

The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale

Andreas Müller¹, Willem Deconinck¹, Christian Kühnlein¹, Gianmarco Mengaldo¹, Michael Lange¹, Nils Wedi¹, Peter Bauer¹, Piotr K. Smolarkiewicz¹, Michail Diamantakis¹, Sarah-Jane Lock¹, Mats Hamrud¹, Sami Saarinen¹, George Mozdzyński¹, Daniel Thiemert¹, Michael Glinton², Pierre Bénard², Fabrice Voitus², Charles Colavolpe², Philippe Marguinaud², Yongjun Zheng², Joris Van Bever³, Daan Degrauwe³, Geert Smet³, Piet Termonia^{3,4}, Kristian P. Nielsen⁵, Bent H. Sass⁵, Jacob W. Poulsen⁵, Per Berg⁵, Carlos Osuna⁶, Oliver Fuhrer⁶, Valentin Clement⁷, Michael Baldauf⁸, Mike Gillard⁹, Joanna Szmelter⁹, Enda O'Brien¹⁰, Alastair McKinstry¹⁰, Oisín Robinson¹⁰, Parijat Shukla¹⁰, Michael Lysaght¹⁰, Michał Kulczewski¹¹, Miłosz Ciznicki¹¹, Wojciech Piątek¹¹, Sebastian Ciesielski¹¹, Marek Błażewicz¹¹, Krzysztof Kurowski¹¹, Marcin Procyk¹¹, Paweł Spychala¹¹, Bartosz Bosak¹¹, Zbigniew P. Piotrowski¹², Andrzej Wyszogrodzki¹², Erwan Raffin¹³, Cyril Mazaurec¹³, David Guibert¹³, Louis Douriez¹³, Xavier Vigouroux¹³, Alan Gray¹⁴, Peter Messmer¹⁴, Alexander J. Macfaden¹⁵, and Nick New¹⁵

¹European Centre for Medium-Range Weather Forecasts (ECMWF), Reading, UK

²Centre National de Recherches Météorologiques, Météo-France, Toulouse, France

³Royal Meteorological Institute (RMI), Brussels, Belgium

⁴Department of Physics and Astronomy, Ghent University, Ghent, Belgium

⁵The Danish Meteorological Institute (DMI), Copenhagen, Denmark

⁶Federal Office of Meteorology and Climatology MeteoSwiss, Zurich, Switzerland

⁷Center for Climate System Modeling, Zurich, Switzerland

⁸Deutscher Wetterdienst (DWD), Offenbach, Germany

⁹Loughborough University, Leicestershire, UK

¹⁰Irish Centre for High-End Computing (ICHEC), National University of Ireland, Galway, Ireland

¹¹Poznan Supercomputing and Networking Center (PSNC), Poznań, Poland

¹²Institute of Meteorology and Water Management - National Research Institute (IMGW-PIB), Warsaw, Poland

¹³ATOS, Bezons, France

¹⁴NVIDIA Switzerland, Zurich, Switzerland

¹⁵Optalysys Ltd., Glasshoughton, UK

Correspondence: Andreas Müller (andreas.mueller@ecmwf.int)

Abstract. In the simulation of complex multi-scale flows arising in weather and climate modelling, one of the biggest challenges is to satisfy strict service requirements in terms of time-to-solution and to satisfy budgetary constraints in terms of energy-to-solution, without compromising the accuracy and stability of the application. These simulations require algorithms that minimise the energy footprint along with the time required to produce a solution, maintain the physically required level of accuracy, are numerically stable and are resilient in case of hardware failure.

The European Centre for Medium Range Weather Forecasts (ECMWF) was leading the ESCAPE (Energy-efficient Scalable Algorithms for weather Prediction on Exascale supercomputers) project, funded by Horizon 2020 (H2020) under the FET-HPC (Future and Emerging Technologies in High Performance Computing) initiative. The goal of ESCAPE was to develop

a sustainable strategy to evolve weather and climate prediction models to the next-generation computing technologies. The project partners incorporate the expertise of leading European regional forecasting consortia, university research, experienced high-performance computing centres and hardware vendors.

This paper presents an overview of the ESCAPE strategy: (i) identify domain-specific key algorithmic motifs in weather prediction and climate models (Weather & Climate dwarfs), (ii) categorise them in terms of computational and communication patterns, while (iii) adapting them to different hardware architectures with alternative programming models, (iv) analyse the challenges in optimising and (v) find alternative algorithms for the same scheme. The participating weather prediction models are: IFS (Integrated Forecasting System); ALARO, a combination of AROME (Application de la Recherche à l'Opérationnel à Meso-Echelle) and ALADIN (Aire Limitée Adaptation Dynamique Développement International); and COSMO-EULAG, a combination of COSMO (Consortium for Small-scale Modeling) and EULAG (Eulerian/semi-Lagrangian fluid solver). For many of the Weather & Climate dwarfs ESCAPE provides prototype implementations on different hardware architectures (mainly Intel Skylake CPUs, NVIDIA GPUs, Intel Xeon Phi, Optalysys optical processor) with different programming models. The spectral transform dwarf represents a detailed example of the co-design cycle of an ESCAPE dwarf.

The dwarfs concept has proven to be extremely useful for rapid prototyping of alternative algorithms and their interaction with hardware; e.g. the use of a domain-specific language (DSL). Manual adaptations have led to substantial accelerations of key algorithms in numerical weather prediction (NWP), but are not a general recipe for performance portability of complex NWP models. Existing DSLs are found to require further evolution but are promising tools for achieving the latter. Measurements of energy- and time-to-solution suggest that a future focus needs to be on exploiting simultaneous use of all available resources in hybrid CPU/GPU arrangements.

20 1 Introduction

Numerical weather and climate prediction capabilities represent substantial socio-economic value in multiple sectors of human society, namely for the mitigation of the impact of extremes, in food production, renewable energy and water management, infrastructure planning, and for finance and insurance where weather sensitive goods and services are traded. Despite significant progress achieved over past decades (Bauer et al., 2015) there are substantial shortcomings in our ability to predict, for example, weather extremes with sufficient lead time and the impact of climate change at regional or national level. Extreme weather events caused over 500 thousand casualties and over 2 trillion US dollars economic damages in the past 20 years (Wallemacq and House, 2018). Another example is the prediction of climate change mitigation and adaptation targets. Failure to meet these targets is ranked among the leading threats to global society (World Economic Forum, 2019).

