Title
Amdahl's Reliability Law: A Simple Quantification of the Weakest-Link Phenomenon

Journal
Computer, 48(7)

ISSN
0018-9162

Author
Parhami, Behrooz

Publication Date
2015

DOI
10.1109/mc.2015.202

Peer reviewed

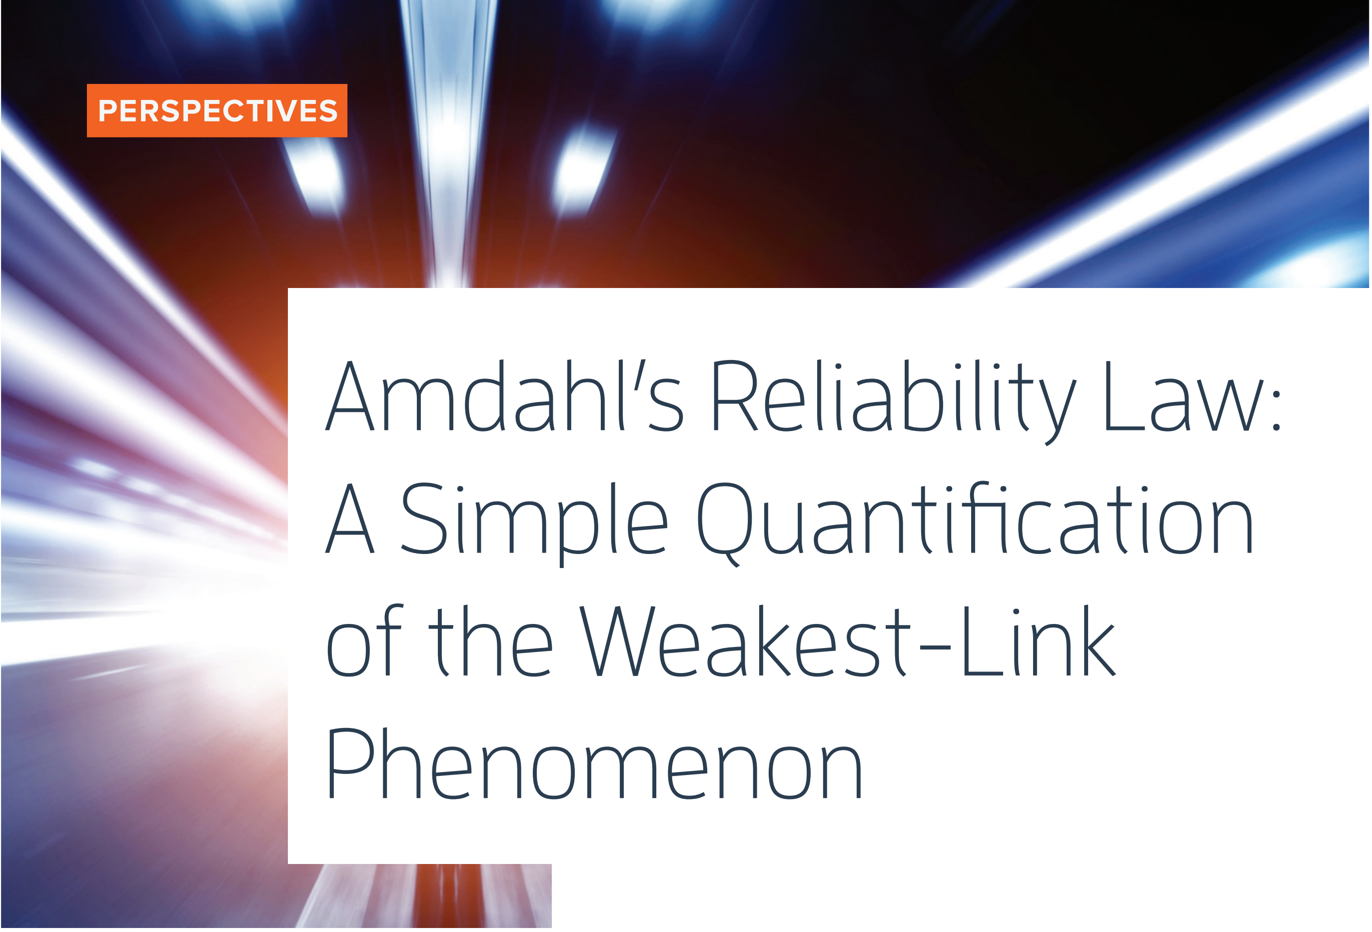

PERSPECTIVES

Amdahl's Reliability Law: A Simple Quantification of the Weakest-Link Phenomenon

Behrooz Parhami, University of California, Santa Barbara

Amdahl's performance law is a useful tool for imparting key ideas about performance limits and the importance of balanced system design to students and practitioners alike. Adapting it to model the weakest-link phenomenon in reliability engineering extends its utility and exposes its expressive power within other engineering disciplines.

