



Using analog computers in today's largest computational challenges

Sven Köppel¹, Bernd Ulmann¹, Lars Heimann¹, and Dirk Killat^{1,2}

¹Anabrid GmbH, Am Stadtpark 3, 12167 Berlin, Germany

²Microelectronics Department, Brandenburg University of Technology, 03046 Cottbus, Germany

Correspondence: Sven Köppel (koeppel@anabrid.com)

Received: 15 February 2021 – Revised: 30 June 2021 – Accepted: 7 July 2021 – Published: 17 December 2021

Abstract. Analog computers can be revived as a feasible technology platform for low precision, energy efficient and fast computing. We justify this statement by measuring the performance of a modern analog computer and comparing it with that of traditional digital processors. General statements are made about the solution of ordinary and partial differential equations. Computational fluid dynamics are discussed as an example of large scale scientific computing applications. Several models are proposed which demonstrate the benefits of analog and digital-analog hybrid computing.

1 Introduction

Digital computing has transformed many — if not close to all — aspects of industry, humanities and science. Turing completeness allows statements to be made about the computability and decidability of problems and *computational power* of machines. Digital storage has undergone numerous technological advances and is available in increasingly vast amounts. Nevertheless, contemporary digital computing is possibly not the last word in computing, despite its dominance in the consumer market for the last 40+ years.

Fundamental research about non-traditional (also referred to as *unconventional* or *exotic*) computing is taking place in material sciences, chemistry but also in more exotic branches such as biology and life sciences. Amongst others, beyond-Turing computing (Siegelmann, 1995), natural computing (Calude et al., 1999), neuromorphic computing (Schuman et al., 2017; Ziegler, 2020) or quantum computing (Zhou et al., 2020; Georgescu et al., 2014; Kendon et al., 2010) are fields of active investigation. Being fundamental research at heart, these disciplines come with technological challenges. For instance, computing with DNA still requires the use of large scale laboratory equipment and machinery (Deaton

et al., 1998). Currently, not only the low-temperature laboratory conditions but also the necessary error correction schemes challenge practical quantum computers (Wilhelm et al., 2017). This currently negates any practical advantage over silicon based digital computing. Furthermore, all of these alternative (or *exotic*) computer architectures share the characteristic that they are fundamentally *non-portable*. This means they will have to be located at large facilities and dedicated special-purpose computing centers for a long time, if not forever. This is not necessarily a practical drawback, since the internet allows for delocalization of systems.

In contrast to this, silicon based *electronic analog computing* is a technology with a rich history, which operates in a normal workplace environment (non-laboratory conditions; Ulmann, 2020). Digital computers overtook their analog counterparts in the last century, primarily due to their ever-increasing digital clock speeds and their flexibility that comes from their algorithmic approach and the possibility of using these machines in a time-shared environment. However, today Moore's law is coming to a hard stop and processor clock speeds have not significantly increased in the past decade. Manycore architectures and vectorization come with their own share of problems, given their fundamental limits as described, for instance, by Amdahl's law (Rodgers, 1985). GPGPUs and specialized digital computing chips concentrate on vectorized, and even data flow-oriented programming paradigms but are still limited by parasitic capacitances which determine the maximum possible clock frequency and provide a noticeable *energy barrier*.

Thanks to their properties, analog computers have attracted the interest of many research groups. For surveys of theory and applications, see for instance Bournez and Pouly (2021) or the works of MacLennan (2004, 2012, 2019). In this paper, we study the usability of analog computers for applications in science. The fundamental properties of ana-

log computers are low power requirements, low resolution computation and intrinsic parallelism. Two very different uses cases/scenarios can be identified: High performance computing (HPC) and low energy portable computing. The energy and computational demands for both scenarios are diametrically-opposed and this paper is primarily focused on HPC.

The paper is structured as follows: In Sect. 2, we review the general assumptions about digital and analog computing. In Sect. 3, small scale benchmark results are presented for a simple ordinary differential equation. In Sect. 4, a typical partial differential equation is considered as an example for a large scale problem. Spatial discretization effects and computer architecture design choices are discussed. Finally, Sect. 5 summarizes the findings.

2 A Simple (Linear) Model for Comparing Analog and Digital Performance

In this paper, we study different techniques for solving differential equations computationally. Due to the different conventions in algorithmic and analog approaches, a common language had to be found and is described in this section. Here, the term *algorithmic approach* addresses the classical Euler method or classical quasi-linear techniques in ordinary or partial differential equations (ODEs/PDEs), *i.e.*, general methods of numerical mathematics. The term *analog approach* addresses the continuous time integration with an operational amplifier having a capacitor in the feedback loop. The fundamental measures of computer performance under consideration are the time-to-solution T , the power consumption P and the energy demand E .

2.1 Time to solution

The time-to-solution T is the elapsed real time (lab time or wall clock time) for solving a differential equation $\partial_t u = f(u)$ from its initial condition $u(t_0)$ at time t_0 to some target simulation time t_{final} , *i.e.*, for obtaining $u(t_{\text{final}})$. The *speed factor* $k_0 := T/t_{\text{final}}$ is the ratio of elapsed simulation time per wall clock time. On analog computers, this allows to identify the *maximum frequency* $\nu = k_0/(2\pi \text{ s})$. On digital computers, the time-to-solution is used as an estimator (in a statistical sense) for the average k_0 . Relating this quantity to measures in numerical schemes is an important discussion point in this paper. Given the simplest possible ODE,

$$dy/dt = f(y) := y, \quad (1)$$

one can study the analog/digital computer performance in terms of the *complexity* of $f(y)$. For a problem M times as big as the given one, the inherently fully *parallel* analog computer exhibits a constant time-to-solution, *i.e.*, in other terms,

$$T_A^M := T_A^1, \text{ i.e., } T_A^M = T_A^1(M) = \mathcal{O}(1), \quad (2)$$

In contrast, a single core (*i.e.*, nonvectorized, nor superscalar architecture) digital computer operates in a *serial* fashion and can achieve a time-to-solution

$$T_D^M = M \cdot T_D^1, \text{ i.e., } T_D^M = T_D^1(M) = \mathcal{O}(M). \quad (3)$$

Here, T^1 refers to the time-to-solution for solving Eq. (1), while T^M refers to the time-to-solution for solving a problem M times as hard. $M \in \mathbb{N}$ is the measure for the algorithmic complexity of $f(y)$. $f(M) = \mathcal{O}(g(M))$ refers to the Bachmann-Landau asymptotic notation. The number of computational elements required to implement $f(y)$ on an analog computer or the number of instructions required for computing $f(y)$ on a digital computer could provide numbers for M . This is because it is assumed that the evaluation of $f(y)$ can hardly be numerically parallelized. For a system of N coupled ODEs $dy_i/dt = f_i(y_1, \dots, y_N)$, the vector-valued f can be assigned an effective complexity $\mathcal{O}(NM)$ with the same reasoning. However, an overall complexity $\mathcal{O}(M)$ is more realistic since parallelism could be exploited more easily in the direction of N (MIMD, multiple instruction, multiple data).

Furthermore, multi-step schemes implementing higher order numerical time integration can exploit digital parallelization (however, in general the serial time-to-solution of a numerical Euler scheme is the limit for the fastest possible digital time integration). Digital parallelization is always limited by the inherently serial parts of a problem (Amdahl's law, Rodgers, 1985), which makes the evaluation of $f(y)$ the hardest part of the problem. Section 4 discusses complex functions $f(y)$ in the context of the *method of lines* for PDEs.

It should be emphasized that, in the general case, this estimate for the digital computer is a most optimistic (best) estimate, using today's numerical methods. It does not take into account hypothetical algorithmic "shortcuts" which could archive solutions faster than $\mathcal{O}(M)$, because they imply some knowledge about the internal structure of $f(y)$ which could probably also be exploited in analog implementations.

2.2 Power and energy scaling for the linear model

For a given problem with time-to-solution T and average power consumption P , the overall energy is estimated by $E = PT$ regardless of the computer architecture.

In general, an analog computer has to grow with the problem size M . Given constant power requirements per computing element and neglecting increasing resistances or parasitic capacitances, in general one can assume the analog computer power requirement P_A^M for a size M problem to scale from a size 1 problem P_A^1 as $P_A^M = P_A^1 \cdot M$. In contrast, a serial single node digital computer in principle can compute a problem of any size serially by relying on dynamic memory (DRAM), *i.e.*, $P_D^M = P_D^1$. That is, the digital computer power requirements for running a large problem (P_D^M) are (at first approximation) similar to running a small problem P_D^1 . Typically, the DRAM energy demands are one to two orders of magnitude

Table 1. A linear model for work: The computational cost C of evaluating $f(u)$ in the ODE $\partial u/\partial t = f(u)$ is expected to grow as $C \in \mathcal{O}(M)$. The effects on time-to-solution T , power P and energy E demands are shown.

