Which scaling rule applies to large artificial neural networks

Technological limitations for biology-imitating computing

János Végh¹

Received: 3 April 2021 / Accepted: 26 August 2021 / Published online: 20 September 2021
© The Author(s) 2021

Abstract

Experience shows that cooperating and communicating computing systems, comprising segregated single processors, have severe performance limitations, which cannot be explained using von Neumann's classic computing paradigm. In his classic "First Draft," he warned that using a "too fast processor" vitiates his simple "procedure" (but not his computing model!); furthermore, that using the classic computing paradigm for imitating neuronal operations is *unsound*. Amdahl added that large machines, comprising many processors, have an inherent disadvantage. Given that artificial neural network's (ANN's) components are heavily communicating with each other, they are built from a large number of components designed/fabricated for use in conventional computing, furthermore they attempt to mimic biological operation using improper technological solutions, and their achievable payload computing performance is conceptually modest. The type of workload that artificial intelligence-based systems generate leads to an exceptionally low payload computational performance, and their design/technology limits their size to just above the "toy" level systems: The scaling of processor-based ANN systems is strongly nonlinear. Given the proliferation and growing size of ANN systems, we suggest ideas to estimate in advance the efficiency of the device or application. The wealth of ANN implementations and the proprietary technical data do not enable more. Through analyzing published measurements, we provide evidence that the role of data transfer time drastically influences both ANNs performance and feasibility. It is discussed how some major theoretical limiting factors, ANN's layer structure and their methods of technical implementation of communication affect their efficiency. The paper starts from von Neumann's original model, without neglecting the transfer time apart from processing time, and derives an appropriate interpretation and handling for Amdahl's law. It shows that, in that interpretation, Amdahl's law correctly describes ANNs.

Keywords Energy efficiency · Computing efficiency · Artificial intelligence · Scaling rule · Neural network · Temporal logic · Time-aware computing

1 Introduction

Given the proliferation of ANN-based devices, applications, and methods, furthermore that even supercomputers are re-targeted for AI applications, the efficacy of such systems is gaining growing importance. Von Neumann in his "First Draft" [51] provided an *approximation*, for (the timing relations of) *vacuum tubes only*. He warned (in his

section 6.3) that because the data transfer time is neglected in his model, using a "too fast processor" vitiates the procedure; furthermore, that using his paradigm for imitating neuronal operations is *unsound*, given that the conduction (transfer) time is longer than the synaptic (processing) time. This limitation means that in today's technology background, it is at least doubly unsound to apply his paradigm to describe scaling ANNs.¹ However, von Neumann did not provide another procedure that can consider the case corresponding to the today's technology,

✉ János Végh
janos.vegh@unideb.hu

¹ Kalimános BT, Debrecen, Hungary

¹ Given that the classic paradigm is unsound for describing neurons, their communicating network has not been touched.

and the case of neural computing, respectively [52]. The question opens, how then *networks of computing objects using neuronal operations*, mimicking or inspired by biology, *can be scaled*.

From computational point of view, an ANN is an adaptive distributed processor-based architecture widely used to utilize its inputs and simulate the human processing in terms of computation, response, and decision making. To analyze their performance, we start from the “first principles” of computing, scrutinizing the terms and omissions. Amdahl’s famous idea for distributed many-processor systems introduced “strong scaling.” He predicted [36]: “*scaling thus put larger machines [the brain inspired computers built from components designed for single-processor approach (SPA) computers] at an inherent disadvantage.*” It was guessed early [36] that the payload performance of processor assemblies does not scale linearly with the number of processors. However, appearance of “massively parallel” systems improved the degree of parallelization so much that a new approximation: The “weak scaling” [19] appeared. Due to the peculiarities of its workload, the AI-related progress (including ANNs) shows up much worse scaling [26] than expected and led to that “*Core progress in AI has stalled in some fields*” [21]. Also, the Gordon Bell Prize jury noticed [6] that “*Surprisingly, there have been no brain inspired massively parallel specialized computers [among the winners].*”

All the reasons listed above have their root essentially in *neglecting the temporal behavior of computing components and methods*. In Sect. 2, we shortly review the considered scaling methods and some of their consequences. In Sect. 2.1, Amdahl’s idea is shortly described (and partly: reinterpreted): His famous formula using our notations is introduced. In Sect. 2.2, we scrutinize the primary purpose of massively parallel processing, Gustafson’s idea. Section 3 discusses different factors affecting computing performance of processor-based (as well as some aspects of other electronic equipments) ANN systems.

2 Common scaling methods

The scaling methods used to model different implementations of parallelized sequential processing (aka “distributed computing”) are approximations to the more general model presented in [44, 45]. As discussed in [46, 49], *parallelized sequential systems have their inherent performance limitation*. Using that formalism and data from the TOP500 database [39], we could estimate performance limits for present supercomputers. It enabled us to comprehend why supercomputers have their inherent performance limit [46]. For the accuracy of the scaling method, see the case

of supercomputer Piz Daint in [46]. We also validated [48] our “time-aware scaling” (as a mostly empirical experience) through applying it, among others, for qualifying load balancing compiler, cloud operation, on-chip communication. Given that experts, with the same background, also build ANN systems, from similar components, we can safely assume that the same scaling is valid for those systems, too. Calibrating our systems for some specific workload (due to the lack of validated data) is not always possible, but one can compare the behavior of systems and draw some general conclusions.

2.1 Amdahl’s law

Amdahl’s law (called also “strong scaling”) is usually formulated as

$$S^{-1} = (1 - \alpha) + \alpha/N \quad (1)$$

where N is the number of parallelized code fragments, α is the ratio of parallelizable fraction to total (so $(1 - \alpha)$ is the “serial percentage”), and S is a measurable speedup. That is, Amdahl’s law considers a *fixed-size problem*, and α portion of the task is distributed to fellow processors.

When calculating the speedup, one calculates

$$S = \frac{(1 - \alpha) + \alpha}{(1 - \alpha) + \alpha/N} = \frac{N}{N \cdot (1 - \alpha) + \alpha}. \quad (2)$$

However, as expressed in [34]: “*Even though Amdahl’s law is theoretically correct, the serial percentage is not practically obtainable.*” That is, concerning S , there is no doubt that it is derived as the ratio of *measured execution times*, for non-parallelized and parallelized cases, respectively. But, what is the exact interpretation of α , and how can it be used?

Amdahl listed performance affecting factors, such as “boundaries are likely to be irregular; interiors are inhomogeneous; computations required may be dependent on the states of the variables at each point; *propagation rates of different physical effects may be quite different*; the rate of convergence or convergence at all may be strongly dependent on sweeping through the array along different axes on succeeding passes, etc.” Amdahl has foreseen issues with “sparse” calculations (or in general: *the role of data transfer*) as well as that the *physical size* of computer and the *interconnection* of its computing units (especially in the case of distributed systems) also matters.

Amdahl used wording “*the fraction of the computational load,*” giving way to his followers to give meaning to that term. This (unfortunately formulated) phrase “*has caused nearly three decades of confusion in the parallel processing community. This confusion disappears when processing times are used in the formulations*” [34]. On the one side, *it was guessed that Amdahl’s law is valid only for*

software (for the number of executed instructions), and on the other side *other affecting factors, he mentioned but did not discuss in detail, were forgotten.*

Expressing Amdahl's speedup is not simple: “*For example, if the following percentage is to be derived from computational experiments, i.e., recording the total parallel elapsed time and the parallel-only elapsed time, then it can contain all overheads, such as communication, synchronization, input/output, and memory access. The law offers no help to separate these factors. On the other hand, if we obtain the serial percentage by counting the number of total serial and parallel instructions in a program, then all other overheads are excluded. However, in this case, the predicted speedup may never agree with the experiments.*” [34] Moreover, the experimental one is always smaller than the theoretical one.

From computational experiments, one can express α from Eq. (1) in terms measurable experimentally as

$$\alpha = \frac{N}{N-1} \frac{S-1}{S}. \quad (3)$$

It is useful to express *computing efficiency* with those experimentally measurable

$$E(N, \alpha) = \frac{S}{N} = \frac{1}{N \cdot (1 - \alpha) + \alpha} = \frac{R_{\text{Max}}}{R_{\text{Peak}}} \quad (4)$$

data. Efficiency is an especially valuable parameter, given that constructors of many parallelized sequential systems (including TOP500 supercomputers) provide the efficiency (as $R_{\text{Max}}/R_{\text{Peak}}$) of their computing system, and of course, the number of processors N in their system. Via reversing Eq. (4), the value of α_{eff} can be expressed with measured data as

$$\alpha_{\text{eff}}(E, N) = \frac{E \cdot N - 1}{E \cdot (N - 1)}. \quad (5)$$

As seen, *the efficiency of a parallelized system is a two-parameter function* (the corresponding parametrical surface is shown in Fig. 1), demonstratively underpinning that “*This decay in performance is not a fault of the architecture, but is dictated by the limited parallelism*” [36]. Furthermore, that its dependence can be correctly described by the properly interpreted Amdahl's law, rather than being an unexplained “empirical efficiency.”

