtful and constructive suggestions during the review process.

REFERENCES

- [1] D.J. Lim, D.F. Kocaoglu, China - Can it move from imitation to innovation?, in: Technol. Manag. Energy Smart World (PICMET), 2011 Proc. PICMET '11, n.d. pp. 1–13.
- [2] D.G. Feitelson, The supercomputer industry in light of the Top500 data, Comput. Sci. Eng. 7 (2005) 42–47.
- [3] J. Dongarra, Visit to the National University for Defense Technology Changsha, China, Tennessee, 2013.
- [4] National Research Council, Getting Up to Speed: The Future of Supercomputing, 1st ed., National Academies Press, Washington, D.C, 2005.
- [5] J. Dongarra, H. Meuer, E. Strohmaier, TOP500 supercomputer sites, Supercomputer. 13 (1997) 89–111.
- [6] M. Peckham, RIKEN Plans Exascale Supercomputer '30 Times Faster' than Today's Fastest in Six Years, Time. (2013).
- [7] S. Ashby, P. Beckman, J. Chen, P. Colella, B. Collins, D. Crawford, et al., The Opportunities and Challenges of Exascale Computing, U.S Dep. Energy (Office Sci. (2010) 1–71.
- [8] H. Simon, T. Zacharia, R. Stevens, Modeling and Simulation at the Exascale for Energy and the Environment, Dep. Energy Tech. Rep. (2007).
- [9] C.-H. Hsu, J.A. Kuehn, S.W. Poole, Towards efficient supercomputing, in: Proc. Third Jt. WOSP/SIPEW Int. Conf. Perform. Eng. - ICPE '12, ACM Press, New York, New York, USA, 2012: p. 157.
- [10] S. Kamil, J. Shalf, E. Strohmaier, Power efficiency in high performance computing, in: 2008 IEEE Int. Symp. Parallel Distrib. Process., IEEE, 2008: pp. 1–8.
- [11] F. Robert, Who will step up to exascale?, Science (80-.). (2013) 264–266.
- [12] T. Geller, Supercomputing's exaflop target, Commun. ACM. 54 (2011) 16.
- [13] H. Simon, Why we need exascale and why we won't get there by 2020, in: Opt. Interconnects Conf., Lawrence Berkeley National Laboratory, Santa Fe, 2013.
- [14] J.D. Owens, M. Houston, D. Luebke, S. Green, J.E. Stone, J.C. Phillips, GPU Computing, Proc. IEEE. 96 (2008) 879–899.
- [15] Sheng Li, J.-H. Ahn, R.D. Strong, J.B. Brockman, D.M. Tullsen, N.P. Jouppi, McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures, in: MICRO-42. 42nd Annu. IEEE/ACM Int. Symp. On. IEEE, IEEE, New York, 2009: pp. 469–480.

- [16] J.P. Panziera, Five challenges for future supercomputers, *Supercomputers*. (2012) 38–40.
- [17] R. Vuduc, K. Czechowski, What GPU computing means for high-end systems, *IEEE Micro*. (2011) 74–78.
- [18] N. Rajovic, P. Carpenter, I. Gelado, N. Puzovic, A. Ramirez, Are mobile processors ready for HPC?, in: *Supercomput. Conf.*, Barcelona Supercomputing Center, Denver, 2013.
- [19] P. Bogetoft, L. Otto, *Benchmarking with DEA, SFA, and R*, 1st ed., Springer, New York, 2010.
- [20] D. Giokas, Bank branch operating efficiency: A comparative application of DEA and the loglinear model, *Omega*. 19 (1991) 549–557.
- [21] J.S. Liu, L.Y.Y. Lu, W.-M. Lu, B.J.Y. Lin, A survey of DEA applications, *Omega*. 41 (2013) 893–902.
- [22] E. Thanassoulis, A. Boussofiane, R.G. Dyson, A comparison of data envelopment analysis and ratio analysis as tools for performance assessment, *Omega*. 24 (1996) 229–244.
- [23] T.R. Anderson, K. Hollingsworth, L. Inman, Assessing the rate of change in the enterprise database system market over time using DEA, in: *PICMET '01. Portl. Int. Conf. Manag. Eng. Technol.*, 2001: p. 203.
- [24] B.F. Cole, *An evolutionary method for synthesizing technological planning and architectural advance*, Georgia Institute of Technology, 2009.
- [25] A.A. Tudorie, *Technology forecasting of electric vehicles using data envelopment analysis*, Delft University of Technology, 2012.
- [26] A. Durmusoglu, T. Dereli, On the technology forecasting using data envelopment analysis (TFDEA), in: *Technol. Manag. Energy Smart World (PICMET)*, 2011 Proc. PICMET '11, PICMET, U.S. Portland, 2011.
- [27] J.-O. Kim, J.-H. Kim, S.-K. Kim, A Hybrid Technological Forecasting Model by Identifying the Efficient DMUs: An Application to the Main Battle Tank, *Asian J. Technol. Innov.* 15 (2007) 83–102.
- [28] D.-J. Lim, T.R. Anderson, J. Kim, Forecast of wireless communication technology: A comparative study of regression and TFDEA Model, in: *Technol. Manag. Emerg. Technol.*, PICMET, Vancouver, Canada, 2012: pp. 1247–1253.
- [29] H.O. Fried, C.A.K. Lovell, S.S. Schmidt, *The measurement of productive efficiency and productivity growth*, 1st ed., Oxford University Press, New York, 2008.
- [30] K. Nakabayashi, K. Tone, Egoist's dilemma: a DEA game, *Omega*. 34 (2006) 135–148.
- [31] J.R. Doyle, R.H. Green, Comparing products using data envelopment analysis, *Omega*. 19 (1991) 631–8.