(Quantity)	Digital	Analog
$T(M)$ [s]	$\mathcal{O}(M)$	$\mathcal{O}(1)$
$P(M)$ [W]	$\mathcal{O}(1)$	$\mathcal{O}(M)$
$E(M)$ [J]	$\mathcal{O}(M)$	$\mathcal{O}(M)$

smaller than those of a desktop or server grade processor and are therefore negligible for this estimate.

Interestingly, this model suggests that the overall energy requirements to solve a large problem on an analog and digital computer, respectively, are both E_D^M and $E_A^M = \mathcal{O}(M)$, *i.e.*, the analog-digital energy ratio remains constant despite the fact that the analog computer computes (runs) linearly faster with increasing problem size M . This can be easily deduced by $E = P \cdot T$. In this model, it is furthermore

$$\frac{E_A^M}{E_D^M} = \frac{M P_A^1 T_A^1}{P_D^1 M T_D^1} = \frac{P_A^1 T_A^1}{P_D^1 T_D^1} = \text{const.} \quad (4)$$

The *orthogonal* performance features of the fully-parallel analog computer and the fully-serial digital computer are also summarized in Table 1.

When comparing digital and analog computer power consumption, the power consumption under consideration should include the *total* computer power including administrative parts (like network infrastructure, analog-to-digital converters or cooling) and power supplies. In this work, data of heterogenous sources are compared and definitions may vary.

2.3 Criticism and outlook

Given that the digital and analog technology (electric representation of information, transistor-based computation) is quite similar, the model prediction of a similarly growing energy demand is useful. Differences are of course hidden in the constants (prefactors) of the asymptotic notation $\mathcal{O}(M)$. Quantitative studies in the next sections examine this prefactor in $\mathcal{O}(M)$.

The linear model is already limited in the case of serial digital processors when the computation gets *memory bound* (instead of *CPU-bound*). Having to wait for data leads to a performance drop and might result in a worsened superlinear T_D^M .

Parallel digital computing as well as *serial analog* computing has not yet been subject of the previous discussion. While the first one is a widespread standard technique, the second one refers to analog-digital hybrid computing which, *inter alia*, allows a small analog computer to be used repeatedly on a large problem, effectively rendering the analog part as an analog accelerator or co-processor for the digital part.

Parallel digital computing suffers from a theoretical speedup limited due to the non-parallel parts of the algorithm (see also Gustafson, 1988), which has exponential impact on T_D^M . This is where the intrinsically parallel analog computer exhibits its biggest advantages. Section 4 discusses this aspect of analog computing.

3 A performance survey on solving ordinary differential equations (ODEs)

In this section, quantitative measurements between contemporary analog and digital computers will be made. We use the *Analog Paradigm Model-1* computer (Ulmann, 2019, 2020), a modern modular academic analog computer and an ordinary Intel© Whiskey Lake “ultra-low power mobile” processor (Core i7-8565U) as a representative of a typical desktop-grade processor. Within this experiment, we solve a simple¹ test equation $d^2y/dt^2 = \lambda y$ (with real-valued y and $\lambda = \pm 1$) on both a digital and analog computer.

3.1 Time to solution

The digital computer solved the simple ordinary differential equation (ODE) with simple text-book level scalar benchmark codes written in C and Fortran and compiled with GCC. Explicit (forward) integrator methods are adopted (Euler/Runge-Kutta). The algorithm computed $N = 2 \times 10^3$ timesteps with timestep size $\Delta t = 5 \times 10^{-4}$ each (see also Sect. 4 for a motivation for this time step size). Therefore, it is $t_{\text{final}} = N\Delta t = 1$. No output² was written during the benchmark to ensure the best performance. The time per element update (per integration step) was roughly (45 ± 35) ns. For statistical reasons, the computation was repeated and averaged 10^5 times. Depending on the order of the integration scheme, the overall wall clock time was determined as $T_D = (75 \pm 45)$ μs in order to achieve the simulation time t_{final} .

In contrast, the equation was implemented with integrating (and negating, if $\lambda = -1$) operational amplifiers on the *Analog Paradigm Model-1*. The machine approached $t_{\text{final}} = 1$ in a wall-clock time $T_A = 1 \text{ s}/k_0$ with $k_0 \in \{1, 10, 10^2, 10^3, 10^4\}$ the available integration speed factors on the machine (Ulmann, 2019). The *Analog Paradigm Model-1* reached the solution of $y'' = y$ at $t_{\text{final}} = 1$ in a wall-clock time $T_A = 100$ μs at best.

Note how $T_A/T_D \approx 1$, *i.e.*, in the case of the smallest possible reasonable ODE, the digital computer (2020s energy

¹This equation is inspired by the Dahlquist and Jeltsch (1979) *test equation* $y' = \lambda y$ used for stability studies. The advantage of using an oscillator is the self-similarity of the solution which can be observed over a long time.

²Both in terms of *dense output* or any kind of evolution tracking. A textbook-level approach with minimal memory footprint is adopted which could be considered an *in-place* algorithm.

efficient desktop processor) is roughly as fast as the *Analog Paradigm Model-1* (modern analog computer with an integration level comparable to the 1970s).

Looking forward, given the limited increase in clock frequency, with a faster processor one can probably expect an improvement of T_D down to the order of $1\ \mu\text{s}$. For an analog computer on a chip, one can expect an improvement of T_A down to the order of $1\ \mu\text{s}$ – $10\ \text{ns}$. This renders $T_A/T_D \approx 10^{-(1\pm 1)}$ as a universal constant.

Summing up, with the given numbers above, as soon as the problem complexity grows, the analog computer outperforms the digital one, and this advantage increases linearly.

3.2 Energy and power consumption

The performance measure codes `likwid` (Hager et al., 2010; Röhl et al., 2017; Gruber et al., 2020) and `perf` (de Melo, 2010) were used in order to measure the overall floating-point operations (FLOP) and energy usage of the digital processor. For the Intel mobile processor, this provided a power consumption of $P_D = 10\ \text{W}$ during computing. This number was derived directly from the CPU performance counters. The overall energy requirement was then $E_D = P_D T_D = (0.9 \pm 0.6)\ \text{mJ}$. Note that this number only takes the processor energy demands into account, not any other auxiliary parts of the overall digital computer (such as memory, main board or power supply). For the overall power consumption, an increase of at least 50% is expected.

The analog computer energy consumption is estimated as $P_A \approx 400\ \text{mW}$. The number is based on measurements of actual *Analog Paradigm Model-1* computing units, in particular 84 mW for a single summer and 162 mW for a single integrator. The overall energy requirement is then $E_A = P_A T_A = 40\ \mu\text{J}$.

Note that $P_D/P_A \approx 25$, while $E_D/E_A \approx (2.25 \pm 1.5)$. The conclusion is that the analog and digital computer require a similar amount of energy for the given computation, a remarkable result given the 40-year technology gap between the two architectures compared here.

For power consumption, it is hard to give a useful projection due to the accumulating administrative overhead in case of parallel digital computing, such as data transfers, non-uniform memory accesses (NUMA) and switching networking infrastructure. It can be assumed that this will change the ratio E_D/E_A further in favor of the analog computer for both larger digital and analog computers. Furthermore, higher integration levels lower E_A : the *Analog Paradigm Model-1* analog computer is realized with an integration level comparable with 1970s digital computers. We can reasonably expect a drop of two to three orders of magnitude in power requirements with fully integrated analog computers.

3.3 Measuring computational power: FLOP per Joule

For the digital computer, the number of computed *floating-point operations* (FLOP³) can be measured. The overall single core nonvectorized performance was measured as $F \approx 1\ \text{GFLOP/s}$. A single computation until t_{final} required roughly $F_D = 3\ \text{kFLOP}$. The ratio $F_D/P_D = 100\ \text{MFLOP/J}$ is a measure of the number of computations per energy unit on this machine. This performance was one to two orders less than typical HPC numbers. This is because an energy-saving desktop CPU and not a high-end processor was benchmarked. Furthermore, this benchmark was by purpose single-threaded.

