A Retrospective Look Back on the Road Towards Energy Proportionality

Daniel Wong
Department of Electrical and Computer Engineering
University of California, Riverside
Email: dwong@ece.ucr.edu

Julia Chen Murali Annavaram
Ming Hsieh Department of Electrical Engineering
University of Southern California
Email: {juliache,annavara}@usc.edu

***Abstract*—In this paper, we take a retrospective look back at the road taken towards improving energy proportionality, in order to find out where we are currently, and how we got here. Through statistical regression of published SPECpower results, were able to identify and quantify the sources of past EP improvements.**

I. Introduction

Energy proportional computing has been a major research focus over the past near-decade [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]. Servers with poor energy proportional behavior, while energy efficient at peak utilization, are highly energy inefficient at intermediate utilizations. In the past, great emphasis has been placed on improving the energy efficiency at peak utilization. On the other end of the spectrum, many idle low power techniques were also introduced to save energy when a server has no work to perform. But unfortunately, most workloads rarely spend time at these two extremes; workloads spend the majority of their time in the low/mid utilization regions [2]. This led to a call for *energy proportional* system design, where servers should consume power proportional to their utilization. Since then, a greater emphasis on intermediate utilization efficiency has emerged.

In this work we take a retrospective look back to identify where past energy proportionality improvements came from, and to explain historical trends. Through statistical regression of published SPECpower results, we're able to identify and quantify the sources of past EP improvements. We quantify that processor and DRAM innovations accounted for 74.5% of the EP variation, with the processor uniquely contributing 4x more than DRAM. Despite the fact that processors account for the largest fraction of power consumption, improvements to all other components (memory, storage, network, motherboard, etc.) will contribute as much as the processor to future server EP improvements.

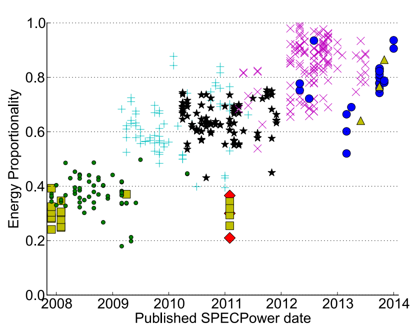

Figure 1. Historical Trends labeled with processor generation. Processors with ˆ label supports DDR2, all others support DDR3. Improvements to EP occurs in spurts.

II. Where are we now and how did we get here?

A. Measuring Energy Proportionality

There has been several works which quantify and measure EP [2, 4, 5, 8, 11]. We will adopt the energy proportionality (EP) metrics presented in [8]. Using this metric, an ideal energy proportional server would have EP = 1, and any EP less than 1 is less than ideal.

B. Methodology

For our retrospective analysis, we use reported SPECpower benchmark [12] results. Our SPECpower dataset contains 398 servers from 2007 to 2014 that are a representative mix of server configurations in current production environments.

Figure 1 shows the historical EP trend. While it is clear that energy proportionality has improved drastically over the past near-decade [4, 5, 8], what is not clear is *what* technological advances contributed to improving energy proportionality and *how much* these advances contributed. We limit our analysis to coarse-grain advances (such as processor generations, memory type, etc.) due to limitations of what we can infer from reported SPECpower server configurations. In order to identify individual component contributions to energy proportionality, we propose applying statistical *stepwise regression*.

1) Feature Annotation: We annotated each server result with various technology features associated with that server. For each server in the SPECpower dataset we identified the following: the processor generation (labelled *gen* in the table), whether the server used SSDs, which DRAM technology was used for memory subsystem, whether fully buffered DIMMs or registered memory DIMMs were used, and finally whether the DRAM was low voltage DDR3 or otherwise. We performed the stepwise regression analysis using SPSS 22 [13].

This work was supported by DARPAPERFECT-HR0011-12-2-0020 and NSF-CAREER-0954211.

Model	R	R Squared	Adjusted R Squared	Variable	Unstandardized Coefficients B	Unstandardized Coefficients Std. Error	t	Sig.	Semi-partial Correlations
1	.841	.708	.707	(Constant)	.189	.016	11.683	.000	
				gen	.122	.004	30.759	.000	.841
2	.863	.745	.744	(Constant)	.202	.015	13.245	.000	
				gen	.088	.006	15.071	.000	.386
				DDR3	.148	.020	7.516	.000	.192
3	.868	.754	.752	(Constant)	.234	.017	13.547	.000	
				gen	.070	.007	9.278	.000	.234
				DDR3	.158	.020	8.085	.000	.204
				LowVoltage	.057	.015	3.760	.000	.095
4	.871	.758	.756	(Constant)	.268	.021	12.468	.000	
				gen	.069	.007	9.209	.000	.230
				DDR3	.128	.022	5.708	.000	.143
				LowVoltage	.058	.015	3.834	.000	.096
				FBDIMM	-.056	.021	-2.586	.010	-.065

Table I
STEPWISE REGRESSION ANALYSIS RESULTS

2) Stepwise multiple regression: Stepwise multiple regression is a method commonly used in educational and psychological research to select a set of predictor variables, from a large pool of predictor variables, that can best explain the dependent variable. In this paper, we use a standard inclusion criterion of $p < 0.05$ and exclusion criterion of $p > 0.05$, giving us a 95% level of confidence. A comprehensive description of the stepwise multiple regression is beyond the scope of this work; [14, 15] provides a good introduction.