One of the key sources of model error is limited spatial (and temporal) resolution (Palmer, 2014) which implies that key physical processes that drive global circulation, like deep convection in the tropics and mesoscale eddies in the ocean, are only crudely represented in models by the contemporary parameterisations. In addition, deficiencies in the representation of process interactions between atmosphere and ocean/sea-ice, as well as atmosphere and land, including a time-varying biosphere, are highly relevant strategic targets to improve the representation of the internal variability of the Earth system in models (Katzav

and Parker, 2015). However, better spatial resolution and enhanced model complexity translate immediately into significant computational challenges (Schulthess et al., 2019), whereby spatial resolution is the most demanding because a doubling of resolution roughly translates into eight times more computations (doubling the horizontal resolution in both directions and a corresponding decrease in the time step). The critical resolution threshold at which global weather & climate models may eventually be able to overcome the bulk of the limiting deficiencies is unclear; however, O(1km) emerges as a defensible intermediate target (Shukla et al., 2010; Neumann et al., 2019).

The ESCAPE project has been motivated by the need for running Earth-system models at much higher resolution and complexity than presently available, but within the same time-critical path of daily production for weather forecasts and with the same production throughput needed for decadal/centennial climate projections as used today. This translates to computing one simulated year per wallclock day. Energy efficiency is a key requirement as power envelopes for HPC systems cannot be scaled up at the same rate as the computing demand. Obviously, this presents a substantial challenge to all aspects of computing, namely the choice and implementation of numerical methods and algorithms, and the programming models to map memory access and communication patterns onto specific hardware (Wehner et al., 2011).

The ECMWF specific strategic target is to provide skilful predictions of high-impact weather up to two weeks ahead by 2025 based on a global 5 km-scale Earth-system model ensemble, complemented by the corresponding, high-resolution data assimilation system for creating the initial conditions. This system needs sufficient model complexity in the atmosphere (physics and chemistry), oceans, sea-ice and land to represent all processes acting on such scales, and a sufficiently large ensemble size so that complex probability distributions of extreme forecast features can be sampled well enough. In terms of computational demand this is comparable to the 1-km resolution target for a single forecast.

The supercomputers used for numerical weather prediction (NWP) have changed dramatically over the past decades, and ECMWF's Integrated Forecasting System (IFS; Wedi et al., 2015, and references therein) has exhibited rather different levels of sustained performance ranging from 40-50% on the parallel vector-processor Cray-XMP, YMP and C90 architectures between 1984 and 1996, to 30% on Fujitsu's similar-type VPP-700 and 5000 systems between 1996 and 2002, down to 5-10% on the scalar multi-core IBM P5-7 and Cray XC30 and 40 architectures operated between 2002 and today. Despite sustained performance declining, overall FLOP performance increased exponentially owing to Moore's law, Dennard scaling and processor pricing (e.g. Flamm, 2018) so that significant upgrades of model resolution, complexity and ensemble size remained affordable in terms of capital cost and power supply. However, as this technological evolution is slowing down, new concepts for designing algorithms and for mapping the associated computational patterns onto the available architectures - from processor to system level - are needed (Lawrence et al., 2018). Many of the algorithms used in NWP were designed well before the multi-core era, but even though they contain highly-tuned shared and distributed memory parallelisation, they only achieve such limited sustained performance because of (i) a poor ratio of data communication to computations (viz. arithmetic density) in some of the compute kernels (hereinafter "kernels" for brevity), (ii) sequential tasking due to strong physical/data dependencies, and (iii) algorithmic intricacies required to solve accurately and efficiently a multi-scale and multi-phase fluid dynamics problems on a rotating sphere.

If the envisioned increase in model fidelity is constrained by only marginally growing power envelopes and decelerating general purpose processor speed, then performance issues need to be addressed at the root, and a more radical redesign of the basic algorithms and their implementations used for weather prediction needs to be considered. This is why ESCAPE investigates both HPC adaptation and alternative numerical formulations of physically and computationally clearly identifiable model components - herein introduced the Weather & Climate dwarfs - followed by a detailed analysis of their respective computational bottlenecks, and subsequent hardware and algorithm dependent optimisations.

The Weather & Climate dwarf idea is introduced in Section 2. In Section 3 the usefulness of the dwarf concept is illustrated with the example of the spectral transform dwarf. The paper ends with conclusions and outlook in Section 4.