In this non-vectorized benchmark, the reduced resolution of the analog computer was ignored. In fact it is slightly lower than an IEEE 754 half precision floating-point, compared to the double precision floating-point numbers in the digital benchmark. One can then assign the analog computer a *time-equivalent* floating-point operation performance

$$F_A := F_D \frac{T_A}{T_D} \approx 10^{(1\pm 1)} F_D = 3 \times 10^{(4\pm 1)} \text{FLOP}. \quad (5)$$

The analog computer FLOP-per-Joule ratio (note that $\text{FLOP/J} = \text{FLOPs/W}$) is

$$\frac{F_A}{E_A} = \frac{3 \times 10^{(4\pm 1)} \text{FLOP}}{40\ \mu\text{J}} = 7.5 \times 10^{8\pm 1} \text{FLOP/J}. \quad (6)$$

That is, the analog computer's "FLOP per Joule" is slightly larger than for the digital one. Furthermore, one can expect an increase of F_A/E_A by 10–100 for an analog computer chip. See for instance Cowan (2005) and Cowan et al. (2005), who claim 20 GFlop/s. We expect 300 GFlop/s to be more realistic, though (Table 2).

Keep in mind that the FLOP/s or FLOP/J measures are (even in the case of comparing two digital computers) always problem/algorithm-specific (*i.e.*, in this case a Runge Kutta solver of $y'' = y$) and therefore controversial as a comparative figure.

4 PDEs and many degrees of freedom

This section presents forecasts about the solution of large scale differential equations. No benchmarks have been carried out, because a suitable integrated analog computer on chip does not yet exist. For the estimates, an analog computer on chip with an average energy consumption of about $P_N = 4\ \text{mW}$ per computing element (*i.e.*, per integration, multiplication, etc.) and maximum frequency $\nu = 100\ \text{Mhz}$, which is

³sic! We either argue with *overall* FLOP and Energy (Joule) or *per second* quantities such as FLOP/s (in short FLOPS) and Power (Watt). In order to avoid confusion, we avoid the abbreviation "FLOPS" in the main text. Furthermore, SI prefixes are used, *i.e.*, $\text{kFLOP} = 10^3\ \text{FLOP}$, $\text{MFLOP} = 10^6\ \text{FLOP}$ and $\text{GFLOP} = 10^9\ \text{FLOP}$.

Table 2. Small scaling summary: Measured benchmark (Intel© processor vs. *Analog Paradigm Model-1*) and expected/projected analog chip results.

	Measured		Projected
	Digital	Analog (<i>MI</i>)	Analog Chip
T [μs]	75 ± 45	100	$10^{-(0.5 \pm 0.5)}$
$k_0 \sim 1/\Delta t$ [Hz]	3×10^4	10^4	$10^{6.5 \pm 0.5}$
P [W]	10	0.4	10^{-2}
$E = P \cdot T$ [μJ]	900 ± 600	40	$10^{-(2.5 \pm 0.5)}$
F [FLOP/s]	10^9	$3 \times 10^{(4 \pm 1)}$	7×10^5
F/E [FLOP/J]	10^8	$7.5 \times 10^{8 \pm 1}$	3×10^{11}

referred to as the *analog* maximum frequency ν^A in the following, was assumed. These numbers are several orders of magnitude better than the $P_N = 160$ mW and $\nu = 100$ kHz of the *Analog Paradigm Model-1* computer discussed in the previous section. For the digital part, different systems than before are considered.

In general, the bandwidth of an analog computer depends on the frequency response characteristics of the elements, such as summers and integrators. The actual achievable performance also depends on the technology. A number of examples shall be given to motivate our numbers: In 65 nm CMOS technology, bandwidths of over 2 GHz are achievable with integrators (Breems et al., 2016). At unity-gain frequencies of 800 MHz to 1.2 GHz and power consumption of less than 2 mW, integrators with a unity-gain frequency of 400 MHz are achievable (Wang et al., 2018).

4.1 Solving PDEs on digital and analog computers

Partial differential equations (PDEs) are among the most important and powerful mathematical frameworks for describing dynamical systems in science and engineering. PDE solutions are usually fields $\mathbf{u} = \mathbf{u}(\mathbf{r}, t)$, i.e., functions⁵ of spatial position \mathbf{r} and time t . In the following, we concentrate on initial value boundary problems (IVBP). These problems are described by a set of PDEs valid within a spatial and temporal domain and complemented with field values imposed on the domain boundary. For a review of PDEs, their applications and solutions see for instance Brezis and Browder (1998). In this text, we use computational fluid dynamics (CFD) as a representative theory for discussing general PDE performance. In particular, classical hydrodynamics (Euler equation) in a flux-conservative formulation is described by hyperbolic conservation laws in the next sections. Such PDEs

⁴Summation will be done implicitly on chip by making use of Kirchhoff’s law (current summing) so that no explicit computing element are required for this operation.

⁵The explicit dependency on \mathbf{r} and t is omitted in the following text.

have a long tradition of being solved with highly accurate numerical schemes.

Many methods exist for the spatial discretization. While finite volume schemes are popular for their conservative properties, finite difference schemes are in general cheaper to implement. In this work, we stick to simple finite differences on a uniform grid with some uniform grid spacing $\Delta \mathbf{r}$. The evolution vector field $\mathbf{u}(\mathbf{r}, t)$ is sampled on G grid points per dimension and thus replaced by $\mathbf{u}_k(t)$ with $0 \leq k < G$. It is worthwhile to mention that this approach works in classical orthogonal “dimension by dimension” fashion, and the number of total grid points is given by G^D . The computational domain is thus bound by $\Omega = [\mathbf{r}_0, \mathbf{r}_G]^D$. A spatial derivative $\partial_i f$ is then approximated by a central finite difference scheme, for instance $\partial_i f_k \approx (f_{k+1} - f_{k-1})/(2\Delta x) + \mathcal{O}(\Delta x^2)$ for a second order accurate central finite difference approximation of the derivative of some function f at grid point k .

Many algorithmic solvers implement numerical schemes which exploit the *vertical method of lines* (MoL) to rewrite the PDE into coupled ordinary differential equations (ODEs). Once applied, the ODE system can be written as $\partial_t \mathbf{u}^k = G^k(\mathbf{u}, \nabla \mathbf{u})$ with \mathbf{u}^k denoting the time evolved (spatial) degrees of freedom and G^k functions containing spatial derivatives ($\partial_i u^j$) and algebraic sources. A standard time stepping method determines a solution $\mathbf{u}(t_1)$ at later time $t_1 > t_0$ by basically integrating $\mathbf{u}^k(t_1) = \int_{t_0}^{t_1} G^k(\mathbf{u}(t)) dt + \mathbf{u}^k(t_0)$. Depending on the details of the scheme, G^k is evaluated (probably repeatedly or in a weak-form integral approach) during the time integration of the system. However, note that other integration techniques exist, such as the arbitrary high order ADER technique (Titarev and Toro, 2002, 2005). The particular spatial discretization method has a big impact on the computational cost of G^i . Here, we focus on the (simplest) finite difference technique, where the number of neighbor communications per dimension grows linearly with the *convergence order* of the scheme.

4.2 Classical Hydrodynamics on analog computers

The broad class of *fluid dynamics* will be discussed as popular yet simple type of PDEs. It is well known for its efficient description of the flow of liquids and gases in motion and is applicable in many domains such as aerodynamics, in life sciences as well as fundamental sciences (Sod, 1985; Chu, 1979; Wang et al., 2019). In this text, the simplest formulation is investigated: the Newtonian hydrodynamics (also referred to as *Euler equations*) with an ideal gas equation of state. It is given by a nonlinear PDE describing the time evolution of a mass density ρ , its velocity v^i , momentum $p^i = \rho v^i$ and energy $e = t + \varepsilon$, with the kinetic contribution $t = \rho v^2/2$ and an “internal” energy ε , which can account for forces on smaller length scales than the averaged scale.