Gene Amdahl presented his now-famous performance—or speedup—law in a short paper he wrote to defend high-performance uniprocessors against highly parallel systems, a debate that was later described as the few-elephants versus army-of-ants approaches to supercomputing. But his original formulation,[1] while quite valid for the specific point he was trying to make, has been contested[2,3] and extended to take factors other than time-to-completion into account.[4–7]

Amdahl's performance law is a very useful tool for imparting the notion that when a computational problem does not have unlimited parallelism, the utility of a massively parallel solution might be quite modest—in other words, indivisible parts of a computational task will lead to all but one of the processors remaining idle for extended time periods. Subsequent to Amdahl's original formulation, researchers have found the law, which quantifies the performance benefits of parallel processing, to be applicable to many other domains of system design as well.

Amdahl observed that if a program needs running time $T = T_1$ on a uniprocessor, of which the fraction Tf (with $f < 1$) is spent on inherently sequential tasks and the remaining fraction $T(1 - f)$ can benefit from parallelism, it will run in time $T_p = Tf + T(1 - f)/p$ on a p-processor parallel system, assuming perfect parallelization with no overhead. In practice, there will be some overhead and resource waste, so the cited formula for T_p represents a best-case scenario. Thus, the speedup achieved will be

$$s = \frac{T_1}{T_p} \leq \frac{1}{f + (1-f)/p} \leq \min\left(p, \frac{1}{f}\right).$$

As Figure 1 shows, the fact that the achieved speedup cannot exceed p with p processors is evident, although

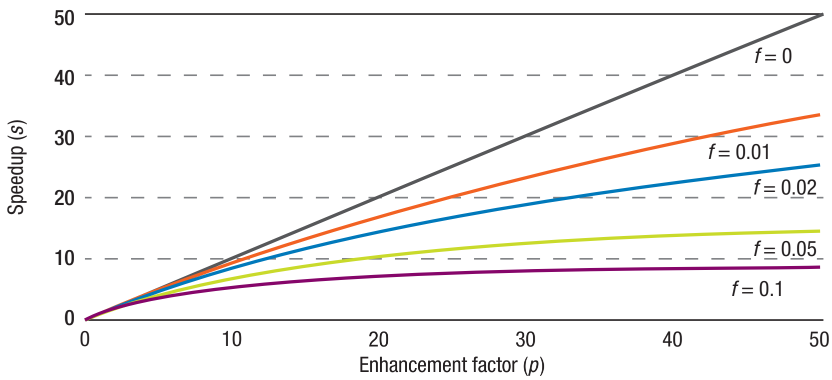

FIGURE 1. Performance enhancement. Speedup s as a function of the number of processors (or enhancement factor) p and the inherently sequential fraction of the computation (or the part of a system unaffected by performance-motivated redesign) f.

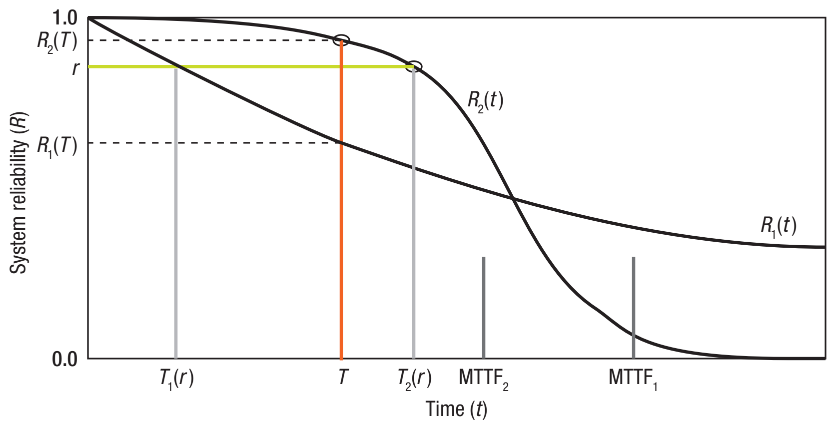

FIGURE 2. Comparing reliabilities. For two candidate systems, mission-time improvement at a fixed reliability r, reliability improvement at a fixed mission time T, and mean-time-to-failure (MTTF) improvement constitute possible comparison yardsticks.

even this seemingly obvious claim has been, rightly, challenged, citing caching effects or using the analogy that moving a heavy sofa can be done much more than four times as fast when using four movers instead of one.[8]

Simply put, it does not pay to speed up one part of a computation when performance is limited by other parts. Amdahl's contribution was the observation that the achieved speedup cannot exceed $1/f$, even if p tends to infinity. His formula suggests that for a program that spends 5 percent of its running time performing inherently sequential tasks, the parallel speedup that can be achieved is upper-bounded by 20, even on a parallel system with many thousands of processors.