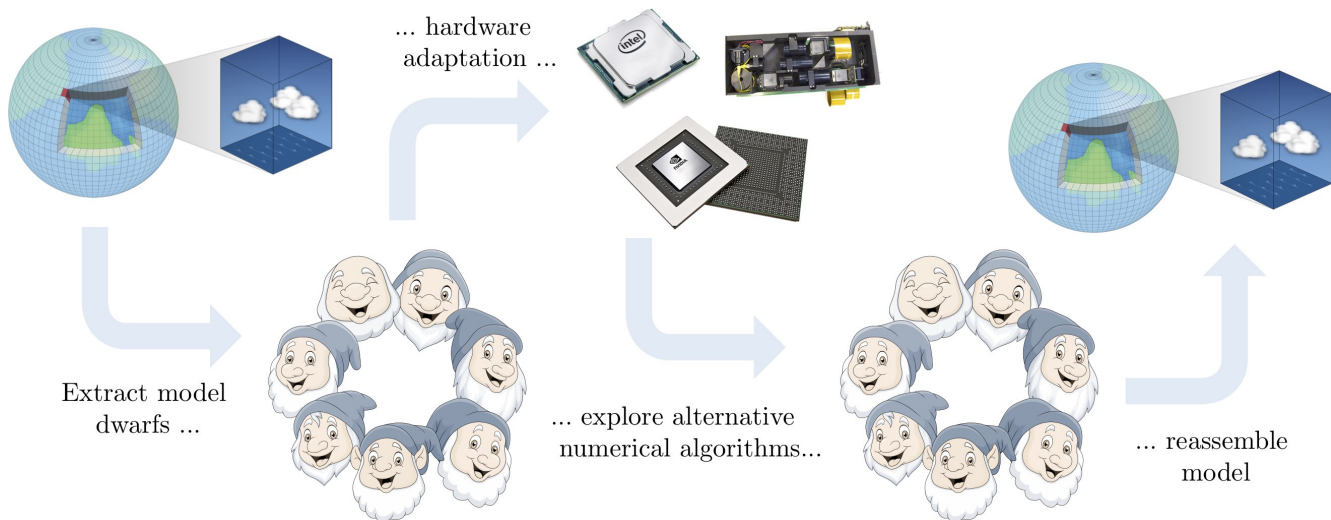
2 Weather & Climate dwarfs

2.1 Motivation

In 2004, Phillip Colella introduced the seven dwarfs of algorithms for high-end simulation in the physical sciences (Colella, 2004, quoted after Kaltofen (2011)). These were later extended to the 13 Berkeley dwarfs (Asanović et al., 2006, 2009) which are meant to represent characteristic patterns of computation and communication. These dwarfs were created to cover the characteristic computational properties of a broad range of scientific applications. These dwarfs are the basis for the OpenDwarfs benchmark suite (Feng et al., 2012; Krommydas et al., 2015; Johnston and Milthorpe, 2018) and were also applied to benchmark cloud computing in Phillips et al. (2011). In a similar fashion Kaltofen (2011) introduced the seven dwarfs of symbolic computation.

Following this idea, we categorise key algorithmic motifs specific to weather prediction and climate models and identify their specific computational and communication patterns, which in return are crucial for the entire model performance. The dwarfs thus represent domain specific mini-applications (Messer et al., 2016) which include direct input from the domain scientist together with documentation, timers for profiling purposes as well as error estimates for verification purposes. In this way the dwarfs facilitate communication of NWP domain specific knowledge on algorithmic motifs with specialists in other domains. Different implementations of the dwarfs can be used as a first step towards optimising and adapting weather prediction and climate models to new hardware architectures and to benchmark current and future supercomputers with simpler but relevant applications. Identifying these key algorithms thus allows better collaboration between operational weather prediction centres, hardware vendors and academia. The concept of dwarfs is different from the existing separation of weather and climate models into different model components, such as atmosphere, land-surface and ocean for which separate dynamical core and physical parameterisation packages already exist. Instead, dwarfs define a standalone and more manageable sub-component in a hierarchy of model complexity for specific targets such as adaptation to GPUs, exploring alternative programming models, and developing performance portable domain specific languages. But dwarfs can also be used by domain scientists for developing alternative algorithms, even across model components.

The fundamental starting point of the ESCAPE project is to identify the dwarfs in the participating weather and climate models (Fig. 1) and to adapt them to different hardware architectures. The knowledge gained in this adaptation is used to



Earth illustration: used under license from GraphicsRF/Shutterstock.com.
 Dwarf illustrations: used under license from Teguh Mujiono/Shutterstock.com

Figure 1. Illustration of the main idea behind the ESCAPE project. The entire model is broken down into smaller building blocks called dwarfs. These are adapted to different hardware architectures. Based on the feedback from hardware vendors and high performance computing centres alternative numerical algorithms are explored. These improvements are eventually built into the operational model.

research alternative numerical algorithms which are better suited for those new architectures, and experiment with alternative programming models, towards improving the overall energy-efficiency of weather prediction applications.

2.2 List of ESCAPE dwarfs

Table 1 gives an overview of the dwarfs defined in the ESCAPE project. These dwarfs have been chosen because they represent
 5 key algorithms that are representative of any weather and climate model, or because their computational and communication patterns represent one of the most runtime consuming parts of weather forecasting systems (Fig. 2). Prototype implementations addressing specific hardware (by default CPU) were developed for selected dwarfs. Each prototype implementation comes with documentation (Mengaldo, 2016; Müller et al., 2017) as well as timers and error measures. The timers enable a quick assessment of the speedup due to the optimisations, whereas the error measures verify to what extent the optimisations affect
 10 the results. Table 1 also lists the models in which each dwarf was identified and from which its prototype implementations have originated. The models include the IFS with the spectral transform method (IFS-ST), the Finite-Volume Module of the IFS (IFS-FVM), ALARO/AROME and COSMO-EULAG. In order to explore alternative discretizations in ESCAPE, a new global shallow water model called GRASS (Global Reduced A-grid Spherical-coordinate System) was developed. Table 1 further shows which dwarfs have prototypes that are based on the data structure framework Atlas (Deconinck et al., 2017), used to ease
 15 adaptation to future architectures and to avoid code duplication across prototypes. Atlas handles both the mesh generation and parallel communication aspects. In the ESCAPE project, Atlas has been improved and extended by adding support for limited

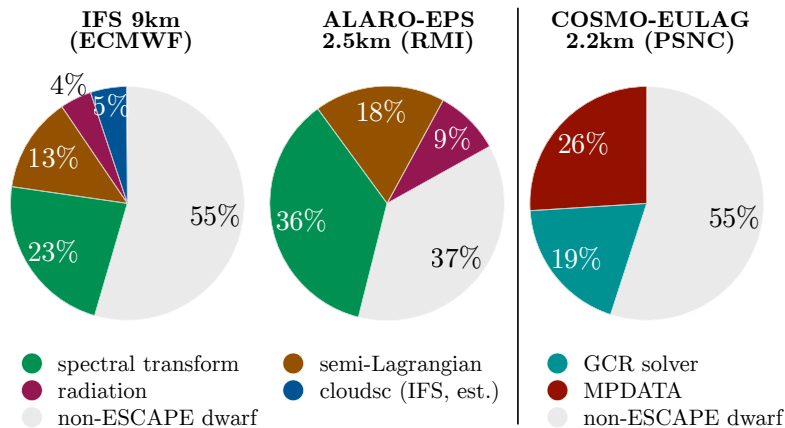


Figure 2. Portion of the forecast runtime spent in ESCAPE dwarfs for the three different models, IFS (left), ALARO ensemble prediction system (EPS) (middle) and COSMO-EULAG (right). The measurements for the IFS were taken during an operational run on 352 nodes (1408 MPI processes, 18 OpenMP threads per process). The limited area models ALARO-EPS and COSMO-EULAG used each 576 MPI processes for the simulations shown here. The IFS and ALARO-EPS are both based on the spectral transform method with latter employing 2D FFTs. For this reason the two left pie charts share one legend. COSMO-EULAG uses different methods and the pie chart has its own legend. The vertical line separates the two different legends.

area grids as well as for accelerators via GridTools (Deconinck, 2017a, b). Table 1 also shows the programming model adopted for each of the dwarf prototypes. In particular, MPI was used for distributed-memory parallelism, OpenMP and OpenACC for shared-memory parallelism, and two domain specific language (DSL) solutions - CLAW and GridTools - for hardware specific optimisations. Apart from DSL, all dwarfs were coded in modern Fortran.