Flux conservative *Newtonian hydrodynamics* with an ideal gas equation of state are one of the most elementary and text-book level formulations of fluid dynamics (Toro, 1998;

Harten, 1997; Hirsch, 1990). The PDE system can be written in a dimension agnostic way in D spatial dimensions (*i.e.*, independent of the particular choice for D) as

$$\frac{\partial \mathbf{u}}{\partial t} - \nabla \cdot \mathbf{f} = \mathbf{S} \text{ with } \nabla \cdot \mathbf{f} = \sum_i^{n_d} \frac{\partial f^i}{\partial x^i}, \quad (7)$$

$$\mathbf{u} = \begin{pmatrix} \rho \\ p^j \\ e \end{pmatrix}, \mathbf{f}^i = \mathbf{f}^i(\mathbf{u}, \mathbf{v}) = \begin{pmatrix} p^i \\ p^i v^j - p \delta^{ij} \\ v^i (e + p) \end{pmatrix}, \quad (8)$$

with $i, j \in [1, D]$. Here, the pressure $p = \rho \varepsilon(\Gamma - 1)$ defines the ideal gas equation of state, with adiabatic index $\Gamma = 2$ and δ^{ij} is the Kronecker delta. A number of vectors are important in the following: The integrated state or *evolved* vector \mathbf{u} in contrast to the primitive state vector or *auxiliary* quantities $\mathbf{v}(\mathbf{u}) = (p, v^i)$, which is a collection of so called *locally reconstructed* quantities. Furthermore, the right hand sides in Eq. (7) do not explicitly depend on the spatial derivative $\partial^i \rho$, thus the conserved flux vector $\mathbf{f} = \mathbf{f}(\nabla \mathbf{q}, \mathbf{v})$ is only a function of the derivatives of the *communicated* quantities $\mathbf{q} = (e, p^i)$ and the auxiliaries \mathbf{v} . Furthermore, \mathbf{q} and \mathbf{v} are both functions of \mathbf{u} only.

$\mathbf{S} = 0$ is a source term. Some hydrodynamical models can be coupled by purely choosing some nonzero \mathbf{S} , such as the popular *Navier Stokes* equations which describe viscous fluids. Compressible Navier Stokes equations can be written with a source term $\mathbf{S} = \nabla \cdot \mathbf{F}^v$, with

$$\text{diffusion fluxes } \mathbf{F}^v = (0, \tau^{ij}, \sum_k \tau^{ik} v^k - q^j)^T, \quad (9)$$

$$\text{viscous stress } \tau^{ij} = \mu(\partial^i v^j + \partial^j v^i - \frac{2}{3}(\partial^k v^k)\delta^{ij}), \quad (10)$$

$$\text{and heat flux } q^i = -(c_p \mu / \text{Pr}) \partial^i T, \quad (11)$$

with specific heats c_p, c_v , viscosity coefficient μ , Prandtl number Pr and temperature T determined by the perfect gas equation of state, *i.e.*, $T = (e - v^2)/(2c_v)$. The computational cost from Euler equation to Navier Stokes equation is roughly *doubled*. Furthermore, the partial derivatives on the velocities and temperatures also *double* the quantities which must be communicated with each neighbor in every dimension. We use Euler equations in the following section for the sake of simplicity.

4.3 Spatial discretization: Trading interconnections vs. computing elements

Schemes of (convergence) order F shall be investigated, which require the communication with F neighbour elements. For instance, a $F = 4$ th order accurate stencil has to communicate and/or compute four neighbouring elements $\mathbf{f}_{k-2}, \mathbf{f}_{k-1}, \mathbf{f}_{k+1}, \mathbf{f}_{k+2}$. Typically, long-term evolutions are carried out with $F = 4$ or $F = 6$. In the following, for simplicity, second order stencil ($F = 2$) is chosen. One identifies three different subcircuits

$$\mathbf{u}_k(\mathbf{f}_{k+1}, \mathbf{f}_{k-1}) := \int (\mathbf{f}_{k+1} - \mathbf{f}_{k-1}) dt / (2\Delta x), \quad (12)$$

with $\mathbf{f}_{k\pm 1} := \mathbf{f}_k(\mathbf{q}_{k\pm 1}, \mathbf{v}_k)$ and $\mathbf{v}_k := \mathbf{v}_k(\mathbf{u}_k)$ according to their previous respective definitions. Figure 1 shows this “building block” for a single grid point, an exemplar for up to $D = 2$ dimensions with an $F = 2$ nd order finite difference stencil. The circuit identifies a number of intermediate expressions which are labeled as these equations:

$$\begin{aligned} & \partial_t \begin{pmatrix} \rho_{i,k} \\ p_{i,k}^x \\ p_{i,k}^y \\ e_{i,k} \end{pmatrix} \\ &= \frac{\overbrace{\begin{pmatrix} p_{i+1,k}^x \\ p_{i+1,k}^x v_{i+1,k}^x - p_{i+1,k} \\ p_{i+1,k}^y v_{i+1,k}^y \\ v_{i+1,k}^x (e_{i+1,k} + p_{i+1,k}) \end{pmatrix}}^{f_{i+1,k}^x} - \overbrace{\begin{pmatrix} p_{i-1,k}^x \\ p_{i-1,k}^x v_{i-1,k}^x - p_{i-1,k} \\ p_{i-1,k}^y v_{i-1,k}^y \\ v_{i-1,k}^x (e_{i-1,k} + p_{i-1,k}) \end{pmatrix}}^{f_{i-1,k}^x}}{2\Delta x} \\ &+ \frac{\overbrace{\begin{pmatrix} p_{i,k+1}^y \\ p_{i,k+1}^y v_{i,k+1}^x \\ p_{i,k+1}^y v_{i,k+1}^y - p_{i,k+1} \\ v_{i,k+1}^y (e_{i,k+1} + p_{i,k+1}) \end{pmatrix}}^{f_{i,k+1}^y}}{2\Delta y} - \frac{\overbrace{\begin{pmatrix} p_{i,k-1}^y \\ p_{i,k-1}^y v_{i,k-1}^x \\ p_{i,k-1}^y v_{i,k-1}^y - p_{i,k-1} \\ v_{i,k-1}^y (e_{i,k-1} + p_{i,k-1}) \end{pmatrix}}^{f_{i,k-1}^y}}{2\Delta y} \\ &= \frac{\begin{pmatrix} \Delta \rho_{i+1,k} \\ \Delta p_{i+1,k}^x \\ \Delta p_{i+1,k}^y \\ \Delta e_{i+1,k} \end{pmatrix} - \begin{pmatrix} \Delta \rho_{i-1,k} \\ \Delta p_{i-1,k}^x \\ \Delta p_{i-1,k}^y \\ \Delta e_{i-1,k} \end{pmatrix}}{2\Delta x} + \frac{\begin{pmatrix} \Delta \rho_{i,k+1} \\ \Delta p_{i,k+1}^x \\ \Delta p_{i,k+1}^y \\ \Delta e_{i,k+1} \end{pmatrix} - \begin{pmatrix} \Delta \rho_{i,k-1} \\ \Delta p_{i,k-1}^x \\ \Delta p_{i,k-1}^y \\ \Delta e_{i,k-1} \end{pmatrix}}{2\Delta y} \end{aligned} \quad (13)$$

Just like in Fig. 1, all expressions which are vanishing in a single spatial dimension are colored in red. Furthermore, note how the index i denotes the x -direction and k the y -direction, and that there are different fluxes \mathbf{f}^j in the particular directions. Equation (13) is closed with the element-local auxiliary recovery

$$\begin{pmatrix} v_{i,k}^x \\ v_{i,k}^y \\ p_{i,k} \end{pmatrix} = \begin{pmatrix} p_{i,k}^x / \rho_{i,k} \\ p_{i,k}^y / \rho_{i,k} \\ e_{i,k} - \rho_{i,k} \left((v_{i,k}^x)^2 + (v_{i,k}^y)^2 \right) / 2 \end{pmatrix}. \quad (14)$$

Note that one can trade neighbor communication (*i.e.*, number of wires between grid points) for local recomputation. For instance, it would be mathematically clean to communicate only the conservation quantities \mathbf{u} and reconstruct \mathbf{v} whenever needed. In order to avoid too many recomputations, some numerical codes also communicate parts of \mathbf{v} . In an analog circuit, it is even possible to communicate parts of the finite differences, such as the $\Delta v_{i,k}$ quantities in Eq. (13).