2.2 Gustafson's law

Partly because of the outstanding achievements of parallelization technology, partly because of issues around practical utilization of Amdahl's law, a “weak scaling” (also called Gustafson's law [19]) was also introduced. Its assumption is that *the computing resources grow*

Fig. 1 The two-parameter efficiency surface (in function of parallelization efficiency measured by benchmark high-performance Linpack (HPL) and number of processing elements) as concluded from Amdahl's law (see Eq. 4), in first-order approximation. Some sample efficiency values for some selected supercomputers are shown, measured with benchmarks HPL and high-performance conjugate gradients (HPCG), respectively. Also, the estimated efficacy of brain simulation using conventional computing is shown

proportionally with the task size, and the speedup (using our notations) is formulated as

$$S = (1 - \alpha) + \alpha \cdot N \quad (6)$$

Similar to Amdahl's law, the efficiency can be derived for Gustafson's law as (compared to Eq. 4)

$$E(N, \alpha) = \frac{S}{N} = \alpha + \frac{(1 - \alpha)}{N}. \quad (7)$$

From these equations immediately follows that speedup (aka parallelization gain) *increases* linearly with the number of processors, *without limitation*; a conclusion that was launched amid much fanfare. They imply, however, some more immediate findings, such as

- the efficiency slightly *increases* with the number of processors N (the more processors, the better efficacy),
- the non-parallelizable portion of the task either shrinks as the number of processors grows, or *despite that it is non-parallelizable, the portion $(1 - \alpha)$ is distributed between the N processors,*
- *executing the extra machine instructions needed to organize the joint work need no time,*
- *all non-payload computing contributions such as communication (including network transfer), synchronization, input/output and memory access take no time.*

However, an error was made in deriving Eq. (6): *The $N - 1$ processors are idle waiting (see the term with subscript*

“idle” below), while the first one is executing the sequential-only portion. Because of this, the time that serves as the base for calculating² the speedup in the case of using N processors is

$$\begin{aligned} T_N &= (1 - \alpha)_{\text{processing}} + \alpha \cdot N + (1 - \alpha) \cdot (N - 1)_{\text{idle}} \\ &= (1 - \alpha) \cdot N + \alpha \cdot N \\ &= N. \end{aligned}$$

That is, before fixing the arithmetic error, strange conclusions follow, after fixing it, the conceptual mistake comes to light: “Weak scaling” assumes that single-processor efficiency can be transferred to parallelized sequential subsystems without loss. Weak scaling assumes that the efficacy of a system comprising N single-thread processors remains the same as that of a single-thread processor. This assumption strongly contradicts the experienced “empirical efficiency” (several hundred-fold deviation from its predicted value) of parallelized systems, not mentioning the “different efficiencies” [46], see also Fig. 1.

However, that “in practice, for several applications, the fraction of the serial part happens to be very, very small thus leading to near-linear speedups” [27], mislead the researchers. Gustafson concluded his “scaling” for several hundred processors only. The interplay of improving parallelization and general hardware (HW) development (including the non-determinism of modern HW [54]), covered for decades that this scaling was used far outside of its range of validity.

That is, Gustafson’s law is simply a misinterpretation of its argument α : a simple function form transforms Gustafson’s law to Amdahl’s law [34]. After making that transformation, the two (apparently very different) laws become identical. However, as suspected by [34]: “Gustafson’s formulation gives an illusion that as if N can increase indefinitely.” Although collective experience showed that it was not valid for the case of systems comprising an ever higher number of processors (an “empirical efficiency” appeared), and later researchers measured “two different efficiencies” [22] for the same supercomputer (under different workloads), the “weak scaling” was not suspected to be responsible for the issues.

“Weak scaling” omits all non-payload (but needed for the operation) activities, such as interconnection time, physical size (signal propagation time), accessing data in an amount exceeding cache size, synchronization of different kinds, that are undoubtedly present when working with ANNs. In other words, “weak scaling” neglects the temporal behavior, a crucial science-based feature of

computing [44]. This illusion led to the moon-shot of targeting to build supercomputers with computing performance well above feasible (and reasonable) size [46] and leads to false conclusions in the case of clouds. Because of this, *except some very few neuron systems*, “weak scaling” cannot be safely used for ANNs, even as a rough approximation.

2.3 Time-aware scaling

The role of α was theoretically established [24], and the phenomenon itself, that the efficiency (in contrast to Eq. 7) decreases as the number of processing units increases, is known since decades [36] (although it was not formulated in the functional form given by Eq. (4)). The “gold rush” for building exascale computers made finally evident that under the extreme conditions represented by the need of millions of processors, “weak scaling” leads to false conclusions. It had to be admitted that it “can be seen in our current situation where the historical ten-year cadence between the attainment of megaflops, teraflops, and petaflops has not been the case for exaflops” [17]. It looks like, however, that in feasibility studies of supercomputing using parallelized sequential systems, and an analysis, whether building computers of such size is feasible (and reasonable) remained (and remains) out of sight either in USA [33, 41] or in EU [15] or in Japan [16] or in China [29].

Figure 1 depicts the two-parameter efficiency surface stemming from the time-aware interpretation of Amdahl’s law. On the parametric surface, described by Eq. (4), some measured efficiencies of present top supercomputers are also depicted, only to illustrate some general rules. The HPL³ efficiencies are sitting on the surface, while the corresponding HPCG⁴ values are much below those values. The conclusion drawn here was that “the supercomputers have two different efficiencies” [22], because that experience cannot be explained in the frame of “classic computing paradigm” and/or “weak scaling.”

Supercomputers *Taihulight*, *Fugaku*, and *K computer* stand out from the “millions core” middle group. Thanks to its 0.3M cores, *K computer* has the best efficiency for HPCG benchmark, while *Taihulight* with its 10M cores the worst one. The middle group follows the rules [46]. For HPL benchmark: the more cores, the lower efficiency. It looks like the community experienced the effect of the two-dimensional efficiency. The top supercomputers run HPL benchmark with using all their cores, but some of them only use a fragment of their cores to measure performance

² For the meaning of the terms, the wording “is the amount of time spent (by a serial processor)” is used by the author in [19].

³ <http://www.netlib.org/benchmark/hpl/>.

⁴ <https://www.epcc.ed.ac.uk/blog/2015/07/30/hpcg>.

with HPCG.⁵ This reduction happens because of the inflection point: As can be concluded from the figure, increasing their nominal performance by an order of magnitude decreases their efficiency (and so: payload performance) by more than an order of magnitude.⁶ For HPCG benchmark: the “roofline” [55] of that communication intensity was already reached, all computers have about the same efficiency. For more discussion on supercomputers, see [46] (and its continuous upgrades on arXiv).

3 Performance limit of processor-based AI systems

3.1 General considerations

As discussed in [46], payload performance $P(N, \alpha)$ of parallelized systems comprising N processors is described⁷ as

$$P(N, \alpha) = \frac{N \cdot P_{\text{single}}}{N \cdot (1 - \alpha) + \alpha} \quad (8)$$

where P_{single} is the single-thread performance of individual processors and α is describing parallelization of the given system for the given workload (i.e., it depends on both of them).

This simple formula explains why *payload performance of a system is not a linear function of its nominal performance* and why in the case of very good parallelization ($(1 - \alpha) \ll 1$) and low N , this nonlinearity cannot be noticed. In contrast to the prediction of “weak scaling,” *payload performance* and *nominal performance* differ by a factor, growing with the number of cores. This conclusion is well known, but forgotten: “*This decay in performance is not a fault of the architecture, but is dictated by the limited parallelism*” [36].

The key issue is, however, that one can hardly calculate the value of α for the present complex HW/software (SW) systems from their technical data, although some estimated values can be derived. For supercomputers, however, one can derive a theoretical “best possible” and already achieved “worst case” values [46]. It gives us reasonable confidence that those values deviate only within a factor of two. We cannot expect similar results for ANNs. There are no generally accepted benchmark computations, and also there are no standard

architectures.⁸ Using a benchmark means a particular workload, and comparing the results of even a standardized ANN benchmark on different architectures is conceptually as little useful as comparing the results of benchmarks HPL and HPCG on the same architecture.

Recall also that at a large number of processors, the internal latency of processor also matters. Following the failure of supercomputer Aurora’18, Intel admitted: “*Knights Hill was canceled and instead be replaced by a “new platform and new microarchitecture specifically designed for exascale”*” [56]. We expect that shortly it shall be admitted that building large-scale AI systems is simply not possible based on the old architectural principles [26]. The potential new architectures, however, require a new computing paradigm (considering both temporal behavior of computing systems [44], and the old truth that “more is different” [2]), that can give a proper reply to power consumption and performance issues of—among others, ANN—computing.

