

A note on the Zipf distribution of Top500 supercomputers

-- And why Grid computing has the wind in its sails --

Matei Ripeanu
(matei@ece.ubc.ca)

Trends inferred from the fastest supercomputers lists for the last 13 years indicate that aggregating the computational power of relatively small machines is becoming increasingly rewarding. It is thus no coincidence that Grid computing, which provides the infrastructure to build these controlled, secure resource aggregations, continues to attract increasing interest.

Most natural and technical phenomena are characterized by highly unbalanced distributions: there are few powerful, devastating earthquakes and countless unnoticeable ones; there are a few people making more than one million dollars a year while much lower pay rates are the norm; and there is only one machine with a peak compute rate larger than 100 TFLOPS (10^{14} floating-point operations per second) while millions of machines work under 1 GFLOPS. Many of these distributions (city sizes, incomes, word frequency) fit a Zipf distribution [1, 2]: after ranking all events in decreasing order, the size of an event is proportional to its rank to a negative constant power k : $E_{size} = c * rank^{-k}$. Visually, as the equation above can be rewritten as $\log(E_{size}) = c - k * \log(rank)$, the indication of a Zipf distribution is a straight line when the distribution is plotted on logarithmic scales on both event size and rank axes.

Adamic et al. [3] conjecture that Zipf distributions characterize many of the entities participating in the Internet, from obvious attributes like CPU speed, available disk space, or network bandwidth, to more elaborate ones such as inter-failure time, node trustworthiness, or reliability. Indeed, existing data support this intuition: Internet's autonomous system size [4], the number of physical links attached to a router [5], node bandwidth for Gnutella network nodes [6, 7], or web-site popularity [3], all follow Zipf distributions, or at least highly heterogeneous distributions that can be well approximated as Zipf.

Interestingly, peak compute rates of world's top supercomputers, as measured by the Top500

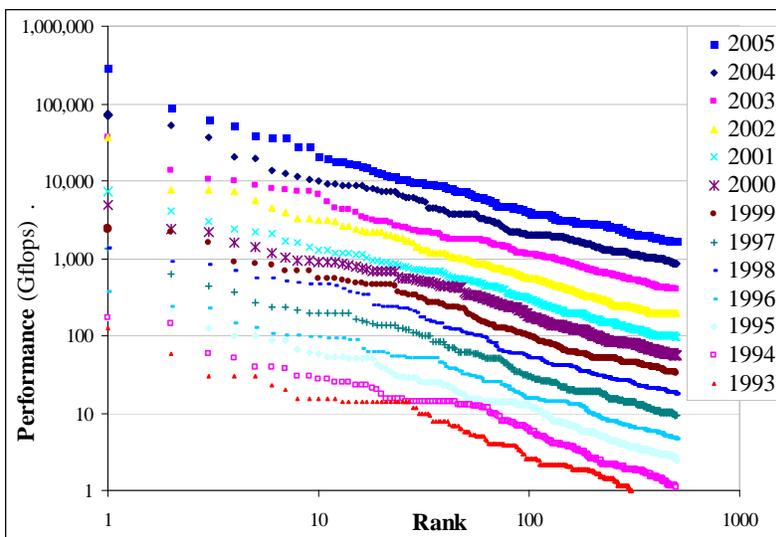


Figure 1: Peak processing rate (GFLOPS) for world's fastest supercomputers as presented by Top500 list from 1993 to 2005. Each series of points represents one year on this plot with log scales on both axes.

list [8], have followed Zipf distributions for each year from 1993 to 2005 too (Figure 1).

It is worth observing that over these 13 years the plots corresponding to these distributions appear as almost parallel and equidistant lines uncovering a yearly constant factor improvement in compute power. Part of this amazing exponential compute power increase over time is explained by performance improvements at the chip level as predicted by More's law. Additionally, this data implies that the other components of the compute stack (e.g., network interconnects, communication libraries, other elements of the software stack) have roughly kept up with chip exponential level improvements. (We note that the average yearly performance increase is about 90% per year).

At a closer look, however, a fascinating trend surfaces: over time *the performance of machines in the tail of these distributions has grown faster than that of top machines*. To explore this trend in a systematic way we compute the exponent of the Zipf distributions for each year, which translates in fact in computing the slope of the lines in the log-log plots of Figure 1. Figure 2 presents the evolution of this exponent over time. We can observe a clear trend: k continuous and steady decrease form 0.93 in 1993 to 0.65 in 2005. Note that the lower the exponent k is, the smaller the differential between machines at the top and at the bottom of the Top500 list for a certain year becomes.