5 In addition to identifying computational and communication patterns in existing models, there was a search for new algorithms that show significant potential to reduce runtime and energy consumption. This led to the developments of a multigrid preconditioner that reduces iterations of the Krylov-subspace solver, employed in the semi-implicit time integration of IFS-FVM (Müller et al., 2017), and a HEVI time-integration scheme with significantly improved stability (Colavolpe et al., 2017). The latter avoids global communications, which also motivated the development of high-order finite-difference methods for
 10 reduced longitude-latitude grids on the sphere (Bénard and Glinton, 2019; Glinton and Bénard, 2019). Furthermore, on the hardware side, ESCAPE explored FFTs and spherical harmonics on optical (co-)processors (Macfaden et al., 2017) to potentially scale these transforms towards higher resolutions at a fixed energy cost. An overview of the work with the optical processor can be found in Sect. 3.5.

As an example of the cycle involved in identifying, isolating, testing, adapting, optimising, and considering alternative
 15 solution procedures for a specific algorithmic motif in the ESCAPE project, the work on the spectral transform dwarf is illustrated in more detail.

Table 1. Overview of dwarfs identified in the ESCAPE project, short description of the computational characteristics of each dwarf and the models in which we identified the dwarf. Standalone applications including documentation, profiling and verification capabilities called prototypes have been released for many of the dwarfs. The column Atlas shows which of the dwarfs have prototypes that are based on the data structure framework Atlas (Deconinck et al., 2017). Dwarfs for which prototypes with MPI and OpenMP/OpenACC is available can be used with both programming models together as hybrid parallelisation. LAITRI is used in the full model in each MPI process individually and therefore does not include MPI parallelisation inside the dwarf. Further explanations can be found in the references given in the last column.

	computational characteristics	IFS-ST	IFS-FVM	ALARO/AROME	COSMO-EULAG	GRASS	prototype released	Atlas	MPI	OpenMP	OpenACC	DSL (Sect. 3.7)	Optalysys (Sect. 3.5)	references
Dynamics dwarfs:														
• time-stepping: spectral transform: – spherical harmonics – biFFT elliptic solver - GCR elliptic solver - multigrid HEVI	matrix multiplication, global transposition 1D and 2D FFT algorithm, global transposition iteration, global reduction interpolation for remapping nearest neighbour column communication	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Mengaldo (2016) Mengaldo (2016) Mengaldo (2016) Müller et al. (2017) Mengaldo et al. (2018); Colavolpe et al. (2017)
• advection: semi-Lagrangian – LAITRI (3d interpolation) MPDATA	large halo communication efficient data access small halo communication, indirect addressing ²	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Mengaldo (2016) Robinson et al. (2016) Mazauric et al. (2017b); Osuna (2018)
• gradient computation: high order finite difference finite volume	large halo communication small halo communication, indirect addressing ²				✓			✓	✓	✓				Bénard and Glimton (2019); Glimton and Bénard (2019) Kühnlein et al. (2019)
Physics dwarfs:														
• radiation scheme: ACRANEB2	many transcendental functions		✓				✓	✓	✓	✓				Müller et al. (2017); Poulsen and Berg (2017)
• cloud microphysics: CLOUDSC	single (vertical) column, many if-statements, many local variables	✓	✓				✓	✓	✓	✓				Xiao et al. (2017); Clement et al. (2018)

¹ The GRASS model is currently only solving the shallow-water equations. Research on HEVI time integration methods was performed to explore future options.

² when using unstructured meshes

3 Dwarf example: spectral transform

Below is a short description of the domain-specific spectral transform with a discussion of computational challenges, primarily focused on the data structures and data access patterns used in IFS. Subsequently, we present the work done on adapting and optimising the dwarf for GPUs, CPUs and optical processors, and finish with a comparison of the results obtained for the 5 different architectures and a discussion on the sustainability of the chosen techniques.

3.1 Background

Each time-step of IFS is split between computations in physical space (i.e. a grid point representation) and computations with respect to spectral space (i.e. spectral coefficients at different wavenumbers). Semi-Lagrangian advection, physical parameterisations and products of terms are computed most efficiently in grid point space while horizontal gradients, semi-implicit 10 calculations and horizontal diffusion are computed more efficiently in spectral space. The transform between these two spaces is performed on the sphere with spherical harmonics, that is computing these results along longitudes in a Fast Fourier transform (FFT) and a Legendre transform (LT) along latitudes. Limited area models replace the Legendre transform with another FFT which leads to the name biFFT.

In spectral transform methods, such as the one used in IFS (Wedi et al., 2013), the specific form of the semi-implicit system 15 facilitating large time-steps (and thus time-to-solution efficiency) is derived from subtracting a system of equations linearised around a horizontally homogeneous reference state. The solution of this specific form is greatly accelerated by the separation of the horizontal and the vertical part, which matches the large anisotropy of horizontal to vertical grid dimensions prevalent in atmospheric and oceanic models. In spectral transform methods one uses the special property of the horizontal Laplacian operator in spectral space on the sphere

$$20 \quad \nabla^2 \psi_n^m = -\frac{n(n+1)}{a^2} \psi_n^m, \quad (1)$$

where ψ symbolises a prognostic variable, a is the Earth radius, and (n, m) are the total and zonal wavenumbers of the spectral discretisation (Wedi et al., 2013). This conveniently transforms the 3D Helmholtz problem into an array (for each zonal wavenumber) of 2D matrix operator inversions with the dimension of the vertical levels square - or vertical levels times the maximum truncation N_T , defined shortly, in the case of treating the Coriolis term implicitly - resulting in a very efficient 25 direct solver. The computational aspects and especially the data layout are illustrated with the inverse spectral transform from spectral to grid-point space.