3.2 Communication-to-computation ratio

As we learned decades ago, “*the inherent communication-to-computation ratio in a parallel application is one of the important determinants of its performance on any architecture*” [36], suggesting that communication can be a dominant contribution to system’s non-payload performance. In the case of neural simulation, a very intensive communication must take place, so the non-payload-to-payload ratio has a significant impact on the performance of ANN-type computations. That ratio and the corresponding workload type are closely related: Using a specific benchmark implies using a specific communication-to-computation ratio. In the case of supercomputing, the same workload is running on (nearly) the same type of architecture, which is not the case for ANNs. Communication is implemented through input/output (I/O) instructions; this is why “*artificial intelligence, ...it’s the most disruptive workload from an I/O pattern perspective.*”⁹

3.3 Computing benchmarks

There are two commonly used benchmarks in supercomputing. Historically, benchmark HPL was used to compare different configurations by executing a standard computing task. However, their performance depends very much on their workload type. Although HPL is excellent for racing

⁵ Beginning with June 2021, the data “Measured cores” are not provided any more, covering this aspect.

⁶ The presumable reason of this new trend is that in this way the measured HPCG/HPL ratio gets much higher, providing the illusion that vast supercomputers became more suitable for real-life tasks.

⁷ At least in a first approximation, see [46].

⁸ Notice that selecting a benchmark also directs the architectural development: The benchmarks HPL and HPCG result in different rankings.

⁹ <https://www.nextplatform.com/2019/10/30/cray-revamps-clusters-for-the-exascale-era/>.

purposes (produces high figures), its behavior on vast configurations strongly deviates from that of “real-life” tasks: The surface in Fig. 1 represents a kind of theoretical upper limit for distributed processing.

The HPL class tasks essentially need communication only at the very beginning and at the very end of the job. Real-life programs, however, usually work in a different way. Because of this reason, a couple of years ago, the community introduced benchmark HPCG: The collective experience shows that the payload performance is much more accurately approximated by HPCG than by HPL, because real-life tasks need much more communication than HPL. Importantly, since their interconnection quality improved considerably, supercomputers show different efficiencies when using various benchmark programs [22]. Their efficiencies differ by a factor of ca. 200–500 (a fact that remains unexplained in the frame of “weak scaling”), when measured by HPL and HPCG, respectively.

In the HPL class, the communication intensity is the lowest possible one: computing units receive their task (and parameters) at the beginning of computation, and they return their result at the very end. The core orchestrating their work must deal with the fellow cores only in these periods, so the communication intensity is proportional to the number of cores in the system. Notice the need to queue requests at the beginning and the end of the task.

In the HPCG class, iteration occurs: The fellow cores return the result of one iteration to the coordinator core, which makes sequential operations: receives and re-sends the parameters, and it needs to compute new parameters before sending them back to the fellow cores. The program repeats the process several times. As a consequence, *the non-parallelizable fraction of benchmarking time grows proportionally to the number of iterations*. The effect of that extra communication decreases the achievable performance roofline [55]: As shown in Fig. 2, the HPCG roofline is about 200 times smaller than the HPL one. As depicted, for benchmark HPL multiple rooflines can be located. The highest roofline can be attributed to a processor having slightly different computing principle [59]. The second highest roofline (dominated by the calculation itself) is the commonly achievable performance gain;¹⁰ their theoretical expectation matches the empirical value [46]. The third roofline shows the effect of the improved connection technology. The fourth roofline (dominated mainly by the internal interconnection) shows how much performance can be achieved without using expensive “racing” parallelization technologies. For benchmark HPCG, only one roofline exists: The calculation dominates.

¹⁰ Notice that even using coherence bus or very clever positioning of L_2 caches cannot help a lot; see the effect of the high number of processors in Fig. 1.

Fig. 2 Performance gain of supercomputers in function of their year of construction, under different workloads. Diagram lines display the measured values derived using HPL and HPCG benchmarks, for the TOP3 supercomputers in the gives years. The small black dots mark the performance data of supercomputers *JUQUEEN* and *K* as of 2014 June, for HPL and HPCG benchmarks, respectively. The big black dot denotes the payload performance of the system used by [42]. The saturation effect can be observed for both HPL and HPCG benchmarks

Neither direct interconnection of processors nor advanced interconnection technology can increase the performance gain. As discussed below, for another types of calculations (non-standard benchmarks), the roofline can be even lower: [10] measured a value as low as 30 for a certain kind of applications. The “roofline for brain simulation” is guessed from one single measurement datum. That datum was measured using a only a small fragment (about 1%) of available cores (see Sect. 3.9), so the real efficiency (see also Sect. 2.3 on benchmarking supercomputers using only a fragment of their available cores) can be about two orders of magnitude lower, i.e., in the order of a dozen of processors.

Unfortunately, an ANN workload, in general, cannot be defined, so one cannot run one single benchmark to guess the performance characteristics of a new application. The performance gain for ANNs can be guessed to be around or above that of the brain simulation,¹¹ depending on ANN’s architecture. The neural communication, using a vast number of simultaneous communication, combined with the idea of using a single high-speed bus, introduces an additional bottleneck: a very different type of workload, see Sect. 3.8. Because of this different workload, the above benchmarks cannot be used to estimate execution time of ANN-type tasks.

¹¹ Our reasoned guess is in good accordance with the experimental evidence [26]: “Strong scaling is stalling after only a few dozen nodes.”

As expressed by Eq. (8), the resulting performance of parallelized computing systems depends on both single-processor performance and performance gain. To separate these two factors, Fig. 3 displays the *performance gain* of supercomputers in the function of their year of construction and ranking in the given year. Two “plateaus” can be localized before the year 2000 and after the year 2010 also, unfortunately; underpinning Amdahl’s law and refuting Gustafson’s law, and also confirming the prediction “*Why we need Exascale and why we won’t get there by 2020*” [35]. The “hillside” reflects the enormous development of interconnection technology between the years 2000 and 2010 (for more details see [46]). For the reason of the “humps” around the beginning of the second plateau, see Sect. 3.5. Unfortunately, different individual factors (such as interconnection quality, using accelerators and clustering, using on-chip memories, or using slightly different computing paradigm, etc.) cannot be separated in this way. However, some limited validity conclusions can be drawn.

3.4 Workload type

The role of the workload came to light after that interconnection technology was greatly improved, and as a consequence, the *benchmarking computation* (defining the *type of workload*) became the dominating contributor, defining value of α (and as a consequence, payload performance), for a discussion see [46]. The overly complex Fig. 4 illustrates the phenomenon, why and how the payload performance of a configuration depends on the application it runs. Notice that at low nominal performance values, the payload performance depends linearly on the nominal performance (the blue diagram line) and only slightly depends on the workload type.

The performance breakdown shown in Fig. 4 was experimentally measured by [10, 36], [23] (Fig. 7) and [42] (Fig. 8), but it used not to be a subject of studies. The fact that speedup diagram line turns back at a critical

Fig. 3 History of supercomputing in terms of performance gain: performance values of the first 25 supercomputers, in function of year of their ranking. Data measured using benchmark HPL

number of processor, was noticed early [36]. After exceeding a critical number of cores, housekeeping gradually becomes the dominating factor of the performance limitation, and leads to a decrease in the payload performance: “*there comes a point when using more Processing Unit (PU)s ...actually increases the execution time rather than reducing it.*” In that paper (at a different workload and architecture), the achievable parallelization gain was about 8, and it was achieved using 20–30 processors. Recently, there are few experimental investigations in this direction. One of the rare exceptions is [10]. The careful systematic investigation of results of running bioinformatics applications pointed out that the speedup curve has a maximum and breaks down for a higher number of processor cores: “*The execution time and the speedup on IPDATA reach the best values within about 90 processors.* Furthermore, that ...”*the parallel version is up to 30 times faster than the serial one.*” For ANNs, it is just a few dozen [26] where “strong scaling” stalls. For different applications (workloads), these figures are of course different, but the conclusion persists.

Figure 4 shows the comparison of three workloads (owing different communication intensity). In the top and middle figures, the communication intensities of the standard supercomputer benchmarks *HPL* and *HPCG* are displayed in the style of AI networks. The “input layer” and “output layer” are the same and comprise the initiating node only, while the other “layers” are again the same: the rest of the cores. Figure 4c depicts an AI network comprising n input nodes and k output nodes; furthermore, the h hidden layers comprise m nodes. The communication-to-computation intensity [36] is, of course, not proportional in the cases of subfigures, but the figure illustrates excellently how the communication need of different computer tasks changes with the type of the workload.

As can be easily seen from the figure, in the case of benchmark *HPL*, the initiating node must issue m communication messages and collect m returned results, i.e., the execution time is $O(2m)$. In the case of benchmark *HPCG*, this execution time is $O(2Nm)$ where N is the number of iterations. (One cannot directly compare the execution times because of the different amounts of computations and the different amounts of sequential-only computations.)

Figure 4a displays the case of minimum communication and Fig. 4b a moderately increased one (corresponding to real-life supercomputer tasks). As nominal performance increases linearly and payload performance decreases reciprocally with the number of cores, at some critical value where an inflection point occurs, the resulting performance starts to fall. The resulting non-parallelizable fraction sharply decreases efficacy (in other words: performance gain or speedup) of the system [50]. This effect was noticed early [36], under different technical

Fig. 4 Different communication/computation intensities of the applications lead to different payload performance values in the same supercomputer system. Left column: models of computing intensities for different benchmarks. Right column: the corresponding payload performances and α contributions in function of the nominal performance of a fictive supercomputer ($P = 1 \text{ Gflop/s}$ @ 1 GHz).

conditions, but somewhat faded due to development of parallelization technology.