- [32] W.D. Cook, K. Tone, J. Zhu, Data envelopment analysis: Prior to choosing a model, *Omega*. 44 (2014) 1–4.
- [33] A.J. Alexander, J.R. Nelson, Measuring technological change: Aircraft turbine engines, *Technol. Forecast. Soc. Change*. 5 (1973) 189–203.
- [34] J.P. Martino, Measurement of technology using tradeoff surfaces, *Technol. Forecast. Soc. Change*. 27 (1985) 147–160.
- [35] D.-J. Lim, T.R. Anderson, O.L. Inman, Choosing effective dates from multiple optima in Technology Forecasting using Data Envelopment Analysis (TFDEA), *Technol. Forecast. Soc. Change*. 88 (2014) 91–97.
- [36] D.-J. Lim, T.R. Anderson, Improving Forecast Accuracy by a Segmented Rate of Change in Technology Forecasting using Data Envelopment Analysis, in: *Infrastruct. Serv. Integr., PICMET (In press)*, Kanazawa, 2014.
- [37] D.-J. Lim, S.R. Jahromi, T.R. Anderson, A.A. Tudorie, Comparing Technological Advancement of Hybrid Electric Vehicles (HEV) in Different Market Segments, *Technol. Forecast. Soc. Change*. (2014) <http://dx.doi.org/10.1016/j.techfore.2014.05.008>.
- [38] J.J. Dongarra, P. Luszczyk, A. Petit, The LINPACK Benchmark: past, present and future, *Concurr. Comput. Pract. Exp.* 15 (2003) 803–820.
- [39] D.-J. Lim, T.R. Anderson, An introduction to Technology Forecasting with a TFDEA Excel add-in, in: *Technol. Manag. Emerg. Technol. (PICMET)*, 2012 Proc. PICMET '12, 2012.
- [40] D.-J. Lim, N. Runde, T.R. Anderson, Applying Technology Forecasting to New Product Development Target Setting of LCD Panels, in: R. Klimberg, K.D. Lawrence (Eds.), *Adv. Bus. Manag. Forecast.*, 9th ed., Emerald Group Publishing Limited, 2013: p. 340.
- [41] T.R. Anderson, L. Inman, Resolving the issue of multiple optima in Technology Forecasting using Data Envelopment Analysis, in: *Technol. Manag. Energy Smart World, PICMET*, 2011: pp. 1–5.
- [42] P. Wood, Cray Inc. replacing IBM to build UI supercomputer, *News Gaz.* (2011) 1.
- [43] L. Lawrence, Cray jumped from AMD for the flexibility and performance of Intel chips, *Inq.* (2012).
- [44] C. Murariu, Intel Signs Deal with Cray, Marks AMD's Departure from Cray's Supercomputers, *Softpedia*. (2012).
- [45] F. Wang, C.-Q. Yang, Y.-F. Du, J. Chen, H.-Z. Yi, W.-X. Xu, Optimizing Linpack Benchmark on GPU-Accelerated Petascale Supercomputer, *J. Comput. Sci. Technol.* 26 (2011) 854–865.
- [46] A. Cole, Intel and InfiniBand: Pure HPC Play, or Is There a Fabric in the Works?, *IT Bus. Edge*. (2012).

- [47] InfiniBand Strengthens Leadership as the High-Speed Interconnect Of Choice, Mellanox Technol. (2012) 1–31.
- [48] Z. Whittaker, Intel buys QLogic’s InfiniBand assets for \$125 million, ZD Net. (2012).
- [49] IntelPR, Intel Acquires Industry-Leading, High-Performance Computing Interconnect Technology and Expertise, Intel Newsroom. (2012).
- [50] T. Morgan, Intel came a-knockin’ for Cray super interconnects, Regist. (2012).
- [51] R. Haring, M. Ohmacht, T. Fox, M. Gschwind, P. Boyle, N. Chist, et al., The IBM Blue Gene/Q Compute Chip, Ieee Micro. 32 (2011) 48–60.
- [52] C. Lazou, Should I Buy GPGPUs or Blue Gene?, HPCwire. (2010).
- [53] M. Feldman, Blue Genes and GPU Clusters Top the Latest Green500, HPCwire. (2011).
- [54] L.J. Tashman, Out-of-sample tests of forecasting accuracy: an analysis and review, Int. J. Forecast. 16 (2000) 437–450.
- [55] H. Kimihiko, Exascale Supercomputer Project, Kobe, 2013.
- [56] L. Tung, After X86 servers, is IBM gearing up to sell its chips business too?, ZDNet. (2014).
- [57] Davis Nick, Cray CS300 Cluster Supercomputers to Feature New Intel(R) Xeon Phi(TM) Coprocessors x100 Family, Cray Media. (2013).
- [58] Davis Nick, Cray to Acquire Appro International, Cray Media. (2012).
- [59] R.V. Aroca, L.M.G. Gonçalves, Towards green data centers: A comparison of x86 and ARM architectures power efficiency, J. Parallel Distrib. Comput. 72 (2012) 1770–1780.

Highlights:

- The challenges to deploy Exascale supercomputer are addressed.
- The segmented rate of change calculation is presented in Technology Forecasting using DEA (TFDEA).
- Technological progresses of hybrid, multicore, and manycore systems are compared.
- The arrival of the Exascale supercomputer is forecasted based on identified rate of changes.
- Innovative engineering efforts may be required to achieve Exascale goal within 2020 timeframe.

Author details:

Mr. Dong-Joon Lim

Department of Engineering and Technology Management

Portland State University

1900 SW 4th Avenue, Suite LL-50-02

Portland, OR 97201

United States

Tel: +1 5036880385

Fax: +1 5037254667

E-mail: tgno3.com@gmail.com