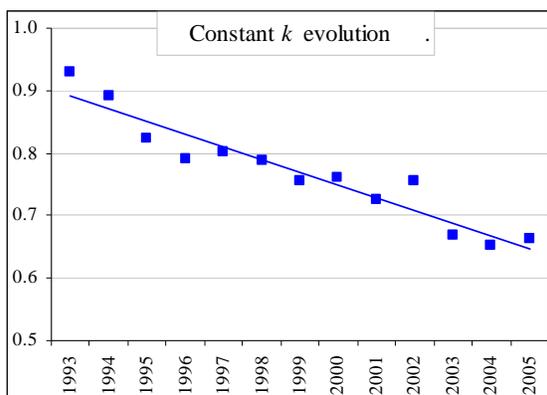


Figure 2: The evolution of the constant k determining the shape of Zipf distributions of compute power of computers in Top500 supercomputers list.

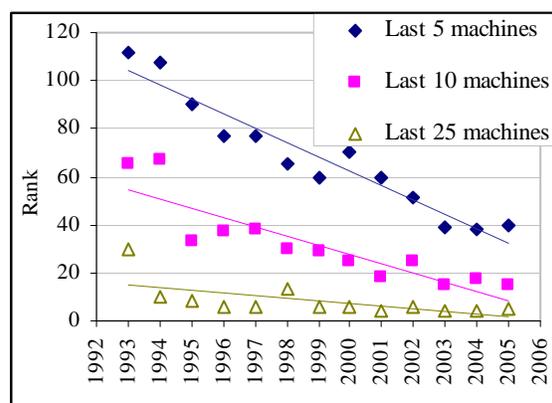


Figure 3: The ranking of a hypothetical machine that aggregates the power of the last 5, 10, or 25 machines in the Top500 list. Note the continuous improvement in ranking over time.

To quickly build intuition on the impact of this trend, Figure 3 presents how a hypothetical machine that aggregates the power of the last five machines in the Top500 list would have ranked. Note the continuous improvement in rankings: from 112th ranking in 1993 to 40th in 2005. The same valid for any distributed machine that aggregates the power available in the tail of these distributions. The aggregate of the last 25 machine would have ranked 30th in 1993 and 5th in 2005. Note that aggregating the power of multiple, geographically distributed, supercomputers is not an implausible scenario: as far back as 2001 Allen et al. [9] demonstrated that an efficient middleware stack and a small number of application-level changes can help hide network latency and the heterogeneity inherent when coupling multiple machines and can lead to achieving high computational efficiencies even for tightly-coupled, ‘stencil’ computations.

To summarize: a long running trend indicates that it is increasingly rewarding to aggregate the computational power of relatively small machines. If this trend persists,

interest will continue to shift from building large machines to large-scale integrations of less powerful systems. The increasing popularity of Grid computing is, partially, a result of this shift in interest as Grids provide computational models and middleware for large-scale resource integration, though with specific assumptions on resource participation and ownership, security requirements, and end-user/application requirements. Given its focus, the trend presented in this article is like the wind in the Grid computing' sails.

Acknowledgements

The author would like to thank Ian Foster and Alex Szalay for their insightful comments.

References

- [1] M. Schroeder, *Fractals, Chaos, Power Laws : Minutes from an Infinite Paradise*: W.H. Freeman and Company, 1991.
- [2] N. Shiode and M. Batty, "Power Law Distributions in Real and Virtual Worlds," INET 2000, Yokohama, Japan, 2000.
- [3] L. A. Adamic and B. A. Huberman, "Zipf's law and the Internet," *Glottometrics*, vol. 3, pp. 143-150, 2002.
- [4] H. Tangmunarunkit, S. Doyle, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Does AS Size Determine Degree in AS Topology?," *ACM Computer Communication Review*, 2001.
- [5] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-Law Relationships of the Internet Topology," SIGCOMM 1999, 1999.
- [6] S. Saroiu, P. K. Gummadi, and S. D. Gribble, "A Measurement Study of Peer-to-Peer File Sharing Systems," Multimedia Computing and Networking Conference (MMCN), San Jose, CA, USA, 2002.
- [7] M. Ripeanu, I. Foster, and A. Iamnitchi, "Mapping the Gnutella Network: Properties of Large-Scale Peer-to-Peer Systems and Implications for System Design," *Internet Computing Journal*, vol. 6, 2002.
- [8] "TOP500 Supercomputer Sites - <http://www.top500.org/>," 2002.
- [9] G. Allen, T. Dramlitsch, I. Foster, T. Goodale, N. Karonis, M. Ripeanu, E. Seidel, and B. Toonen, "Supporting Efficient Execution in Heterogeneous Distributed Computing Environments with Cactus and Globus," SC'2001, Denver Colorado, 2001.