The non-parallelizable fraction (denoted in the figure by α_{eff}^X) of a computing task comprises components X of different origins. As already discussed, and was noticed decades ago, “the inherent communication-to-computation ratio in a parallel application is one of the important determinants of its performance on any architecture” [36], suggesting that *communication can be a dominant contribution to distributed system’s performance*.

The blue diagram lines refer to the right-hand scale (R_{Max} values) and all others ($(1 - \alpha_{eff}^X)$ contributions) to the left-hand scale. The figure is purely illustrating the concepts; the displayed numbers are somewhat similar to real ones

The workload in ANN systems comprises components of type “computation” and “communication” (this time also involving data access and synchronization, i.e., everything that is “non-computation”). As logical inter-dependence between those contributions is strictly defined, payload performance of the system is limited by both factors, and the same system (maybe even within the same workload, case by case) can be either computing bound and communication bound, or both.

Notice that supercomputers showing the breakdown depicted in Fig. 4 are not included in the history depicted

in Fig. 3. Aurora'18 failed, Aurora'21 semi-failed, Gyokou was withdrawn, “Chinese decision-makers decided to withhold the country’s newest Shuguang supercomputers even though they operate more than 50 percent faster than the best current US machines.”¹² Also, *Fugaku* stalled [14] at some 40% of its planned capacity.

Notice that a similar hillside cannot be drawn for benchmark HPCG, because of two reasons. On the one side, HPCG measurement started only a few years ago. On the other side, top supercomputers publish data measured with cores less than the number of cores used to measure HPL. Recall that efficiency is a two-parameter function (Fig. 1). For “real-life” programs, represented by HPCG, this critical number is much lower than in the case of HPL. There is a real competition between the different contributions to dominate system’s performance. As demonstrated in Fig. 6 in [46], before 2010, running both benchmarks HPL and HPCG on a top supercomputer was a *communication-bound* task, since 2010 HPL is a computing-bound task, while HPCG persisted to be a *communication-bound* task. This is why some supercomputers provide their HPCG efficiency measured only with a fragment of their cores: The HPCG “roofline” is reached at that lower number of cores. Adding more cores does not increase their payload performance, but decreases their efficiency.

3.5 Accelerators

As a side effect of “weak scaling,” it is usually presumed that decreasing the time needed for the payload contribution affects the efficiency of ANN systems linearly. However, it is not so. As discussed in [47], we also change the non-payload-to-payload ratio that defines system’s efficiency. We mention two prominent examples here: using shorter operands (move less data and perform less bit manipulations) and to mimic the operation of a neuron in an entirely different way: using quick analog signal processing rather than slow digital calculation.

It is a common fallacy that benchmark *HPL-AI* is benchmark for AI systems. Actually, it means “*The High-Performance LINPACK for Accelerator Introspection*” (HPL-AI) and that benchmark seeks to highlight the convergence of HPC and artificial intelligence (AI) workloads.¹³ It has not much to do with AI, except that it uses the operand length common in AI tasks. HPL, similarly to AI, is a *workload type*.¹⁴ Researchers succeeded to achieve more than three times better performance gain (3.01 for

Summit and 3.42 for *Fugaku*), that (as correctly stated in the announcement) “*Achieving a 445 petaflops mixed-precision result on HPL (equivalent to our 148.6 petaflops DP result)*” [20], i.e., the peak DP performance did not change. However, this naming convention suggests the illusion that when using supercomputers for AI tasks and using half-precision, one can expect this payload performance.

Unfortunately, this achievement comes from accessing less data in memory and using quicker operations on shorter operands rather than reducing communication intensity. For AI applications, limitations remain the same as described above, except that when using mixed precision, the power efficiency shall be better by a factor of nearly four, compared to the power efficiency measured using double precision operands.¹⁵

We expect that when using half-precision (FP16) rather than double precision (FP64) operands in the calculations, four times less data are transferred and manipulated by the system. The measured power consumption data underpin the statement. However, system’s computing performance is only slightly more than three times higher than using 64-bit (FP64) operands. The nonlinearity has its effect even in this simple case. (Recall that HPL uses minimum communication.) In the benchmark, the housekeeping activity (data access, indexing, counting, addressing) also takes time. Concerning the temporal behavior [44] of the operation, in this case, the data transmission time T_t is the same, the data processing time (due to the shorter length of operands) changes, and so the apparent speed changes nonlinearly. Even, the measured performance data enabled us to estimate execution time with zero precision.¹⁶ (FP0) operands, see [46].

Another plausible assumption is that if we use quick analog signal processing to replace the slow digital calculation, as proposed in [9, 37, 53], our computing system gets proportionally quicker. Presumably, on systems comprising just a few neurons, one can measure a considerable, but less than expected, speedup. The housekeeping becomes more significant than in the case of purely digital processing. In a hypothetical measurement, the speedup would be much less than the ratio of the corresponding analog/digital processing times, even in the case of HPL benchmark. Recall that here the workload is of AI type, with much worse parallelization (and nonlinearity). As a consequence, one cannot expect a considerable speedup in

¹² <https://www.scmp.com/tech/policy/article/3015997/china-has-decided-not-fan-flames-super-computing-rivalry-amid-us>.

¹³ <https://www.icl.utk.edu/hpl-ai/>.

¹⁴ Even <https://www.top500.org/lists/top500/2020/06/> mismatches operand length and workload: “In single or further reduced precision, which are often used in machine learning and AI applications, *Fugaku*’s peak performance is over 1000 petaflops (1 exaflops).”

¹⁵ Similarly, exchanging data directly between processing units [59] (without using the global memory) also enhances α (and payload performance) [3], but it represents a (slightly) different computing paradigm.

¹⁶ Without dedicated measurement, no more accurate estimations are possible.

large neuromorphic systems. For a detailed discussion of introducing new effect/technologies/materials, see [44].

It sounds good that “*The analog memristor array is effectively the neural network laid out in the form of a crossbar, which can perform the entire operation in one clock cycle*” [25]. In brackets, however fairly added, that “(not counting the clock cycles that may be required to fetch and store the input and output data).” Yes, all operands of the memristor array must be transferred to its input section (and previously, they must be computed or otherwise produced); furthermore, the results must be transferred from its output section to their destination. The effective computing time of the memristor-related operations shall be compared to conventional operations’ effective time, from the beginning to the end of the computing operation, to make a fair comparison. (The problem persists even if continuous-time data representation [53] is used.)

The temporal behavior of components and their materials can easily be misidentified in the time-unaware model. Five decades ago, even *memristance* has been introduced [38] as a fundamental electrical component, meaning that the memristor’s electrical resistance is not constant but depends on the history of current that had previously flowed through the device. There are, however, some serious doubts as to whether a genuine memristor can actually exist in physical reality [1]. In light of our analysis, some temporal behavior definitely exists; the question is how much it is related to material or biological features, if our time-aware computing method is followed.

Besides, adding analog components to a digital processor has its price. Given that a digital processor cannot handle resources outside of its world, one must call the operating system (OS) for help. That help, however, is expensive in terms of execution time. The required context switching takes time in the order of executing 10^4 instructions [11, 40], which dramatically increases the time of housekeeping and the total execution time. This effect makes the systems’ non-payload-to-payload ratio much worse than it was before introducing that enhancement.

3.6 Timing of activities

In ANNs, the data transfer time must be considered seriously. In both biological and electronic systems, both the distance between entities of the network, and the signal propagation speed are finite. Because of this, in physically large-sized and/or intensively communicating systems, the “idle time” of processors defines the final performance that a parallelized sequential system can achieve. In conventional computing systems, “data dependence” limits achievable parallelism: We must compute data before

using it as an argument for another computation. Although, of course, also in conventional computing, computed data must be delivered to the place of their second utilization, thanks to “weak scaling” [19], this “communication time” is neglected. For example, scaling of matrix operations and “sparsity,” mentioned in [25], work linearly *only* if data transfer time is neglected.

Timing plays an important role in all levels of computing, from gate-level processing to clouds connected to the Internet. In [44, 52], the example describing temporal operation of a one-bit adder provides a nice example, that although the *line-by-line compiling (sequential programming, called also Neumann-style programming [4])*, formally introduces only logical dependence, through its technical implementation it implicitly and inherently introduces a temporal behavior, too.

In neuromorphic computing, including ANNs, the transfer time is a vital part of information processing. A biological brain must deploy a “speed accelerator” to ensure that the control signals arrive at the target destination before the arrival of the controlled messages, despite that the former derived from a distant part of the brain [8]. *This aspect is so vital in biology that the brain deploys many cells with the associated energy investment to keep the communication speed higher for the control signal.* Computer technology cannot speed up communication selectively, as in biology. It is also impossible to keep part of the system for a lower speed selectively: The propagation speed of electromagnetic waves is predefined. However, as discussed in [47], *handling data timing adequately, is vital, especially for bio-mimicking ANNs.*

3.7 The layer structure

The bottom part of Fig. 4 depicts how ANNs are supposed to operate. The life begins in several input channels (rather than one as in HPL and HPCG cases) that would be advantageous. However, the system must communicate its values to *all* nodes in the top hidden layer: The more input nodes and the more nodes in the hidden layer(s), the many *times* more communication is required for the operation. The same situation also happens when the first hidden layer communicates data to the second one, except that here *the square of the number of nodes* is to be used as a weight factor of communication.