The inverse spectral transform begins with the spectral data $\mathbf{D}(f, i, n, m)$ which is a function of field index f (for the variable surface pressure at a single level and for wind vorticity, wind divergence and temperature at each height level), real and imaginary part i and wave numbers (zonal wave number $m = 0, \dots, N_T$ and total wave number $n = 0, \dots, N_T - m$ where 30 N_T is the spectral truncation). This deviates from the usual notation where total wavenumber goes from m to N_T , because it simplifies the separation between even and odd n , using column-major order like in Fortran, i.e. the field index f is the fastest moving index and the zonal wave number m is the slowest moving index. Typical dimensions can be seen in the operational

high-resolution (9km) forecast run at ECMWF: the number of fields is in this case 412 for the direct transform and 688 for the inverse transform and the number of zonal wave numbers is given by the truncation $N_T = 1279$. The number of latitudes is $2N_T + 2 = 2560$ and the number of longitudes increases linearly from 20 next to the poles to $4N_T + 20 = 5136$ next to the equator.

- 5 We take advantage of the symmetry of the Legendre polynomials for even n and anti-symmetry for odd n . The coefficients of the Legendre polynomials are pre-computed and stored in $\mathbf{P}_{e,m}(n, \phi)$ for even n and $\mathbf{P}_{o,m}(n, \phi)$ for odd n , where ϕ stands for the latitudes of our Gaussian mesh. Only latitudes on the northern hemisphere are computed. Latitudes on the southern hemisphere are reconstructed from the northern latitudes. In the same way the spectral data is split for each m into even part $\mathbf{D}_{e,m}(f, i, n)$ and odd part $\mathbf{D}_{o,m}(f, i, n)$. Fields (indexed f) are all independent in the parallel transform. The inverse Legendre
10 transform is performed by computing the following matrix multiplications using BLAS:

$$\begin{aligned}\mathbf{S}_m(f, i, \phi) &= \sum_n \mathbf{D}_{e,m}(f, i, n) \cdot \mathbf{P}_{e,m}(n, \phi), \\ \mathbf{A}_m(f, i, \phi) &= \sum_n \mathbf{D}_{o,m}(f, i, n) \cdot \mathbf{P}_{o,m}(n, \phi).\end{aligned}\tag{2}$$

The resulting array for the symmetric and antisymmetric parts are now combined into the Fourier coefficients on the northern and southern hemisphere:

$$\begin{aligned}\phi > 0 : \mathbf{F}(i, m, \phi, f) &= \mathbf{S}_m(f, i, \phi) + \mathbf{A}_m(f, i, \phi), \\ 15 \quad \phi < 0 : \mathbf{F}(i, m, \phi, f) &= \mathbf{S}_m(f, i, \phi) - \mathbf{A}_m(f, i, \phi).\end{aligned}$$

These Fourier coefficients are finally used to compute the fields in grid point space at each latitude ϕ and longitude λ via FFT:

$$\mathbf{G}_{\phi, f}(\lambda) = \text{FFT}(\mathbf{F}_{\phi, f}(i, m)).\tag{3}$$

3.2 Computational challenges

- The computations in grid point space require all fields to be on the same computational node. The summation over the total
20 wavenumber n in the Legendre transform (2) makes it most efficient to have all total wavenumbers (for a given field) on the same node and the Fourier transform (3) over (i, m) makes it most efficient to have all of the zonal wavenumbers m with real and imaginary part on the same node. This is only possible if the data is transposed before and after the spectral transform as well as in between Legendre and Fourier transform. These transpositions produce substantial communication which increases the contribution of the spectral transform to the overall runtime for future resolutions (Fig. 3).

- 25 Simplified simulations of the MPI communications performed in ESCAPE indicate that the strong scalability of the communication time for the spectral transform transpositions is better than for halo communication required by (large) semi-Lagrangian advection halos and global norm computation commonly used in semi-implicit methods (Zheng and Marguinaud, 2018).

- These results indicate that halo communication will become almost as costly as the transpositions in the spectral transform
30 method if a very large number of MPI processes is involved. An alternative which avoids transpositions and halo communication is given by the spectral element method shown in Müller et al. (2018) with explicit time integration in the horizontal




Figure 3. Cost profiles of the significant components of the IFS NWP model in percent of CPU time, at the operational 9km horizontal resolution with 137 vertical levels (left), the anticipated future horizontal resolution of 5km with 137 vertical levels (middle) and an experimental resolution of 1.45km with 62 vertical levels (right). The grid point dynamics represents the advection and gridpoint computations related to the dynamical core, semi-implicit solver represents the computations and communications internal to spectral space, spectral transform relates to the communications and computations in the transpositions from gridpoint to spectral and reverse as well as the FFT and DGEMM computations (see also spectral transform schematic below), physics+radiation relates to the cost of physical parametrisations including radiation, and finally accounting for the additional components of the wave and the ocean models - nonetheless excluded in the 1.45km simulation. All of these profiles have been obtained through measurements on ECMWF’s Cray XC40 supercomputer.

direction. This leads to a very small amount of data that is communicated in each time-step because this method only communicates the values that are located along the interface between different processor domains. This method, however, requires much smaller time-steps which leads overall to an even larger communication volume (Fig. 4). Figure 4 is based on the model comparison presented in Michalakes et al. (2015) and does not include all of the optimisations for the spectral element method presented in Müller et al. (2018). The spectral transform results are based on the operational version of IFS and do not contain the optimisations presented in this paper. Both models have significant potential for optimisation and it is not obvious which method will have the lowest communication volume when fully optimised, but key for both is an approach that allows reasonably large time steps. Regardless, the only way to avoid waiting time during communication is to enable concurrent execution of various model components such that useful computation can be done during data communication (Mozdzynski et al., 2015; Dziekan et al., 2019).

3.3 GPU optimisation

For the GPU version, the prototype is restructured to: a) allow the grid-based parallelism to be fully exposed to (viz. exploited by) the GPU in a flexible manner; b) ensure that memory coalescing is achieved; and c) optimise data management.

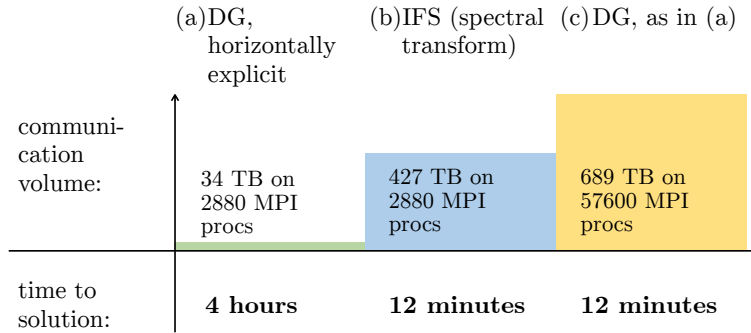


Figure 4. Overall communication volume comparing spectral element (SEM) from Müller et al. (2018) and the global spectral transform methods. The SEM requires a substantially lower amount of communication at the same number of cores, but due to the smaller timestep requires a much higher frequency of repeated communications for the given 2-day simulation. Extrapolating the number of MPI processes to achieve the same time to solution would result in a larger amount of communication for the SEM. Here we assume SEM $\Delta t = 4s$; IFS $\Delta t = 240s$; communication volume is calculated for a 48 hour forecast SEM as 290.4kBytes per MPI task and Δt ; IFS as 216mBytes per MPI task and Δt ; as a result the SEM time-to-solution is estimated as 20 x IFS based on the performance results in Michalakes et al. (2015).