The number of analog computing elements required to solve the Euler equation on a single grid point is determined

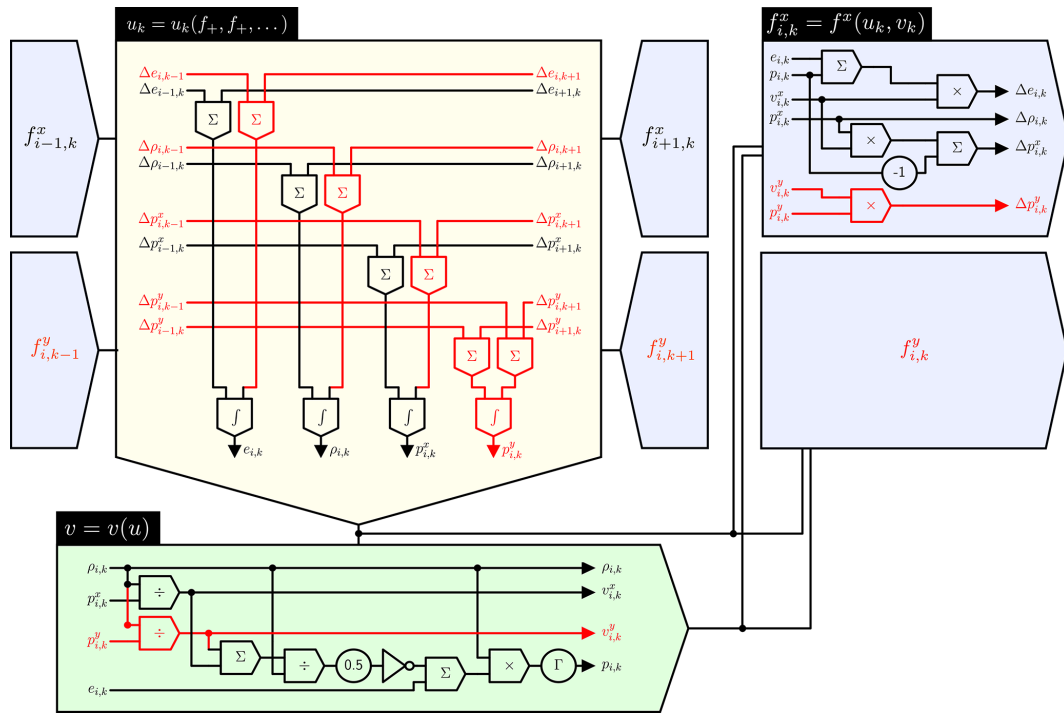


Figure 1. Overview circuit showing the blocks f , u and v . The three labeled blocks are distinguished by colour. Information flow is indicated with arrows. The overall circuit is given for lowest order (RK1) and in one spatial dimension. The red circuitry is the required addition for two spatial dimensions. All computing elements are drawn “abstractly” and could be directly implemented with (negating) operational amplifiers on a very large *Analog Paradigm Model-1* analog computer.

as $N_{\text{single}} = 5D + 5F(D + 2) + 9$, with D being the number of spatial dimensions and F the convergence order (*i.e.*, basically the finite difference stencil size). Typical choices of interest are convergence orders of $F \in [2, 6]$ in $D \in [1, 3]$ spatial dimensions. Inserting the averaged $F = 3 \pm 1$ and $D = 2 \pm 1$ into N_{single} yields an averaged $N_{\text{single}} \approx (84 \pm 40)$ computing elements per spatial degree of freedom (grid point) required for implementing Euler equations.

Unfortunately, this circuit is too big to fit on the *Analog Paradigm Model-1* computer resources available. Consequently the following discussion is based on a future implementation using a large number of interconnected analog chips. It is noteworthy that this level of integration is necessary to implement large scale analog computing applications. With $P_N = 4 \text{ mW}$ per computing element, the average power per spatial degree of freedom (*i.e.*, single grid point) is $P_{ND} = (336 \pm 160) \text{ mW}$.

4.4 Time to solution

Numerical PDE solvers are typically benchmarked using a *wall-clock time per degree of freedom update* measure T_{DOF} , where *element update* typically means a time integration timestep. In this measure, the overall wall clock time is normalized (divided) by the number of spatial degrees of freedom as well as the number of parallel processors involved.

The fastest digital integrators found in literature carry out a time per degree of freedom update $T_{\text{DOF}} = 10^{1 \pm 1} \mu\text{s}$. Values smaller than $1 \mu\text{s}$ require already the use of sophisticated communication avoiding numerical schemes such as discontinuous Galerkin (DG) schemes.⁶ For instance, Dumbser et al. (2008) demonstrate the superiority of so called $P_N P_M$ methods (polynomial of degree N for reconstruction and M for time integration, where the limit $P_0 P_M$ denotes a standard high-order finite volume scheme) by reporting $T_{\text{DOF}} = 0.8 \mu\text{s}$ for a $P_2 P_2$ method when solving two-dimensional Euler equations. Diot et al. (2013) report an adaptive scheme which performs no faster than $T_{\text{EU}} = 30 \mu\text{s}$ when applied to three-dimensional Euler equations. The predictor-corrector arbitrary-order ADER scheme applied by Köppel (2018) and Fambri et al. (2018) to the *general-relativistic magnetodynamic* extension of hydrodynamics reported $T_{\text{DOF}} = 41 \mu\text{s}$ as the fastest speed obtained. The non-parallelizable evaluation of more complex hydrodynamic models is clearly reflected in the increasing times T_{DOF} .

Recalling the benchmark result of $T_{\text{DOF}} \sim 45 \text{ ns}$ from Sect. 3.1, the factor of 1000 is mainly caused by the inevitable communication required for obtaining neighbor val-

⁶ $h - p$ methods, which provide both mesh refinement in grid spacing h as well as a “local” high order description typically in some function base expansion of order p . For reviews, see for instance Cockburn and Shu (2001) or Shu (2016).

ues when solving $f(y, \nabla y)$ in $\partial_t y = f(y)$. Switched networks have an intrinsic communication latency and one cannot expect T_{DOF} to shrink significantly, even for newer generations of supercomputers. A key advantage of analog computing is that grid neighbor communication happens continuously in the same time as in the grid-local circuit. That is, no time is lost for communication.

One can do a comparison with the analog computer without knowing the simulation time step size Δt . The reasoning is based on the maximum frequency, *i.e.*, the shortest wavelength which can be resolved with a (first order in time⁷) numerical scheme is $f_{\text{sim}} := 1/(10\Delta t)$, *c.f.*, Fig. 2. The factor $10 = 2 \cdot 5$ includes a factor of 2 due to the Nyquist-Shannon sampling theorem, while the factor of 5 is chosen to take into account that a numerical scheme can marginally reconstruct a wave at frequency $f = 1/(2\Delta t)$ by two points while it can be obtained perfectly by the analog computer (down to machine precision without any artifacts). The integration of signals beyond the maximum frequency results in a nonlinear response which heavily depends on the electrical details of the circuit (which are beyond the scope of the analog computer architecture discussed in this paper). One can demand that the numerical integrator time resolution is good enough to reconstruct a signal *without* prior knowledge on the wave form even at the maximum frequency.⁸ This drives the demand for 5 additional sampling points per half-wave, in order to make analog and digital outcome comparable (see also Fig. 2).

It is noted that this argument is relevant as long as one is interested in obtaining and preserving the *correct* time evolution (of a system described by the differential equation) with an analog or digital computer, respectively. In general, it is not valid to reduce the computational correctness within the solution domain of an initial value problem as this will invalidate any later solution.

By assigning the numerical PDE solver a maximum frequency identical to the highest frequency which can be evolved by the scheme in a given time, one introduces an *effective digital computer maximum frequency*

$$\nu^{\text{D}} = 1/(10T_{\text{DOF}}) = 10^{1\pm 1} \text{ kHz}. \quad (15)$$

Note how the mapping of simulation time (interval) Δt to wall-clock time (interval) T_{DOF} results in a mapping of simulation frequency f_{sim} to wall-clock (or real-time) frequency ν^{D} (Fig. 2).

⁷For a high order time integration scheme, the cutoff increases formally linearly as $f_0 \sim p/(10T_{\text{DOF}})$. That is, for a fourth order scheme, the digital computer is effectively four times faster in this comparison.

⁸Note that on a digital computer, the maximum frequency is identical to a *cutoff frequency* (also referred to as *ultraviolet cutoff*). On analog computers, there is no such *hard* cutoff as computing elements tend to be able to compute with decreased quality at higher frequencies.

The calculated $\nu^{\text{D}} = 10^{-2\pm 1}$ MHz has to be contrasted with $\nu^{\text{A}} = 100$ MHz of the analog computer chip. One can conclude that analog computers can solve large scale high performance computing at least $\nu^{\text{A}}/\nu^{\text{D}} = 10^{3\pm 1}$ times faster than the digital ones, when T_{A} and T_{D} are the analog and digital time to solution. Since $T \sim 1/\nu$, the resolution time reduces accordingly and $T_{\text{A}}/T_{\text{D}} = 10^{-3\pm 1}$.