Initially, n input nodes issue messages, each one m messages (queuing#1) to nodes in the first hidden layer, i.e., altogether nm messages. If one uses a commonly used shared bus to transfer messages, these nm messages must be queued (queuing#2). Also, every single node in the hidden layer receives (and processes) m input messages (queuing#3). Between hidden layers, the same queuing is repeated (maybe several times) with nm messages, and

finally, km messages are sent to the output nodes. During this process, the system queues messages (at least) three times. Notice that using a single high-speed bus, because of the needed arbitration, drastically increases the transfer time of the individual messages and furthermore changes their timing, see Sect. 3.8.

To make a fair comparison with benchmarks *HPL* and *HPCG*, let us assume one input and one output node. In this case, the AI execution time is $O(h \times m^2)$, provided that the AI system has h hidden layers. (Here, we assumed that messaging mechanisms between different layers are independent. It is not so if they share a global bus.)

For a numerical example: Let us assume that in supercomputers, 1M cores are used. In AI networks, 1K nodes are present in the hidden layers, and only one input and output nodes are used. In that case, all execution times are $O(1M)$. (Again, the amount of computation is sharply different, so the scaling can be compared, but not the execution times.) This communication intensity explains why in Fig. 3 the *HPCG* “roofline” falls hundreds of times lower than that of the *HPL*: The increased communication need strongly decreases systems’ achievable performance gain.

Notice that the number of *computation* operations increases with m , while number of *communication* operations with m^2 . In other words, the more nodes in the hidden layers, the higher is their communication intensity (communication-to-computation ratio), and because of this, the lower is the efficiency of the system. Recall that since AI nodes perform simple computations, compared to the functionality of supercomputer benchmarks, their communication-to-computation ratio is much higher, making their efficacy even worse. The conclusions are underpinned by experimental research [26]:

- “Strong scaling is stalling after only a few dozen nodes.”
- “The scalability stalls when the compute times drop below the communication times, leaving compute units idle, hence becoming a communication bound problem.”
- “The network layout has a large impact on the crucial communication/compute ratio: shallow networks with many neurons per layer ...scale worse than deep networks with less neurons.”

3.8 Using high-speed bus(es)

As discussed in connection with the reasoning, why the internal temporal ratio between transporting and processing data has significantly changed [44]: Moore’s observation is (was) valid for electronic density only, but not valid for

connecting technology, such as buses. On the one side, because of the smaller physical size and the quicker clock signal, on the other side the unchanged cm -long bus cable, using a serial bus means spending the overwhelming majority of apparent processing time with arbitration (see the temporal diagram and case study in [44]), so using a sequential bus is at least questionable in large-scale systems: *The transfer time is limited by the needed arbitration (increases with the number of neurons!) rather than by the bus speed. “The idea of using the popular shared bus to implement the communication medium is no longer acceptable, mainly due to its high contention” [13].*

The massively “bursty” nature of the data (when imitating biological neural network, the different nodes of the layer want to communicate simultaneously) also makes the case harder. Communication circuits receive the task of sending data to N other nodes. What is worse, bus arbitration, addressing and latency prolong the transfer time (and decrease the system’s efficacy). This type of communicational burst may easily lead to a “communicational collapse” [31], but it may also produce unintentional “neuronal avalanches” [5].

The fundamental issue is replacing the private communication channel between biological neurons with the mandatory use of some kind of shared media in technological neurons. As discussed in [44], at a large number of communicating units, sharing the medium becomes the dominant contributor to the time consumption of computing. Its effect can be mitigated using different technological implementations. Still, the conclusion persists: Its technical implementation of neuronal communication defines *the payload computational efficiency of an ANN, and the computing performance of its nodes has only marginal importance.* The lengthy queueing also leads to unrealistic timings: as discussed in Sect. 3.9, some physically (not biologically or logically) delayed signals must be dropped to provide a seemingly acceptable computing performance. Another (wrong) approach to solving the same problem is introducing conditional computation [7] as discussed in Sect. 3.12. The origin of the issue is that using a shared medium makes computing system’s temporal behavior much more emphasized and that *the temporal behavior cannot be compensated using methods developed having a timeless behavior in mind.*

The ANN-type workload introduces an extra handicap: The neurons must communicate the result of an elementary operation after every single operation, in addition through a single shared bus. In Fig. 5, the inset shows a simple neuromorphic use case: One input neuron and one output neuron communicate through a hidden layer, comprising only two neurons. Figure 5a mostly shows *the biological implementation: All neurons are directly wired to their partners, i.e., a system of “parallel buses” (axons) exists.*

Notice that the operating time also comprises two “non-payload” times (T_i): data input and data output, which coincide with the non-payload time of the other communication party. The diagram displays the logical and temporal dependencies of the neuronal functionality. The payload operation (“the computing”) can only start after its data are delivered (by the, from this point of view, non-payload functionality: input-side communication), and the output communication can only begin when the computing is finished. Importantly, *communication and calculation mutually block each other*. Two important points that neuromorphic systems must mimic noticed immediately: i/ *the communication time is an integral part of the total execution time* and ii/ *the ability to communicate is a native functionality of the system*. In such a parallel implementation, *the performance of the system, measured as the resulting total time (processing + transmitting), scales linearly with increasing either the non-payload communication speed or the payload processing speed*.

Figure 5b shows a *technical implementation of a high-speed shared bus* for communication. To the right of the grid, the activity that loads the bus at the given time is shown. A double arrow illustrates the communication bandwidth, the length of which is proportional to the number of packages the bus can deliver in a given time unit. We assume that the input neuron can send its information in a single message to the hidden layer; furthermore, the processing by neurons in the hidden layer both starts and ends simultaneously. However, the neurons must compete for accessing the bus, and only one of them can send its message immediately, but the other(s) must wait

until the bus gets released. The output neuron can only receive the message when the first neuron completed its sending. Furthermore, the output neuron must first acquire the second message from the bus, and the processing can only begin after having both input arguments. *This constraint results in sequential bus delays both during non-payload processing in the hidden layer and payload processing in the output neuron*. Adding one more neuron to the layer introduces one more delay.

At this point, two wrong solutions can be taken: Either the second neuron must wait until the second input arrives (in biology, a spike also carries a synchronization signal, and triggers its integration), or (in “technical neurons,” using continuous levels rather than pulses, this synchronization facility is omitted) changes its output continuously, as the inputs arrive, and its processing speed enables. In the latter case, however, *until the second input arrives (and gets processed) the neuron provides an output signal, differing from the one expected based on the mathematical dependence*. As discussed in detail in [44], this temporarily may be wrong, output signal is known in the electronics, and those “glitches” are eliminated via using a “worst-case” delay for the output signal. However, including a serial bus in that computation would enormously prolong the needed “worst-case” delay.

Using the formalism introduced above, $T_t = 2 \cdot T_B + T_d + X$, i.e., the bus must be reached in time T_B (not only the operand delivered to the bus, but also waiting for arbitration: the right to use the shared bus), twice, plus the physical delivery T_d through the bus. The X denotes “foreign contribution”: If the bus is not

Fig. 5 Implementing neuronal communication in different technical approaches. For legend’s details, see text. **a** (the biological implementation): the parallel bus; **b, c** (the technical implementation): the shared serial bus, before and after reaching the communication “roofline” [55]

dedicated for “neurons in this layer only,” any other traffic also loads the bus: both messages from different layers and the general system messages may make processing slower (and add their contribution to faking the imitated biological effect).

Even if only one single neuron exists in the hidden layer, it must use the mechanisms of sharing the bus, case by case. The physical delivery to the bus takes more time than a transfer to a neighboring neuron. (Both the arbiter and the bus are in *cm* distance range, meaning several *nsec* transfer times, while the direct transfer between the connected gates may be in the *psec* range.) If we have more neurons (such as a hidden layer) on the bus and work in parallel, they must all wait for the bus. The high-speed bus is very slightly loaded when only a couple of neurons are present. Its load increases linearly with the number of neurons in the hidden layer (or, maybe, all neurons in the system). The temporal behavior of the bus, however, is different.

Under a biology-mimicking workload, the second neuron must wait for all its inputs originating in the hidden layer. If we have L neurons in the hidden layer, the transmission time of the neuron behind the hidden layer is $T_t = L \cdot 2 \cdot T_B + T_d + X$. This temporal behavior explains why “*shallow networks with many neurons per layer ...scale worse than deep networks with less neurons*” [26]: *The physical bus delivery time T_d , as well as the processing time T_p , becomes marginal if the layer forces to make many arbitrations to reach the bus*: The number of the neurons in the hidden layer defines the transfer time. In deeper networks, the system sends its messages at different times in its different layers (and even they may have independent buses between the layers), although *the shared bus persists in limiting the communication*. Notice that there is no way to organize the message traffic: Only one bus exists.