To achieve performance, the two-dimensional grid over the sphere is augmented with a third dimension that represents multiple fields. The computations are inherently parallel across this grid, so all this parallelism can be exposed to the GPU for maximal performance. In contrast, the original implementation had a sequential loop over one of the dimensions (at a high level in the call tree of the application). By restructuring the code for each operation, loops over the three dimensions become tightly nested. These loops are mapped to the GPU via OpenACC directives to collapse to a single loop.

Similarly, for library calls it is important to maximise exposure of parallelism through batching computations using provided interfaces. On the GPU we perform all of the matrix multiplications in the Legendre transform (2) with a single batched call of the cuBlasDgemm library. The different matrices in (2) have different sizes because the total wavenumber goes from 0 to N_T . To use the fully batched matrix multiplication each matrix is padded with zeroes up to the largest size, since the library currently does not support differing sizes within a batch. This step results in a tenfold increase of the overall number of floating point operations but still improves the overall performance (Fig. 5). The FFT in (3) is performed with the cuFFT library, batching over the vertical (field) dimension. Multiple calls are still needed for each latitude since FFTs cannot be padded in a similar way to matrix multiplications. Therefore the implementation remains suboptimal. There would be scope for further improvements if an FFT batching interface supporting differing sizes were to become available.

Sometimes matrix transposes are necessary, but where possible these were pushed into the DGEMM library calls, which have much higher-performing implementations of transposed data accesses. There remain transpose patterns within kernels involved in transposing grid point data from column structure to latitudinal (and inverse) operations, which naturally involve




Figure 5. Roofline plot for the spectral transform dwarf at 125km resolution ($N_T = 159$) on the NVIDIA Tesla P100 GPU. The full time-step of the original prototype is represented by the solid red triangle. The corresponding time-step for the optimised prototype is represented by the red circle. Also included are partial results for kernels only (open green diamond atop the red circle) and matrix multiplication only (green dash). Each point is positioned in the plot according to its operational intensity: points under the sloping region of the roofline are limited by available memory bandwidth, and points under the horizontal region are limited by peak computational performance.

transposes and are thus harder to fix through restructuring. However, using the “tile” OpenACC clause, which instructs the compiler to stage the operation through multiple relatively small tiles, can perform the transpose operations within fast on-chip memory spaces such that the access to global memory is more regular.

Data allocation on the GPU is expensive, as is data movement between the CPU and GPU. Hence the code is structured such that the fields stay resident on the GPU for the whole timestep loop: all allocations/frees have been moved outside the timestep loop with re-use of temporary arrays, and thus all data transfer has been minimized. Arranging the model data consecutive in memory ensures that multiple threads on the GPU can cooperate to load chunks of data from memory in a “coalesced” manner. This allows a high percentage of available memory throughput.

The restructured algorithm achieves an overall speedup factor of $23\times$ compared to the initial version which also used cuBlas and cuFFT but followed the CPU version more closely. Matrix multiplication performance is higher than the overall performance (in flops) and the operational intensity is increased into the compute-bound regime. Note that matrix multiplication is associated with $O(N^3)$ computational complexity for $O(N^2)$ memory accesses. The extra padding operations lead to larger N and therefore also to increased operational intensity. More details about the single GPU optimisations can be found in Mazauric et al. (2017b).

For multiple GPU arrangements there is a large benefit from high-bandwidth interconnects (i.e. NVLink, NVSwitch) due to the relative importance of communication for the transpositions, cf. Sect. 3.2. For full bandwidth connectivity when using




Figure 6. Computational performance of the spectral transform dwarf at 18km resolution ($N_T = 639$) on 4 NVIDIA V100 GPUs of the DGX1 with the original MPI implementation (left), CUDA-aware MPI communication (middle) and NVLink optimised communication (right). This resolution is currently used operationally for the members of the ensemble forecast at ECMWF.

more than 4 GPUs, the NVSwitch interconnect on the DGX-2 server transfers up to 300 GB/s between any pair of GPUs in the system, or equivalently 2.4 TB/s total throughput.

When running a single application across multiple GPUs, it is necessary to transfer data between the distinct memory spaces. Traditionally, such transfers needed to be realised via host memory and required the participation of the host CPU. Not only did this introduce additional latency, but also limited the overall bandwidth to the bandwidth offered by the PCIe bus connecting CPU and GPUs. However, more recent MPI implementations are CUDA-aware. This means that pointers to GPU memory can be passed directly into the MPI calls, avoiding unnecessary transfers (both in the application and in the underlying MPI implementation). This is particularly useful when using a server that features high-bandwidth connections between GPUs, in which case CUDA-aware MPI will use these links automatically. Moving our dwarf to CUDA-aware MPI gave us a speedup of $12\times$ (Fig. 6). However, even with this optimisation the all-to-all operations remained inefficient because communication between different GPUs was not exchanged concurrently. Perfect overlap was achieved by implementing an optimised version of the all-to-all communication phase directly in CUDA using the Inter Process Communication (IPC) API. Using memory handles, rather than pointers, CUDA IPC allows to share memory spaces between multiple processes, thus allowing one GPU to directly access memory on another GPU. This allowed another speedup of about 30% (Fig. 6).

Figure 7 compares a DGX-2 with NVSwitch and the DGX-1 for a spectral transform at 18km resolution ($N_T = 639$). The NVIDIA Multi Process Service allows oversubscription of GPUs; e.g., the 8 GPU result on DGX-2 uses 16 MPI tasks. Oversubscription, a technique common to other architectures, can be beneficial to spread out load imbalances resulting from the grid decomposition and hide latencies.