APPLYING AMDAHL'S LAW

It turns out that Amdahl's law has a much broader domain of applicability than merely for parallel processing. Several authors, including this one,[9] have used examples such as the following to show this generality.

Let us assume that certain programs of interest spend 30 percent of their execution time on floating-point addition, 25 percent on floating-point multiplication, and 10 percent on floating-point division. Which of the three redesign options—making the adder twice as fast, making the multiplier three times as fast, or making the divider 10 times as fast—has the greatest impact on overall performance, if we ignore the costs of the three options? Solving the engineering problem in this example entails a straightforward application of Amdahl's performance formula, using 1 – 0.3, 1 – 0.25, and 1 – 0.1 as the unaffected fraction f of the computation, along with improvement factors 2, 3, and 10, respectively:

$$s_{\text{Improved-flp-adder}} = \frac{1}{0.7 + 0.3\,/\,2} = 1.18$$

$$s_{\text{Improved-flp-multiplier}} = \frac{1}{0.75 + 0.25\,/\,3} = 1.20$$

$$s_{\text{Improved-flp-divider}} = \frac{1}{0.9 + 0.1\,/\,10} = 1.10.$$

Note that improving the more extensively used operations can have a greater performance impact, even if the improvement is by a smaller factor.

The lesson to take away from studying Amdahl's law and examples such as the one above is that achieving high performance requires a balanced system: improving one aspect of system performance (such as certain floating-point operations in the example above) might have a limited impact on the overall performance.

In reliability, too, a similar maxim is often cited, namely, that reliability is a weakest-link phenomenon: a chain does not get stronger if you improve the strength of some links but make no change to the other links. Before embarking on a discussion of the reliability counterpart to Amdahl's speedup formula, we need to know how to compare reliabilities.

ABOUT THE AUTHOR

BEHROOZ PARHAMI is a professor of electrical and computer engineering, and a faculty member in the computer engineering program, at the University of California, Santa Barbara. His research interests include computer arithmetic, parallel processing, and dependable computing. Parhami received a PhD in computer science from University of California, Los Angeles. He is an IEEE Life Fellow, and a Fellow of the Institution of Engineering and Technology (IET) and the British Computer Society. Contact him at parhami@ece.ucsb.edu.

COMPARING SYSTEM RELIABILITIES

As Figure 2 indicates, there are many ways to compare reliabilities, but we focus on two in this discussion. One is the mean time to failure (MTTF), which in the special case of a system having failure rate λ with exponential reliability $R = e^{-\lambda t}$ (often assumed) is $1/\lambda$. For two systems having failure rates λ_1 and λ_2, the MTTF improvement factor (MTTFIF) of system 2 over system 1 is

$$\text{MTTFIF}_{2/1} = \text{MTTF}_2/\text{MTTF}_1 = \lambda_1/\lambda_2.$$

Thus, halving the failure rate doubles the MTTF and leads to an MTTFIF of 2.0.

A second comparative measure[10] is the reliability improvement index (RII), defined as

$$\text{RII}_{2/1} = \ln R_1/\ln R_2.$$

The motivation behind defining the RII is that for systems of practical interest, the reliability is always very close to 1 so that neither reliability difference (RD) $R_2 - R_1$ nor reliability ratio (RR) R_2/R_1 provides good discrimination among systems. Consider, for example, three systems with reliabilities $R_1 = 0.99$, $R_2 = 0.999$, and $R_3 = 0.9999$. We have $\text{RD}_{2/1} = 0.009$, $\text{RD}_{3/2} = 0.0009$, and $\text{RD}_{3/1} = 0.0099$, which are all fairly small numbers with insignificant variation. Similarly, $\text{RR}_{2/1} = 1.0091$, $\text{RR}_{3/2} = 1.0009$, and $\text{RR}_{3/1} = 1.0100$ are all close to 1, again providing insufficient discrimination. On the other hand, $\text{RII}_{2/1} \approx 10.05$, $\text{RII}_{3/2} \approx 10.00$, and $\text{RII}_{3/1} \approx 100.5$ properly reflect significant differences in reliabilities.

Note that the two measures MTTFIF and RII coincide in the case of systems with exponential reliability formulas because for such systems, $\ln R = -t/\text{MTTF}$.

AMDAHL'S RELIABILITY LAW

According to the arguments to follow, Amdahl's reliability law states that when you improve the failure rate for parts of the system, accounting for a fraction $1 - f$ of system complexity, by a factor p, then the overall improvement achieved is not only less than p, as expected, but it cannot exceed $1/f$ also. For example, if only half the system is improved, then the overall improvement cannot exceed 2.0, regardless of the improvement factor p.