At this point comes into picture the role of the workload on the system: The two neurons in the hidden layer want to use the single shared bus, at the same time, for communication. As a consequence, *the apparent processing time is several times higher, than the physical processing time, and it increases linearly with the number of neurons in the hidden layer* (and, maybe, with also the total number of neurons in the system, if a single high-speed bus is used).

The ratio of the time spent with forwarding data on the high-speed bus gradually decreases as the system’s size increases. *In vast systems, especially when attempting to mimic neuromorphic workload, the speed of the bus is getting marginal*. Notice that the times shown in the figure are not proportional: The (temporal) distances between cores are in the several picoseconds range, while the bus (and the arbiter) is at a distance well above nanoseconds, so *the actual temporal behavior (and the idle time stemming from it) is much worse than the figure suggests*. “*The idea*

of using the popular shared bus to implement the communication medium is no longer acceptable, mainly due to its high contention.” [13]. The extraordinary workload of AI makes it much harder to operate the systems.

3.9 The “quantal nature of computing time”

One of the famous cases demonstrating existence and competition of those limitations in the fields of AI is the research published in [42]. The systems used in the study were a HW simulator [18] explicitly designed to simulate 10^9 neurons (10^6 cores and 10^3 neurons per core) and many-thread simulation running on a supercomputer [28] able to simulate 2×10^8 neurons (the authors mention 2×10^3 neurons per core and supercomputers having 10^5 cores), respectively. The experience, however, showed [42] that scaling stalled at 8×10^4 neurons, i.e., about four orders of magnitude less than expected. *They experienced stalling about the same number of neurons, for both the HW and the SW simulator*.

Given that supercomputers have a performance limit [46], one can comprehend the former experience: The brain simulation needs massive communication (the authors estimated that $\approx 10\%$ of the execution time was spent with non-payload activity), that sharply decreases their achievable performance, so their system reached the maximum payload performance that their $(1 - \alpha)$ enables: The sequential portion was too high. But why the purpose-built brain simulator cannot reach its maximum expected performance? Is it just an accident that they both stalled at the same value, or some other limiting factor came into play? Paper [43] gives the detailed explanation.

The short reply is that digital systems, including brain simulators, have a central clock signal representing an inherent performance limit: No action in the system can happen in a shorter time. The total time divided by the clock period’s length defines maximum performance gain [46] of a system. If the size of the clock period is the commonly used 1 ns , and measurement time (in the case of supercomputers) is in the order of several hours, clocking does not mean a limitation.

Computational time and biological time are not only not equal, but they are also not proportional. To synchronize the neurons periodically, a “time grid,” commonly with 1 ms integration time, was introduced. The systems use this grid time to put the free-running artificial neurons back to the biological time scale, i.e., they act as a clock signal: simulation of the next computation step can only start when this clock signal arrives. This action is analogous with introducing a clock signal for executing machine instructions: The processor, even when it is idle, cannot begin the execution of its next machine instruction until this clock

signal arrives. That is, *in this case, the clock signal is 10^6 times longer than the clock signal of the processor*. Just because neurons must work on the same (biological) time scale, when using this method of synchronization, the (commonly used) 1 millisecond “grid time” has a noticeable effect on the payload performance.¹⁷

The brain simulation measurement [42] enables us to guess the efficacy of ANNs. Given that using more cores only increased the *nominal* performance (and, correspondingly, its power consumption) of their system, the authors decided to use only a small fragment of their resources, only 1% of the cores available in the HW simulator. In this way, we can place the efficiency of brain simulation on a supercomputer benchmarking scale, see Fig. 2. Under those circumstances, as witnessed by Fig. 2, a performance gain about 10^3 was guessed for brain simulation. Notice that the large-scale supercomputers use about 10% of their cores in HPCG measurement, see also Fig. 1. The difference in the efficiency values of HPL, HPCG, and the brain comes from the different workloads, which is the reason for the issue. For ANNs, the efficiency can be similar to that of brain simulation, somewhat above the performance gain of HPCG, because the measurement time is shorter than that of HPCG on supercomputers, so the corresponding inherent non-parallelizable portion is higher.

Recall also the “communicational collapse” from the previous section: Even if communication packages are randomized in time, it represents a colossal peak traffic, mainly if a single global (although high speed) bus is used. This effect is so strong in large systems that emergency measures must have been introduced, see Sect. 3.12. In smaller ANNs, it was found [42] that *only a few dozens of thousands of neurons can be simulated on processor-based brain simulators*. This experience includes both many-thread software simulators and a purpose-built brain simulator.¹⁸ Recall also from [26] that “*strong scaling is stalling after only a few dozen nodes*.” For a discussion on the effect of serial bus in ANNs, see [44].

3.10 Rooflines of ANNs

As all technical implementations, computing also has technological limitations. The “roofline” model [55] successfully describes that until some needed resource exceeds

its technical limitation, utilization of that resource shows a simple linear dependency. Exceeding that resource is impossible: Usage of the resource is stalling at the maximum possible level; the two lines form a “roofline.” In a complex system, such as a computing system, the computing process uses different resources, and under other conditions, various resources may dominate in defining the “roofline(s) of computing process,” see Fig. 2. An example is running benchmarks HPL and HPCG: As discussed in [46], either computing or interconnection dominates the payload performance.

As Sect. 3.9 discusses, in some cases, a third competitor can also appear on the scene, and even it can play a significant role. That is, it is not easy at all to describe an ANN system in terms of the “roofline” [55] model: Depending on the actual conditions, the dominant player (the one that defines the top level of the roofline) may change. Anyhow, it is sure that the contribution of the component, representing the lowest roofline, shall dominate. Still, the competition of parts may result in unexpected issues (for example, see how computation and interconnection changed their dominating rule, in [46]). Because of this, Fig. 2 has limited validity. It provides, however, a feeling that 1/ for all workflow types a performance plateau exists and already reached; 2/ what value of payload performance gain value can be achieved for different workloads; and 3/ where the payload efficiency of the particular kinds of ANNs, brain simulation on supercomputers, is located compared to those of the standard benchmarks (a reasoned guess).

3.11 Role of parameters of computing components

As discussed in [46], *different components of computing systems mutually block each other's operation*. Datasheet parameters of components represent a hard limit, valid for ideal, stand-alone measurements, where the utilization is 100%. When they must cooperate with other components, the way as they cooperate (aka the workload of the system) defines their soft limit (degrades utilization of units): until its operand(s) delivered, computing cannot start; until computed, result transmission cannot start, mainly when several computing units compete for the shared medium. This competition is the reason why ANNs represent a very specific workload, where weaknesses of principles of computing systems are even more emphasized.

3.12 Training ANNs

One of the most shocking features of ANNs is their weeks-long training time, even for (compared to the functionality of brain) simple tasks. The mathematical methods, of

¹⁷ This periodic synchronization shall be a limiting factor in large-scale utilization of processor-based artificial neural chips [12, 32], although thanks to their ca. thousand times higher “single-processor performance,” only when approaching the computing capacity of (part of) the brain, or when the simulation turns to be communication bound.

¹⁸ Despite this, Spinnaker2, this time with 10M processors is under construction [30].

course, do not comprise time dependence, the technical implementation of ANNs, however, does: As their time dependence is discussed in [44], the delivery times of new neuronal outputs (that serve as new neuronal inputs at the same time) are only loosely coupled: The assumption that producing an output means at the same time producing an input for some other neuron, works only in the timeless “classic computing” (and in biology using parallel axons), see discussing the temporal behavior of the serial bus in [44].

To comprehend what change of considering temporal behavior means, consider the temporal diagram of a 1-bit adder in [44]. When using adders, we have a fixed time when we read out the result. We are not interested in “glitches,” so we set a maximum time until all bits relaxed, and (at the price of losing some performance), we will receive the final result only; the adder is synchronized.

The case of ANNs, however, is different. In the adder, outputs of bit n are the inputs of bit $n + 1$ (there is no feedback), and the bits are wired directly. In ANNs, the signals are delivered via a bus, and the interconnection type and sequence depend on many factors (ranging from the kind of task to the actual inputs). During training ANNs, their feedback complicates the case. The only fixed thing in timing is that the neuronal input arrives inevitably only after a partner produced it. The time ordering of delivered events, however, is not sure: It depends on technical parameters of delivery, rather than on the logic that generates them. Time stamping cannot help much. There are two bad choices. Option one is that neurons should have a (biological) sending time-ordered input queue and begin processing when all partner neurons have sent their message. That needs a synchronous signal and leads to severe performance loss (in parallel with the one-bit adder). Option two is that they have a (physical) arrival-time ordered queue, and they are processing the messages as soon as they arrive. This technical solution enables us to give feedback to a neuron which fired later (according to its timestamp), and set a new neuronal variable state, which is a “future state” when processing a message received physically later, but with a timestamp referring to a biologically earlier time. A third, maybe better, option would be to maintain a biological-time ordered queue, and either in some time slots (much shorter than the commonly used “grid time”) send out output and feedback, individually process the received events, and send back feedback and output immediately. In both cases, it is worth to consider if their effect is significant (exceeds some tolerance level compared to the last state) and mitigate the need for communication also in this way.