This is a remarkable result as it already assumes the fastest numerical integration schemes on a *perfectly* scaling parallel digital computer. In practical problems, these assumptions are hardly ever met: The impossibility of (ideal) parallelization is one of the major drawbacks of digital computing. Nevertheless, the above results show that even without these drawbacks, the analog computer is orders of magnitude faster. Notably, while it needs careful adjustment both the problem and the code for a high-performance computer to achieve acceptable parallel performance, when using an analog computer these advantages come effortless. The only way to reduce the speed or timing advantage is to choose a disadvantageous or unsuitable number scaling.

In this study the low resolution of an analog computer (which is effectively IEEE 754 half precision floating-point) has been neglected. In fact, high order time integration schemes can invest computing time in order to achieve machine level accuracy which a typical error $\Delta f_{\text{digital}} \sim 10^{-10}$ on some evolved function or field f and an error definition $\Delta f_{\text{simulation}} := (f_{\text{simulation}} - f_{\text{exact}})/f_{\text{exact}}$. An analog computer is limited by its intrinsic accuracy with a typical error $\Delta f_{\text{analog}} \sim 10^{-(4\pm 1)}$ (averaging over the *Analog Paradigm Model-1* and future analog computers on chip).

4.5 Energy and power consumption

One expects the enormous speedup $T_{\text{A}}/T_{\text{D}}$ of the analog computer to result in a much lower energy budget $E_{\text{D}} = (T_{\text{D}}/T_{\text{A}})E_{\text{A}} = 10^{3\pm 1}E_{\text{A}}$ for a given problem. However, as the power requirement is proportional to the analog computer size, $P_{\text{A}} = NP_{\text{ND}}$, the problem size (number of grid points) which can be handled by the analog computer is limited by the overall power consumption. For instance, with a typical high performance computer power consumption of $P_{\text{A}} = 20$ MW, one can simultaneously evolve a grid with $N = P_{\text{A}}/P_{\text{ND}} = 10^{11\pm 0.5}$ points. This is in the same order of magnitude as the largest scale computational fluid dynamics simulations evolved on digital high performance computer clusters (*c.f.*, Green 500 list, Subramaniam et al., 2013, 2020). Note that in such a setup, the solution is obtained on average $10^{3\pm 1}$ times faster with a purely analog computer and consequently also the energy demand is $10^{3\pm 1}$ times lower.

Just to depict an analog computer of this size: Given 1000 computing elements per chip, 1000 chips per rack unit, 40 units per rack still requires 2500 racks to build such a computer in a traditional design. This is one order of magnitude larger than the size of typical high performance cen-

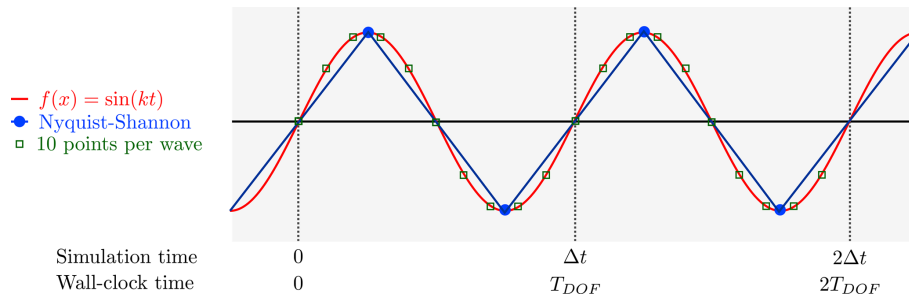


Figure 2. Analog signal sampling vs. numerical time integration: The time evolved sine with *maximum frequency* $k = 2\pi/\lambda$ has just the wavelength $\lambda = \Delta t$, with Δt being the timestep size of the explicit Euler scheme. The Nyquist-Shannon theorem allows to determine wave length and phase position with two sampling points per wave length. However, a first order reconstruction of numerical data shows a triangle (zigzag) function. In contrast, the full wave is clearly visible at analog integration. More sampling points close the gap between analog and numerical representation.

ters. Clearly, at such a size the interconnections will also have a considerable power consumption, even if the monumental engineering challenges for such a large scale interconnections can be met. On a logical level, interconnections are mostly wires and switches (which require little power, compared to computing elements). This can change dramatically with level converters and an energy estimate is beyond the scope of this work.

4.6 Hybrid techniques for trading power vs. time

The analog computers envisaged so far have to grow with problem size (*i.e.*, with grid size, but also with equation complexity). Modern chip technology could make it theoretically possible to build a computer with 10^{12} analog computing elements, which is many orders of magnitude larger than any analog computer that has been built so far (about 10^3 computing elements at maximum). The idea of combining an analog and a digital computer thus forming a *hybrid computer* featuring analog and digital computing elements is not new. With the digital memory and algorithmically controlled program flow, a small analog computer can be used repeatedly on a larger problem under control of the digital computer it is mated to. Many attempts at solving PDEs on hybrid computers utilized the analog computer for computing the element-local updated state with the digital computer looping over the spatial degrees of freedom. In such a scheme, the analog computer fulfils the role of an *accelerator* or *coprocessor*. Such attempts are subject of various historical (such as Nomura and Deiters, 1968; Reihing, 1959; Vichnevetsky, 1968, 1971; Volynskii and Bukham, 1965; Bishop and Green, 1970; Karplus and Russell, 1971; Feilmeier, 1974) and contemporary studies (for instance Amant et al., 2014; Huang et al., 2017).

A simple back-of-the-envelope estimation with a modern hybrid computer tackling the $N = 10^{11}$ problem is described below. The aim is to trade the sheer number of computing elements with their electrical power P , respectively,

against solution time T . It is assumed that the analog-digital hybrid scheme works similarly to numerical parallelization: The simulation domain with N degrees of freedom is divided into Q parts which can be evolved independently to a certain degree (for instance in a predictor-corrector scheme). This allows to use a smaller analog computer which only needs to evolve N/Q degrees of freedom at a time. While the power consumption of such a computer is reduced to $P_A \rightarrow P_A/Q$, the time to solution increases to $T_A \rightarrow QT_A$. Of course, the overall required energy remains the same, $E_A = P_A T_A = (P_A/Q)(QT_A)$.

In this simple model, energy consumption of the digital part in the hybrid computer as well as numerical details of the analog-digital hybrid computer scheme have been neglected. This includes the time-to-solution overhead introduced by the numerical scheme implemented by the digital computer (negligible for reasonably small Q) and the power demands of the ADC/DAC (analog-to-digital/digital-to-analog) converters (an overhead which scales with $(D+2)G^D/Q$, *i.e.*, the state vector size per grid element).

Given a fixed four orders of magnitude speed difference $v^D/v^A = 10^4$ and a given physical problem with grid size $N = 10^{11}$, one can build an analog-digital hybrid computer which requires less power and is reasonably small so that the overall computation is basically still done in the analog domain and digital effects will not dominate. For instance, with Q chosen just as big as $Q = v^D/v^A$, the analog computer would evolve only $N/Q = 10^7$ points in time, but run 10^4 times “in repetition”. The required power reduces from cluster-grade to desktop-grade $P_A = (N/Q)P_{ND} = 3.3$ kW. The runtime advantage is of course lost, $T_D/T_A = (Qv^A)/v^D = 1$.

Naturally, this scenario can also be applied to solve larger problems with a given grid size. For instance, given an analog computer with the size of $N = 10^{11}$ grid points, one can solve a grid of size QN by successively evolving Q parts of the computer with the same power P_A as for a grid of size N . Of course, the overall time to solution and energy will grow

with Q . In any case, time and energy remain (3 ± 1) orders of magnitude lower than for a purely digital computer solution.

5 Summary and outlook

In Sect. 2, we have shown the time and power needs of analog computers are *orthogonal* to those of digital computers. In Sect. 3, we performed an actual miniature benchmark of a commercially available *Analog Paradigm Model-1* computer versus a mobile Intel© processor. The results are remarkable in several ways: The modern analog computer *Analog Paradigm Model-1*, uses integrated circuit technology which is comparable to the 1970s digital integration level. Nevertheless it achieves competitive results in computational power and energy consumption compared to a mature cutting-edge digital processor architecture which has been developed by one of the largest companies in the world. We also computed a problem-dependent effective FLOP/s value for the analog computer. For the key performance measure for energy-efficient computing, namely *FLOP-per-Joule*, the analog computer again obtains remarkable results.

Note that while FLOP/s is a popular measure in scientific computing, it is always application- and algorithm-specific. Other measures exist, such as transversed edges per second (TEPS) or synaptic updates per second (SUPS). Cockburn and Shu (2001) propose for instance to measure the *efficiency* of a PDE solving method by computing the inverse of the product of the (spatial-volume integrated) L^1 -error times the computational cost in terms of time-to-solution or invested resources.