The scaling on DGX-1V is limited because with an increasing number of GPUs some messages go through the lower-bandwidth PCIe and QPI links and/or Infiniband when scaling across multiple servers. On DGX-2 all 16 GPUs have full




Figure 7. Computational performance of the spectral transform dwarf at 18km resolution ($N_T = 639$) on up to 16 GPUs on one DGX-2 and up to 32 DGX-1 servers connected with Infiniband. The DGX-1V uses MPI for ≥ 8 GPUs (due to lack of AlltoAll links), all others use CUDA IPC. DGX-2 results use pre-production hardware. The points were connected with lines for the purpose of improving readability.

connectivity; i.e. the maximum peak bandwidth of 300 GB/s between each pair of GPUs is available. The performance scales well up to the full 16 GPUs on DGX-2, and is $2.4\times$ faster than the result obtained with the 16 GPU (2-server) DGX-1V. The speedup going from 4 to 16 GPUs on DGX-2 is $3.2\times$, whereas the ideal speedup would be $4\times$. Preliminary investigations reveal that this deviation from ideal scaling is not primarily due to communication overhead but due to load imbalance between the

5 MPI tasks for a given spherical grid decomposition. This indicates that better scaling might be observed with a more balanced decomposition. More details about the multi-node optimisation of the spectral transform dwarf can be found in Douriez et al. (2018).

First results on the supercomputer Summit of the Oak Ridge National Laboratory are shown in Fig. 8. At this large scale we observe a speedup between $12.5\times$ (on 480 nodes) and $23.7\times$ (on 1920 nodes) when comparing the optimised GPU version with the initial CPU version of the spectral transform dwarf on the same number of nodes. These simulations were run at 2.5km

10 resolution ($N_T = 3999$) and use 240 fields. In operational application the number of fields should be larger (see Sect. 3.1). It has been verified that the GPU optimised dwarf from ESCAPE is capable of running on a huge number of GPUs on Summit with a significant speedup.

3.4 CPU optimisation

15 The spectral transform dwarf is based on the operational implementation used in IFS that has been continuously optimised over multiple decades. According to profiling results, it clearly appeared that the main computational intensive kernels are the FFT and matrix multiplication executed by a dedicated highly tuned library (as Intel Mathematics Kernel Library, called MKL). In support of this work different data scope analysis tools were explored; see Mazaauric et al. (2017a). The first optimisation

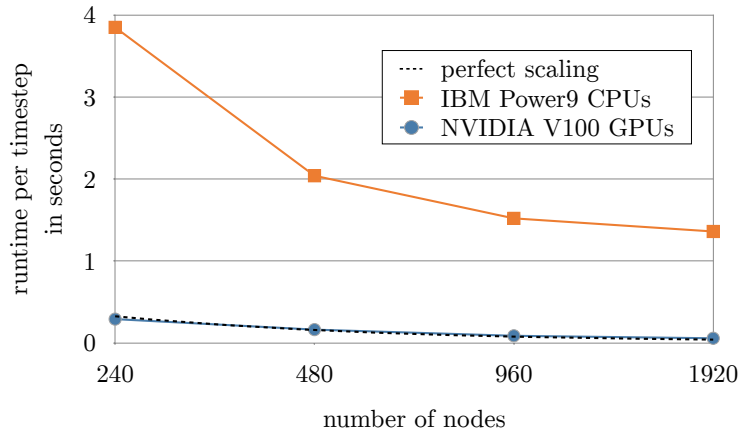


Figure 8. Strong scaling comparison between GPU optimised version and initial CPU version of the spectral transform dwarf on Summit at the Oak Ridge National Laboratory. All of these simulations are using 2.5km resolution ($N_T = 3999$) with 240 fields. The GPU version uses 6 V100 GPUs per node which leads to a maximum number of 11520 GPUs. The perfect scaling was chosen such that it goes through the GPU result at 480 nodes to illustrate the scaling efficiency of the GPU version. The points were connected with lines for the purpose of improving the readability of the plot.

strategy concentrated the effort on non-intrusive optimisations which have the advantage of being portable and maintainable. Among these optimisations, the use of extensions to the x86 instruction set architecture (ISA) as SSE, AVX, AVX2, AVX-512 is notable, because it indicates how much of the source code can be vectorised by the compiler. When the compiler failed at vectorising some loops or loop nests, a deeper investigation of how to use compiler directives followed. As the different instruction sets are not supported by all processors, the study proposed an intra node scalability comparison study among several available systems (at the time of benchmarking).

System tuning using Turbo frequency (TUR), Transparent Huge Page (THP), memory allocator (MAP) can be done without modifying the source code. This exposes both performance gains and interesting information on dwarf behaviour. Indeed, on ATOS BullSequana x400 supercomputer with Intel® Xeon Broadwell E5-2690v4 processors, enabling turbo offers a gain equal to 11%, enabling THP gives 22%, MAP 27%, and finally the best performance (35% of performance gain) is achieved by the combination of MAP and TUR. This shows that memory management is a key point. More details about the single-node CPU optimisation can be found in Mazauric et al. (2017b).

Multi-node optimisation for CPUs focused on improving the MPI communication. The largest potential for optimisations was found to be in the preparation phase of point-to-point communications. During the preparation phase the sender side gathers the local data into a contiguous buffer (Pack operation) and hands it off to the MPI library. On the receiver side, data is then scattered from a contiguous user buffer to its correct location (Unpack operation). Pack and Unpack are nearly inevitable with scattered data because Remote Direct Memory Access (RDMA) with no gather-scatter operations are known to be often less effective, notably due to the memory pinning latency (Wu et al., 2004). It also means that sender and receiver