We start showing an input during training, and the system begins to work, using the synaptic weights valid before showing that input. Those weights may be

randomized, or maybe that they correspond to the previous input data. The signals that the system sends, are correct, but a receiver does not know the future: It processes a signal only after it was physically delivered,¹⁹ meaning that it (and its dependents) may start to adjust their weights to a state which is still undefined. In [44], the first AND gate has quite a short indefinite time, but the OR has a long one.

When playing chess against its opponent, a faster computer can be advantageously used to analyze the forthcoming moves. Even it can compute all future moves before its opponent makes the next move. However, with publishing its next move, it must wait the next move made by its opponent; otherwise, it may publish a wrong move. Sending the feedback to its opponent as soon as the next move is computed—i.e., without synchronization—results in a quickly computed but maybe wrong move. Without synchronization, the faster is the computer, the worse is its performance as chess player.

At the beginning of their operation, some component neurons of the network may have undefined states and weights. Their operation is essentially an iteration, where—without synchronization—*the actors mostly use mostly wrong input signals, and surely adjust their weights to false signals initially and with significant time delay at later times*. If we are lucky (and consider that we are working with unstable states in the case of more complex systems), the system will converge, but painfully slowly. Or not at all. *Not considering the temporal behavior of the network leads to painfully slow and doubtful convergence.*

Synchronization is a must, even in ANNs. We must take care when using accelerators, feedback, and recurrent networks. The time matters. Computing neuronal results faster, to provide feedback more quickly, cannot help much, if at all. Delivering feedback information also needs time and uses the same shared medium, with all its disadvantages. In biology, the “computing time” and the “communication time” are in the same order of magnitude. In computing, the communication time is very much longer than computation, that is, *the received feedback refers to a time (and the related state variables) that was valid a very long time ago*. In biology, spiking is also a “look at me” signal: the feedback shall be sent to *that* neuron, reflecting the change its output caused.²⁰ Without this, neurons receive feedback about “the effect of all fellow neurons, including me.” Receiving a spike defines the time of the beginning of signal’s validity; “leaking” also defines their “expiration time.” When using spiking networks, their temporal behavior is vital.

¹⁹ Even if the message envelope contains a time stamp.

²⁰ See the Hebbian learning: neuron uses its inputs and output, exclusively.

In excessive systems, some result/feedback events must be dropped because of long queuing to provide seemingly higher performance. The logical dependence that the feedback is computed from the results of the neuron that receives the feedback, the physical implementation of the computing system converts to time dependence [44]. Because of this time sequence, the feedback messages will arrive at the neuron later (even if at the same biological time, according to their time stamp they carry), so they stand at the end of the queue. Because of this, it is highly probable that they “*are dropped if the receiving process is busy over several delivery cycles*” [42]. In vast systems, the feedback in the learning process involves results based on undefined inputs, and the calculated and (maybe correct) feedback may be neglected.

An excellent “experimental proof” of the claims above is provided in [7]. With the words of that paper: “*Yet the task of training such networks remains a challenging optimization problem. Several related problems arise: very long training time (several weeks on modern computers, for some problems), the potential for over-fitting (whereby the learned function is too specific to the training data and generalizes poorly to unseen data), and more technically, the vanishing gradient problem.*” “*The immediate effect of activating fewer units is that propagating information through the network will be faster, both at training as well as at test time.*” This effect also means that the computed feedback, based maybe on undefined inputs, reaches the previous layer’s neurons faster. A natural consequence is that (see their Fig. 5): “*As λ_s increases, the running time decreases, but so does performance.*” Similarly, introducing the spatiotemporal behavior of ANNs, even in its simple form, using separated (i.e., not connected in the way proposed in [44]) time and space contributions to describe them, significantly improved the efficacy of video analysis [57].

The role of time (mismatching) is confirmed directly, via making investigations in the time domain. “*The CNN models are more sensitive to low-frequency channels than high-frequency channels*” [58]: The feedback can follow the slow changes with less difficulty compared to the faster changes.

4 Summary

The strongly simplified computing paradigm, proposed by von Neumann, surely has severe limitations when applied to today’s technology. According to von Neumann, it is doubly *unsound* if one attempts to mimic neural operation based on a paradigm that is *unsound* for that goal, on a technological base (other than vacuum tubes) for which the paradigm is *vitiating*. As predicted by Amdahl, large (many-

processor) machines have inherent disadvantage in computing. Hennessy pointed out that the efficiency of distributed systems is strongly limited and heavily depends on their workload.

The operating characteristics of ANNs are practically unknown, mainly because of their mostly proprietary design/documentation. We reviewed some general features of ANNs, with the goal to provide help in designing new systems and to understand their scaling behavior. The existing theoretical predictions and measured results show good agreement, but dedicated measurements using well-documented benchmarks and a variety of well-documented architectures are needed. The low efficacy of our designs forces us to change our design methods. On the one side, it requires a careful design method when using existing components (i.e., to select the “least wrong” configuration; millions of devices shall work with low energy and computational efficacy!). On the other side, it urges working out a different computing paradigm (and architecture based on it).

Acknowledgements The paper is the extended version of the presentation held at the 22nd Int’l Conf on Artificial Intelligence (ICAI’20), Las Vegas, USA, as ICA2246.

Author contributions Not applicable.

Funding Information Project No. 136496 has been implemented with the support provided from the National Research, Development and Innovation Fund of Hungary, financed under the K funding scheme.