C. Stepwise Regression Output

Table I shows the results of the stepwise regression. The stepwise regression produced 4 models. The table column headers of importance are defined below.

Semi-partial Correlation: The squared semi-partial correlation gives us the percentage of the total variance in y that is uniquely accounted for by a particular predictor variable beyond that accounted for by other predictor variables. For example, in model 2, 14.9% of the total variance of EP is uniquely accounted for by *gen* that is not accounted for by *DDR3* (0.386^2).

1) Selecting a parsimonious model: A *parsimonious model* is a model that achieves a desired level of explanation or prediction with as few predictor variables as possible. The stepwise regression stopped after generating model 4 due to no other predictor variables being significant. According to the results of the stepwise regression, only *gen*, *DDR3*, *LowVoltage*, and *FBDIMM* are significant in predicting a server's EP. Although *LowVoltage* and *FBDIMM* are statistically significant in contributing to the model, they both uniquely help explain less than 1% of the total variation in EP based on the squared semi-partial correlation. Therefore, the main contributing factors are *gen* and *DDR3*. Therefore, we will use model 2, which uses just *gen* and *DDR3*, because it is the more parsimonious model.

D. Stepwise Regression Results

1) Findings from Stepwise Regression: The goal of this section is to answer the question, how did we get here? The regression model showed that processor generations and DRAM switching from DDR2 to DDR3 are the main driving forces behind historical EP growth. From the regression, we are able to quantify the effect of each component through our statistical analysis. Using model 2, the switch to DDR3 is credited with improving EP by 0.148 and every new generation of processor is expected to improve EP by 0.088. Together, processor generations and switch to DDR3 accounts for 74.5% of EP variation. We can also observe that *gen* can uniquely explain 14.9% of EP variation, while *DDR3* can uniquely explain 3.7% of EP variation. *Thus processor generation can uniquely explain 4x more variation to EP than the switch to DDR3.* The remaining 55.9% is contributed to a combination of both processor and memory.

III. CONCLUSION

Energy proportionality of current servers has improved to the point where servers are now approaching the traditional definition of "ideal" proportional servers. Through statistical regression of published SPECpower results, we are able to quantify the contributions of various components towards the growth of energy proportionality. We identified that processors and memory are the two biggest driver of historical energy proportionality growth.

REFERENCES

[1] P. Ranganathan, *et.al*, "Ensemble-level power management for dense blade servers," *ISCA '06*, 2006.
[2] L. A. Barroso and U. Holzle, "The case for energy-proportional computing," *Computer*, vol. 40, no. 12, dec 2007.
[3] N. Tolia, *et.al*, "Delivering energy proportionality with non energy-proportional systems: optimizing the ensemble," in *HotPower'08*, 2008.
[4] G. Varsamopoulos and S. K. S. Gupta, "Energy proportionality and the future: Metrics and directions," in *ICPP Workshops*, 2010.
[5] F. Ryckbosch, *et.al*, "Trends in server energy proportionality," *Computer*, vol. 44, no. 9, Sep. 2011.
[6] K. T. Malladi, *et.al*, "Towards energy-proportional datacenter memory with mobile DRAM," in *ISCA '12*, 2012.
[7] K. T. Malladi, *et.al*, "Rethinking DRAM Powermodes for Energy Proportionality ," in *MICRO '12*, 2012.
[8] D. Wong and M. Annavaram, "Knightshift: Scaling the energy proportionality wall through server-level heterogeneity," in *MICRO '12*, 2012.
[9] ——, "Implications of high energy proportional servers on cluster-wide energy proportionality," in *HPCA '14*, 2014.
[10] S. Kanev, *et.al*, "Tradeoffs between power management and tail latency in warehouse-scale applications," in *IISWC '14*, 2014.
[11] B. Subramaniam and W.-c. Feng, "Towards energy-proportional computing for enterprise-class server workloads," in *ICPE '13*, 2013.
[12] www.spec.org/power_ssj2008/, "Spec power_ssj2008," 2012.
[13] IBM SPSS Statistics, *Version 22.0*.
[14] J. Cohen, *et.al*, *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences*, 3rd ed. Routledge, 2002.
[15] R. A. Gordon, *Regression Analysis for the Social Sciences*, 1st ed. Routledge, 2010.