Let a system with an overall failure rate λ, or exponential reliability equation $R_{\text{Original}} = e^{-\lambda t}$, consist of parts 1 and 2 having failure rates λf and $\lambda(1 - f)$, respectively. If part 1 is left unchanged and part 2 is improved to have a lower failure rate $\lambda(1 - f)/p$, the new overall failure rate will be $\gamma = \lambda[f + (1 - f)/p]$, with an associated reliability equation $R_{\text{Improved}} = e^{-\lambda[f + (1 - f)/p]t}$. The RII is thus

$$\text{RII} = \ln R_{\text{Original}} / \ln R_{\text{Improved}} = \frac{1}{f + (1-f)/p}.$$

As noted earlier, the MTTF improvement factor is characterized by the same formula:

$$\text{MTTFIF} = (1/\gamma)/(1/\lambda) = \frac{1}{f + (1-f)/p}.$$

It is quite interesting that a formula developed to expose the pitfalls of parallel processing when applications exhibit limited parallelism can be used to broadly assess the importance of balance in attaining high performance or reliability—as well as in formulating other design, scheduling, and management paradigms.

Just as we have a generalized form of Amdahl's performance law that envisages performance improvements of various magnitudes in different parts of the system, we can formulate a generalized form of Amdahl's reliability law. If n different parts of our system accounting for the fractions λf_i of the total failure rate λ, with $f_1 + f_2 + \ldots + f_n = 1$, are improved by the corresponding factors p_i, then the overall improvement is

$$\text{RII}_{\text{General}} = \frac{1}{\frac{f_1}{p_1} + \frac{f_2}{p_2} + \ldots + \frac{f_n}{p_n}}.$$

We would use the performance version of this formula if we carried out all three improvements in our example with floating-point operations at once, leading to

$$s_{\text{Improved-all-flp-units}} = \frac{1}{\frac{0.35}{1} + \frac{0.3}{2} + \frac{0.25}{3} + \frac{0.1}{10}} = 1.69.$$

Similarly, if the failure rates for various parts of a system are improved by different factors, the formula above for $\text{RII}_{\text{General}}$ yields the resulting reliability improvement index.

Formulating the age-old maxim that reliability is a weakest-link phenomenon via the introduction of Amdahl's reliability formula is beneficial for putting the maxim on a quantitative footing and making it more believable to students and practitioners of reliability engineering. This is another example of the extremely wide applicability of Amdahl's law beyond computer performance modeling and parallel system speedup.

A particularly promising area for future work would be bringing the cost of improved reliability into consideration alongside the generalized form of Amdahl's reliability formula. There are many possibilities for the cost function $C(f_i, p_i)$, depending on how cost rises with the improved part's complexity and the extent of improvement—for example, linear and superlinear functions of p_i, as well as general functions specified in tabular form. Standard optimization procedures would be useful in determining how to allocate a reliability improvement budget for each of the n parts. ■

REFERENCES

1. G. Amdahl, "Validity of the Single Processor Approach to Achieving Large-Scale Computing Capabilities," *Computer*, vol. 46, no. 12, 2013, pp. 483–485.
2. A.Y. Grama, A. Gupta, and V. Kumar, "Isoefficiency: Measuring the Scalability of Parallel Algorithms and Architectures," *IEEE Parallel & Distributed Technology*, vol. 1, no. 3, 1993, pp. 12–21.
3. J.L. Gustafson, "Reevaluating Amdahl's Law," *Comm. ACM*, vol. 31, no. 5, 1988, pp. 532–533.
4. S. Cho and R. G. Melhem, "Corollaries to Amdahl's Law for Energy," *Computer Architecture Letters*, vol. 7, no. 7, 2008, pp. 25–28.
5. M.D. Hill and M.R. Marty, "Amdahl's Law in the Multicore Era," *Computer*, vol. 41, no. 7, 2008, pp. 33–38.
6. I. Onyukdel and S.H. Hosseini, "Amdahl's Law: A Generalization under Processor Failures," *IEEE Trans. Reliability*, vol. 44, no. 3, 1995, pp. 455–462.
7. D.H. Woo and H.-H.S. Lee, "Extending Amdahl's Law for Energy-Efficient Computing in the Many-Core Era," *Computer*, vol. 41, no. 12, 2008, pp. 24–31.
8. S.G. Akl, "Secure File Transfer: A Computational Analog to the Furniture Moving Paradigm," *Parallel and Distributed Computing Practices*, vol. 5, no. 2, 2002, pp. 193–203.
9. B. Parhami, *Computer Architecture: From Microprocessors to Supercomputers*, Oxford Univ. Press, 2005.
10. T.F. Klaschka, "A Method for Redundancy Scheme Performance Assessment," *IEEE Trans. Computers*, vol. 20, no. 11, 1971, pp. 1371–1376.