In Sect. 4, large scale applications were discussed on the example of fluid dynamics and by comparing high performance computing results with a prospected analog computer chip architecture. Large scale analog applications can become power-bound and thus require the adoption of analog-digital hybrid architectures. Nevertheless, with their $\mathcal{O}(1)$ runtime scaling, analog computers excel for time integrating large coupled systems where algorithmic approaches suffer from communication costs. We predict outstanding advantages in terms of time-to-solution when it comes to large scale analog computing. Given the advent of chip-level analog computing, a *gigascale* analog computer (a device with $\sim 10^9$ computing elements) could become a *game changer* in this decade. Of course, major obstacles have to be addressed to realize such a computer, such as the interconnection topology and realization in an (energy) efficient manner.

Furthermore, there are a number of different approaches in the field of partial differential equations which might be even better suited to analog computing. For instance, solving PDEs with artificial intelligence has become a fruitful research field in the last decade (see for instance Michoski et al., 2020; Schenck and Fox, 2018), and analog neural networks might be an interesting candidate to challenge digital approaches. Number representation on analog computers

can be nontrivial when the dynamical range is large. This is frequently the case with fluid dynamics, where large density fluctuations are one reason why perturbative solutions fail and numerical simulations are carried out in the first place. One reason why *indirect* alternative approaches such as neural networks could be better suited than *direct* analog computing networks is that this problem is avoided. Furthermore, the demand for high accuracy in fluid dynamics can not easily fulfilled by low resolution analog computing. In the end, it is quite possible that a small-sized analog neural network might outperform a large-sized classical pseudo-linear time evolution in terms of time-to-solution and energy requirements. Most of these engineering challenges have not been discussed in this work and are subject to future studies.

Code availability. The software code is available in the Supplement.

Data availability. No data sets were used in this article.

Supplement. The supplement related to this article is available online at: <https://doi.org/10.5194/ars-19-105-2021-supplement>.

Author contributions. BU performed the analog simulations. SK carried out the numerical simulations and the estimates. DK implements the concept for an integrated analog computer. All authors contributed to the article.

Competing interests. The authors declare that they have no conflict of interest.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Special issue statement. This article is part of the special issue "Kleinheubacher Berichte 2020".

Acknowledgements. We thank our anonymous referees for helpful comments and corrections. We further thank Chris Giles for many corrections and suggestions which improved the text considerably.

Review statement. This paper was edited by Madhu Chandra and reviewed by two anonymous referees.

References

- Amant, R., Yazdanbakhsh, A., Park, J., Thwaites, B., Esmaeilzadeh, H., Hassibi, A., Ceze, L., and Burger, D.: General-purpose code acceleration with limited-precision analog computation, in: ACM/IEEE 41st International Symposium on Computer Architecture (ISCA), vol. 42, pp. 505–516, <https://doi.org/10.1109/ISCA.2014.6853213>, 2014.
- Bishop, K. and Green, D.: Hybrid Computer Implementation of the Alternating Direction Implicit Procedure for the Solution of Two-Dimensional, Parabolic, Partial-Differential Equations, *AICHE Journal*, 16, 139–143, <https://doi.org/10.1002/aic.690160126>, 1970.
- Bournez, O. and Pouly, A.: A Survey on Analog Models of Computation, in: *Handbook of Computability and Complexity*, Springer, Cham, pp. 173–226, https://doi.org/10.1007/978-3-030-59234-9_6, 2021.
- Breems, L., Bolatkale, M., Brekelmans, H., Bajoria, S., Niehof, J., Rutten, R., Oude-Essink, B., Fritschij, F., Singh, J., and Lassche, G.: A 2.2 GHz Continuous-Time Delta Sigma ADC With –102 dBc THD and 25 MHz Bandwidth, *IEEE J. Solid-St. Circ.*, 51, 2906–2916, <https://doi.org/10.1109/jssc.2016.2591826>, 2016.
- Brezis, H. and Browder, F.: Partial Differential Equations in the 20th Century, *Adv. Math.*, 135, 76–144, <https://doi.org/10.1006/aima.1997.1713>, 1998.
- Calude, C. S., Pačun, G., and Tačar, M.: A Glimpse into natural computing. Centre for Discrete Mathematics and Theoretical Computer Science, The University of Auckland, New Zealand, available at: <https://www.cs.auckland.ac.nz/research/groups/CDMTCS/researchreports/download.php?selected-id=93> (last access: 2 August 2021), 1999.
- Chu, C.: Numerical Methods in Fluid Dynamics, *Adv. Appl. Mech.*, 18, 285–331, [https://doi.org/10.1016/S0065-2156\(08\)70269-2](https://doi.org/10.1016/S0065-2156(08)70269-2), 1979.
- Cockburn, B. and Shu, C.-W.: Runge–Kutta Discontinuous Galerkin Methods for Convection-Dominated Problems, *J. Sci. Comput.*, 16, 173–261, <https://doi.org/10.1023/A:1012873910884>, 2001.
- Cowan, G., Melville, R. C., and Tsividis, Y. P.: A VLSI analog computer/math co-processor for a digital computer, *ISSCC Dig. Tech. Pap. I*, San Francisco, CA, USA, 10–10 February 2005, vol. 1, pp. 82–586, <https://doi.org/10.1109/ISSCC.2005.1493879>, 2005.
- Cowan, G. E. R.: A VLSI analog computer/math co-processor for a digital computer, PhD thesis, Columbia University, available at: <http://www.cisl.columbia.edu/grads/gcowan/vlsianalog.pdf> (last access: 2 August 2021), 2005.
- Dahlquist, G. and Jeltsch, R.: Generalized disks of contractivity for explicit and implicit Runge-Kutta methods, Royal Institute of Technology, Stockholm, Sweden, available at: https://www.sam.math.ethz.ch/sam_reports/reports_final/reports2008/2008-20.pdf (last access: 2 August 2021), 1979.
- Deaton, R., Garzon, M., Rose, J., Franceschetti, D., and Stevens, S.: DNA Computing: A Review, *Fund. Inform.*, 35, 231–245, <https://doi.org/10.3233/FI-1998-35123413>, 1998.
- de Melo, A. C.: The New Linux ‘perf’ Tools, Tech. rep., available at: <http://vger.kernel.org/~acme/perf/lk2010-perf-paper.pdf> (last access: 27 July 2021), 2010.
- Diot, S., Loubère, R., and Clain, S.: The MOOD method in the three-dimensional case: Very-High-Order Finite Volume Method for Hyperbolic Systems, *Int. J. Numer. Meth. Fl.*, 73, 362–392, <https://doi.org/10.1002/fld.3804>, 2013.
- Dumbser, M., Balsara, D. S., Toro, E. F., and Munz, C.-D.: A unified framework for the construction of one-step finite volume and discontinuous Galerkin schemes on unstructured meshes, *J. Comput. Phys.*, 227, 8209–8253, <https://doi.org/10.1016/j.jcp.2008.05.025>, 2008.
- Fambri, F., Dumbser, M., Köppel, S., Rezzolla, L., and Zanotti, O.: ADER discontinuous Galerkin schemes for general-relativistic ideal magnetohydrodynamics, *Mon. Not. R. Astron. Soc.*, 477, 4543–4564, <https://doi.org/10.1093/mnras/sty734>, 2018.
- Feilmeier, M.: *Hybridrechnen*, Springer, Basel, <https://doi.org/10.1007/978-3-0348-5490-0>, 1974.
- Georgescu, I. M., Ashhab, S., and Nori, F.: Quantum simulation, *Rev. Mod. Phys.*, 86, 153–185, <https://doi.org/10.1103/revmodphys.86.153>, 2014.
- Gruber, T., Eitzinger, J., Hager, G., and Wellein, G.: LIKWID 5: Lightweight Performance Tools, Zenodo [data set], <https://doi.org/10.5281/zenodo.4275676>, 2020.
- Gustafson, J. L.: Reevaluating Amdahl's law, *Commun. ACM*, 31, 532–533, <https://doi.org/10.1145/42411.42415>, 1988.
- Hager, G., Wellein, G., and Treibig, J.: LIKWID: A Lightweight Performance-Oriented Tool Suite for x86 Multicore Environments, in: 2010 39th International Conference on Parallel Processing Workshops, IEEE Computer Society, Los Alamitos, CA, USA, 13–16 September 2010, pp. 207–216, <https://doi.org/10.1109/ICPPW.2010.38>, 2010.
- Harten, A.: High resolution schemes for hyperbolic conservation laws, *J. Computat. Phys.*, 135, 260–278, <https://doi.org/10.1006/jcph.1997.5713>, 1997.
- Hirsch, C.: Numerical computation of internal and external flows, in: *Computational Methods for Inviscid and Viscous Flows*, vol. 2, John Wiley & Sons, Chichester, England and New York, 1990.
- Huang, Y., Guo, N., Seok, M., Tsividis, Y., Mandli, K., and Sethumadhavan, S.: Hybrid analog-digital solution of nonlinear partial differential equations, in: Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture, ACM, 665–678, <https://doi.org/10.1145/3123939.3124550>, 2017.
- Karplus, W. and Russell, R.: Increasing Digital Computer Efficiency with the Aid of Error-Correcting Analog Subroutines, *IEEE T. Comput.*, C-20, 831–837, <https://doi.org/10.1109/TC.1971.223357> 1971.
- Kendon, V. M., Nemoto, K., and Munro, W. J.: Quantum analogue computing, *Philos. T. R. Soc. A*, 368, 3609–3620, <https://doi.org/10.1098/rsta.2010.0017>, 2010.
- Köppel, S.: Towards an exascale code for GRMHD on dynamical spacetimes, *J. Phys. Conf. Ser.*, 1031, 012017, <https://doi.org/10.1088/1742-6596/1031/1/012017>, 2018.
- MacLennan, B. J.: Natural computation and non-Turing models of computation, *Theor. Comput. Sci.*, 317, 115–145, <https://doi.org/10.1016/j.tcs.2003.12.008>, 2004.
- MacLennan, B. J.: Analog Computation, in: *Computational Complexity*, edited by: Meyers, R., Springer, New York, pp. 161–184, https://doi.org/10.1007/978-1-4614-1800-9_12, 2012.
- MacLennan, B. J.: Unconventional Computing, University of Tennessee, Knoxville, Tennessee, USA, available at: <http://web.eecs>.