Availability of data and material Not applicable.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Code availability (software application or custom code) Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abraham I (2018) The case for rejecting the memristor as a fundamental circuit element. *Sci Rep* 8:10972. <https://doi.org/10.1038/s41598-018-29394-7>
- Anderson PW (1972) More is different. *Science* 177:393–396. <https://doi.org/10.1126/science.177.4047.393>
- Ao Y, Yang C, Liu F, Yin W, Jiang L, Sun Q (2018) Performance Optimization of the HPCG Benchmark on the Sunway TaihuLight Supercomputer. *ACM Trans Archit Code Optim* 15(1):11.1–11:20
- Backus J (1978) Can programming languages be liberated from the von Neumann style? A functional style and its algebra of programs. *Commun ACM* 21:613–641
- Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. *J Neurosci* 23(35):11167–11177. <https://doi.org/10.1523/JNEUROSCI.23-35-11167.2003>
- Bell G, Bailey DH, Dongarra J, Karp AH, Walsh K (2017) A look back on 30 years of the Gordon bell prize. *Int J High Perform Comput Appl* 31(6):469–484. <https://doi.org/10.1177/1094342017738610>
- Bengio E, Bacon PL, Pineau J, Precu D (2016) Conditional Computation in Neural Networks for faster models. In: *ICLR'16*. arXiv:1511.06297
- Buzsáki György G, Wang X-J (2012) Mechanisms of gamma oscillations. *Ann Rev Neurosci* 3(4):19:1–19:29. <https://doi.org/10.1146/annurev-neuro-062111-150444>
- Chicca E, Indiveri G (2020) A recipe for creating ideal hybrid memristive-CMOS neuromorphic processing systems. *Appl Phys Lett* 116(12):120501. <https://doi.org/10.1063/1.5142089>
- D'Angelo G, Rampone S (2014) Towards a HPC-oriented parallel implementation of a learning algorithm for bioinformatics applications. *BMC Bioinf* 15(Suppl 5):S2 <http://www.biomedcentral.com/1471-2105/15/S5/S2>
- David FM, Carlyle JC, Campbell RH (2007). Context switch overheads for Linux on ARM platforms. In: *Proceedings of the 2007 workshop on experimental computer science, ExpCS '07*. ACM, New York. <https://doi.org/10.1145/1281700.1281703>
- Davies M et al (2018) Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38:82–99
- de Macedo Mourelle L, Nedjah N, Pessanha FG (2016) Reconfigurable and adaptive computing: theory and applications, chap. 5 interprocess communication via crossbar for shared memory systems-on-chip. CRC Press, London
- Dongarra J (2016) Report on the Fujitsu Fugaku System. Tech Report ICL-UT-20-06, University of Tennessee Department of Electrical Engineering and Computer Science
- European Commission: Implementation of the Action Plan for the European High-Performance Computing strategy. http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=15269 (2016)
- Extremtech: Japan Tests Silicon for Exascale Computing in 2021. <https://www.extremtech.com/computing/272558-japan-tests-silicon-for-exascale-computing-in-2021> (2018)
- Feldman M (2019) Exascale is not your grandfather's HPC. <https://www.nextplatform.com/2019/10/22/exascale-is-not-your-grandfathers-hpc/>
- Furber SB, Lester DR, Plana LA, Garside JD, Painkras E, Temple S, Brown AD (2013) Overview of the SpiNNaker system architecture. *IEEE Trans Comput* 62(12):2454–2467
- Gustafson JL (1988) Reevaluating Amdahl's Law. *Commun. ACM* 31(5):532–533. <https://doi.org/10.1145/42411.42415>
- Haidar A, Tomov S, Dongarra J, Higham NJ (2018) Harnessing GPU tensor cores for fast FP16 arithmetic to speed up mixed-precision iterative refinement solvers. In: *Proceedings of the international conference for high performance computing, networking, storage, and analysis, SC '18*. IEEE Press, pp 47:1–47:11
- Hutson M (2020) Core progress in AI has stalled in some fields. *Science* 368:6494/927. <https://doi.org/10.1126/science.368.6494.927>
- IEEE Spectrum: Two Different Top500 Supercomputing Benchmarks Show Two Different Top Supercomputers (2017). <https://spectrum.ieee.org/tech-talk/computing/hardware/two-different-top500-supercomputing-benchmarks-show-two-different-top-supercomputers>
- Ippen T, Eppler JM, Plesser HE, Diesmann M (2017) Constructing neuronal network models in massively parallel environments. *Front Neuroinform* 11:30
- Karp AH, Flatt HP (1990) Measuring parallel processor performance. *Commun ACM* 33(5):539–543. <https://doi.org/10.1145/78607.78614>
- Kendall JD, Kumar S (2020) The building blocks of a brain-inspired computer. *Appl Phys Rev* 7:011305. <https://doi.org/10.1063/1.5129306>
- Keuper J, Pfreundt FJ (2016). Distributed training of deep neural networks: theoretical and practical limits of parallel scalability. In: *2nd Workshop on machine learning in HPC environments (MLHPC)*. IEEE, pp 1469–1476. <https://doi.org/10.1109/MLHPC.2016.006>. <https://www.researchgate.net/publication/308457837>
- Krishnaprasad S (2001) Uses and abuses of Amdahl's law. *J Comput Sci Coll* 17(2):288–293
- Kunkel S, Schmidt M, Eppler JM, Plesser HE, Masumoto G, Igarashi J, Ishii S, Fukai T, Morrison A, Diesmann M, Helias M (2014) Spiking network simulation code for petascale computers. *Front Neuroinform* 8:78. <https://doi.org/10.3389/fninf.2014.00078>
- Liao X et al (2018) Moving from exascale to zettascale computing: challenges and techniques. *Front Inf Technol Electron Eng* 19(10):1236–1244. <https://doi.org/10.1631/FITEE.1800494>
- Liu C, Bellec G, Vogginger B, Kappel D, Partzsch J, Neumärker F, Höppner S, Maass W, Furber SB, Legenstein R, Mayr CG (2018) Memory-efficient deep learning on a SpiNNaker 2 prototype. *Front Neurosci* 12:840. <https://doi.org/10.3389/fnins.2018.00840>
- Moradi S, Manohar R (2018) The impact of on-chip communication on memory technologies for neuromorphic systems. *J Phys D Appl Phys* 52(1):014003
- Sawada J et al (2016) TrueNorth ecosystem for brain-inspired computing: scalable systems, software, and applications. In: *SC '16: proceedings of the international conference for high performance computing, networking, storage and analysis*, pp 130–141
- Service RF (2018) Design for U.S. exascale computer takes shape. *Science* 359:617–618
- Shi Y (1996) Reevaluating Amdahl's law and Gustafson's law. https://www.researchgate.net/publication/228367369_Reevaluating_Amdahl's_law_and_Gustafson's_law
- Simon H (2014) Why we need Exascale and why we won't get there by 2020. In: *Exascale Radioastronomy Meeting, AASCTS2*. https://www.researchgate.net/publication/261879110_Why_we_need_Exascale_and_why_we_won't_get_there_by_2020
- Singh JP, Hennessy JL, Gupta A (1993) Scaling parallel programs for multiprocessors: methodology and examples. *Computer* 26(7):42–50. <https://doi.org/10.1109/MC.1993.274941>
- Strukov D et al (2019) Building brain-inspired computing. *Nat Commun* 10(12):4838. <https://doi.org/10.1038/s41467-019-12521-x>
- Strukov DB, Snider GS, Stewart DR, Williams RS (2008) The missing memristor found. *Nature* 453(7191):80–83

39. TOP500.org: The top 500 supercomputers. <https://www.top500.org/> (2019)
40. Tsafirir D (2007) The context-switch overhead inflicted by hardware interrupts (and the enigma of do-nothing loops). In: Proceedings of the 2007 workshop on experimental computer science, ExpCS '07. ACM, New York, pp 3–3
41. US Government NSA and DOE: A report from the NSA-DOE technical meeting on high performance computing (2016). https://www.nitrd.gov/nitrdgroups/images/b/b4/NSA_DOE_HPC_Tech_MeetingReport.pdf
42. van Albada SJ, Rowley AG, Senk J, Hopkins M, Schmidt M, Stokes AB, Lester DR, Diesmann M, Furber SB (2018) Performance comparison of the digital neuromorphic hardware SpiN-Naker and the neural network simulation software NEST for a full-scale cortical microcircuit model. *Front Neurosci* 12:291
43. Végh J (2019) How Amdahl's law limits performance of large artificial neural networks. *Brain Inf* 6, 1–11. <https://braininformatics.springeropen.com/articles/10.1186/s40708-019-0097-2/metrics>
44. Végh J (2020) Introducing temporal behavior to computing science. In: 2020 CSCE, fundamentals of computing science, pp FCS2930, in print. IEEE. [arXiv:2006.01128](https://arxiv.org/abs/2006.01128)
45. Végh J (2021). A model for storing and processing information in technological and biological computing systems. In: The 2021 international conference on computational science and computational intelligence; foundations of computer science FCS'21: in print. IEEE, pp FCS4404
46. Végh J (2020) Finally, how many efficiencies the supercomputers have? *J Supercomput* 76(12):9430–9455
47. Végh J, Berki AJ (2020) Do we know the operating principles of our computers better than those of our brain? <https://arxiv.org/abs/2005.05061> <https://american-cse.org/sites/csci2020proc/pdfs/CSCI2020-6ScvvdzjqC7bKupZxFmCoA/762400a668/762400a668.pdf> (in print)
48. Végh J, Molnár P (2017) How to measure perfectness of parallelization in hardware/software systems. In: 18th Internat. Carpathian Control Conf. ICCCC, pp 394–399
49. Végh J, Tisan A (2019). The need for modern computing paradigm: science applied to computing. In: Computational intelligence CSCI The 25th international conference on parallel and distributed processing techniques and applications. IEEE, pp 1523–1532. <https://doi.org/10.1109/CSCI49370.2019.00283>. [arXiv:1908.02651](https://arxiv.org/abs/1908.02651)
50. Végh J, Vászárhelyi J, Drótos D (2019) The performance wall of large parallel computing systems. In: Lecture notes in networks and systems 68. Springer, pp 224–237. https://link.springer.com/chapter/10.1007/978-3-030-12450-2_21
51. von Neumann J (1993) First draft of a report on the EDVAC. *IEEE Ann Hist Comput* 15(4):27–75. <https://doi.org/10.1109/85.238389>
52. von Neumann's missing "Second Draft": what it should contain. In: Proceedings of the 2020 international conference on computational science and computational intelligence (CSCI'20: December 16–18, 2020, Las Vegas. IEEE Computer Society, p CSCI2019 (2020). <https://doi.org/10.1109/CSCI51800.2020.00235>
53. Wang C, Liang SJ, Wang CY, Yang ZZ, Ge Y, Pan C, Shen X, Wei W, Zhao Y, Zhang Z, Cheng B, Zhang C, Miao F (2021) Scalable massively parallel computing using continuous-time data representation in nanoscale crossbar array. *Nat Nanotechnol* <https://doi.org/10.1631/FITEE.1800494>
54. Weaver V, Terpstra D, Moore S (2013) Non-determinism and overcount on modern hardware performance counter implementations. In: Performance analysis of systems and software (ISPASS), 2013 IEEE international symposium on, pp 215–224. <https://doi.org/10.1109/ISPASS.2013.6557172>
55. Williams S, Waterman A, Patterson D (2009) Roofline: an insightful visual performance model for multicore architectures. *Commun ACM* 52(4):65–76
56. www.top500.org: Intel dumps knights hill, future of xeon phi product line uncertain (2017). <https://www.top500.org/news/intel-dumps-knights-hill-future-of-xeon-phi-product-line-uncertain/>
57. Xie S, Sun C, Huang J, Tu Z, Murphy K (2018) Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In: Ferrari V, Hebert M, Sminchisescu C, Weiss Y (eds) Computer vision-ECCV 2018. Springer, Cham, pp 318–335
58. Xu K, Qin M, Sun F, Wang Y, Chen YK, Ren F (2020) Learning in the frequency domain. [arXiv:2002.12416](https://arxiv.org/abs/2002.12416)
59. Zheng F, Li HL, Lv H, Guo F, Xu XH, Xie XH (2015) Cooperative computing techniques for a deeply fused and heterogeneous many-core processor architecture. *J Comput Sci Technol* 30(1):145–162

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.