- utk.edu/~bmaclenn/Classes/494-594-UC/handouts/UC.pdf (last access: 27 July 2021), 2019.
- Michoski, C., Milosavljević, M., Oliver, T., and Hatch, D. R.: Solving differential equations using deep neural networks, *Neurocomputing*, 399, 193–212, <https://doi.org/10.1016/j.neucom.2020.02.015>, 2020.
- Nomura, T. and Deiters, R.: Improving the analog simulation of partial differential equations by hybrid computation, *Simulation*, 11, 73–80, <https://doi.org/10.1177/003754976801100207>, 1968.
- Reihing, J.: A time-sharing analog computer, in: Proceedings of the western joint computer conference, San Francisco, CA, USA, 3–5 March 1959, 341–349, <https://doi.org/10.1145/1457838.1457904>, 1959.
- Rodgers, D. P.: Improvements in Multiprocessor System Design, *SIGARCH Comput. Archit. News*, 13, 225–231, <https://doi.org/10.1145/327070.327215>, 1985.
- Röhl, T., Eitzinger, J., Hager, G., and Wellein, G.: LIKWID Monitoring Stack: A Flexible Framework Enabling Job Specific Performance monitoring for the masses, 2017 IEEE International Conference on Cluster Computing (CLUSTER), Honolulu, HI, USA, 5–8 September 2017, 781–784, <https://doi.org/10.1109/CLUSTER.2017.115>, 2017.
- Schenck, C. and Fox, D.: Spnets: Differentiable fluid dynamics for deep neural networks, in: Conference on Robot Learning, PMLR, 87, 317–335, available at: <http://proceedings.mlr.press/v87/schenck18a.html> (last access: 27 July 2021), 2018.
- Schuman, C. D., Potok, T. E., Patton, R. M., Birdwell, J. D., Dean, M. E., Rose, G. S., and Plank, J. S.: A Survey of Neuromorphic Computing and Neural Networks in Hardware, arXiv [preprint], arXiv:1705.06963v1, 19 May 2017.
- Shu, C.-W.: High order WENO and DG methods for time-dependent convection-dominated PDEs: A brief survey of several recent developments, *J. Comput. Phys.*, 316, 598–613, <https://doi.org/10.1016/j.jcp.2016.04.030>, 2016.
- Siegelmann, H. T.: Computation Beyond the Turing Limit, *Science*, 268, 545–548, <https://doi.org/10.1126/science.268.5210.545>, 1995.
- Sod, G.: Numerical Methods in Fluid Dynamics: Initial and Initial Boundary-Value Problems, Cambridge University Press, Cambridge, UK, 1985.
- Subramaniam, B., Saunders, W., Scogland, T., and Feng, W.-c.: Trends in Energy-Efficient Computing: A Perspective from the Green500, in: Proceedings of the International Green Computing Conference, Arlington, VA, USA, 27–29 June 2013, <https://doi.org/10.1109/IGCC.2013.6604520>, 2013.
- Subramaniam, B., Scogland, T., Feng, W.-c., Cameron, K. W., and Lin, H.: Green 500 List, 2020, available at: <http://www.green500.org> (last access: 28 July 2021), 2020.
- Titarev, V. A. and Toro, E. F.: ADER: Arbitrary High Order Godunov Approach, *J. Sci. Comput.*, 17, 609–618, <https://doi.org/10.1023/A:1015126814947>, 2002.
- Titarev, V. A. and Toro, E. F.: ADER schemes for three-dimensional non-linear hyperbolic systems, *J. Comput. Phys.*, 204, 715–736, <https://doi.org/10.1016/j.jcp.2004.10.028>, 2005.
- Toro, E. F.: Primitive, Conservative and Adaptive Schemes for Hyperbolic Conservation Laws, in: Numerical Methods for Wave Propagation. Fluid Mechanics and Its Applications, edited by: Toro, E. F. and Clarke, J. F., vol. 47, Springer, Dordrecht, the Netherlands, pp. 323–385, https://doi.org/10.1007/978-94-015-9137-9_14, 1998.
- Ulmann, B.: Model-1 Analog Computer Handbook/User Manual, available at: <http://analogparadigm.com/downloads/handbook.pdf> (last access: 28 July 2021), 2019.
- Ulmann, B.: Analog and Hybrid Computer Programming, De Gruyter Oldenbourg, Berlin, Boston, <https://doi.org/10.1515/9783110662207>, 2020.
- Vichnevetsky, R.: A new stable computing method for the serial hybrid computer integration of partial differential equations, in: Spring Joint Computer Conference, Atlantic City New Jersey, 30 April–2 May 1968, 143–150, <https://doi.org/10.1145/1468075.1468098>, 1968.
- Vichnevetsky, R.: Hybrid methods for partial differential equations, *Simulation*, 16, 169–180, 1971.
- Volynskii, B. A. and Bukham, V. Y.: Analogues for the Solution of Boundary-Value Problems, 1st edn., in: International Tracts in Computer Science and Technology and Their Application, Oxford, London, 1965.
- Wang, W., Zhu, Y., Chan, C.-H., and Martins, R. P.: A 5.35-mW 10-MHz Single-Opamp Third-Order CT Delta Sigma Modulator With CTC Amplifier and Adaptive Latch DAC Driver in 65-nm CMOS, *IEEE J. Solid-St. Circ.*, 53, 2783–2794, <https://doi.org/10.1109/jssc.2018.2852326>, 2018.
- Wang, Y., Yu, B., Berto, F., Cai, W., and Bao, K.: Modern numerical methods and their applications in mechanical engineering, *Adv. Mech. Eng.*, 11, 1–3, <https://doi.org/10.1177/1687814019887255>, 2019.
- Wilhelm, F., Steinwandt, R., Langenberg, B., Liebermann, P., Messinger, A., Schuhmacher, P., and Misra-Spieldenner, A.: Status of quantum computer development, Version 1.2, BSI Project Number 283, Federal Office for Information Security, available at: https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/Publikationen/Studien/Quantencomputer/P283_QC_Studie-V_1_2.html (last access: 28 July 2021), 2020.
- Zhou, Y., Stoudenmire, E. M., and Waintal, X.: What Limits the Simulation of Quantum Computers?, *Phys. Rev. X*, 10, 041038, <https://doi.org/10.1103/physrevx.10.041038>, 2020.
- Ziegler, M.: Novel hardware and concepts for unconventional computing, *Sci. Rep.*, 10, 11843, <https://doi.org/10.1038/s41598-020-68834-1>, 2020.