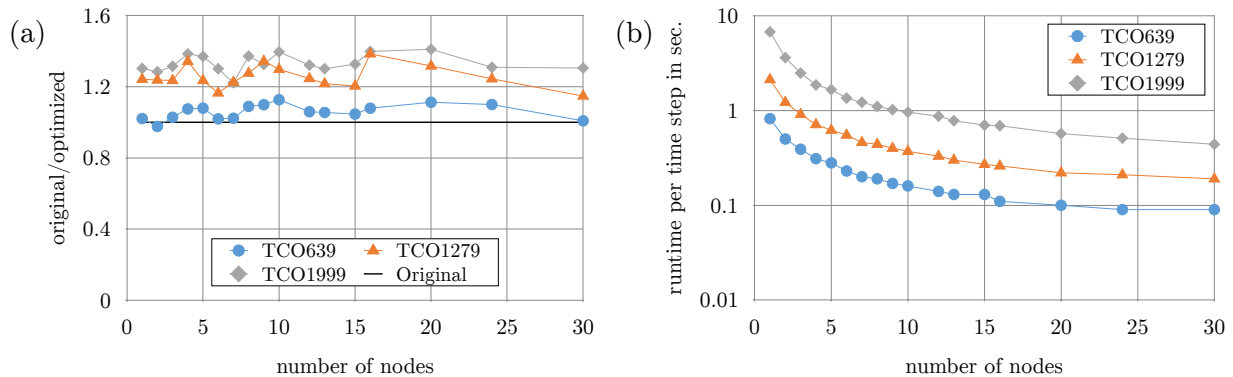


Figure 9. Performance measurements on up to 30 Intel® Xeon Skylake 6148 nodes. Subpanel (a) shows the speed-up at 18km resolution ($N_T = 639$) with regards to different optimisations of the communication preparation phase. Subpanel (b) shows the computational performance for three different resolutions of 18km ($N_T = 639$), 9km ($N_T = 1279$) and 5km ($N_T = 1999$) representative of current and future operational requirements. Note the use of log-scale. The points were connected with lines for the purpose of improving readability. All of these results have been obtained through measurements on the ATOS BullSequana x400 supercomputer of ATOS internal HPC resources.

must exchange their memory layout as they may differ. The performance improvement came from reordering the loops for both pack and unpack following the memory layout of the scattered buffer. This optimisation decreases the number of tests (i.e. copy or not copy) and avoids scanning memory multiple times, leading to a global performance gain on the whole dwarf of about 20% (Fig. 9a). The computational performance up to 30 ATOS BullSequana x400 nodes equipped with Intel® Xeon Skylake 6148 processors is shown in Fig. 9b. The work has been performed on ATOS internal HPC resources. This optimisation can be directly applied to the operational model IFS due to its non-intrusiveness. More details about the multi-node optimisation on CPUs can be found in Douriez et al. (2018).

The speedup of the GPU optimisations came from using highly optimised GPU libraries for batched DGEMM and FFT and avoiding the transposition of temporary arrays. On the CPU, one can equally apply these improvements in the handling of temporary arrays, which lead to the development of a new serial spectral transform inside Atlas, used operationally at ECMWF for post-processing purposes. Post-processing is run in serial mode due to the large number of concurrent post-processing jobs. Compared to the previous serial transform employed in ECMWF product generation, there is a speedup $2\times$ to $3\times$, in particular for limited area domains (Fig. 10).

3.5 Optical processors

The fundamental idea behind optical processors is to encode information into a laser beam by adjusting the magnitude and phase in each point of the beam. Optalysys have been investigating an optical implementation of the spectral transform dwarf - biFFT for limited area models as well as spherical harmonics for global models. This information becomes the Fourier transform of the initial information in the focal plane of a lens. The information can be encoded into the optical beam by using




Figure 10. Speedup of the spectral transform by porting optimisations introduced in the GPU version back to the CPUs. The base line for this comparison is the current operational post-processing library used at ECMWF. This version of the spectral transform allows the computation on limited area domains. The speedup is given here for the global transform and three examples of limited domains (Europe, UK and Hungary).




Figure 11. Illustration of the fundamental idea behind the optical processor. The laser beam is emitted on the bottom left. Two spatial light modulators (SLM) are used together to input the complex function. The system uses beamsplitters (BS) and an optical relay to image one reflective SLM onto another, followed by a lens assembly which approximates an ideal thin lens and renders the optical Fourier transform on a camera sensor. The half-waveplate (WP) before the second SLM is used to rotate linearly polarised light onto the axis of SLM action (the direction in which the refractive index switches), thus causing it to act as a phase modulator.

spatial light modulators (SLMs) as illustrated in Fig. 11. The result of the Fourier transform can be recorded by placing a camera in the focal plane of a lens. A photo of an early prototype is shown in Fig. 12.




Figure 12. Photo of the first prototype of the optical processor. The final product is built into an enclosure of similar size like a GPU.

SLMs are optical devices with an array of pixels which can modulate an optical field as it propagates through (or is reflected by) the device. These pixels can modulate the phase (essentially applying a variable optical delay) or polarisation state of the light. Often they modulate a combination of the two. When combined with polarisers, this polarisation modulation can be converted into an amplitude modulation. Hence the modulation capability of a given SLM as a function of ‘grey level’ can be expressed by a complex vector, which describes an operating curve on the complex plane. Each pixel of the SLM is generally a 1-parameter device; arbitrary complex modulation is not offered by the SLM, only some sub-set. This is one of the key issues with regards to exploiting the optical Fourier transform.

Sensor arrays - essentially common camera sensors - are used to digitise the optical field. They are in general sensitive to the intensity of the light, which is the magnitude of the amplitude squared. This poses a difficulty to sensitively measuring the amplitude. Moreover, sensor arrays are not sensitive to optical phase. The latter can be overcome with a method that allows the optical phase to be derived from the intensity-only measurements. This is achieved by perturbing transformed functions with a specified disturbance, measuring the resulting change in the output, and deducing the phase from the properties of the Fourier transforms.

Each pixel of the SLM is addressable with an 8-bit value (256 levels). The SLM is not capable of independently modulating the magnitude and phase of the optical field. In the Optalysys processing system, the SLMs are configured to modulate both the amplitude and phase in a coupled manner, such that optimal correlation performance is achieved. The optical Fourier transform and all of the functions are inherently two dimensional. The propagating light beam can be thought of as a 2D function propagating and transforming along a third direction. The system is most naturally applied to 2D datasets, and many problems can be mapped to an appropriate representation.

A critical aspect to realizing the potential of optical processing systems is the interface to a traditional computing platform. Bridging this gap has been a significant undertaking for Optalysys, and has resulted in the development of a custom PCIe drive board. This board interfaces to a host machine over PCIe and has direct memory access (DMA) to the system memory (RAM). It provides an interface to 4 SLMs and 2 cameras. The cameras are 4K (4096x3072). Initially, they operate at 100 Hz, but a
5 future firmware upgrade will unlock 300 Hz operation, and 600 Hz half-frame operation, dramatically increasing the potential data throughput.

There are currently two options for the SLMs. One option is using high speed binary (bright-dark) SLMs which operate at 2.4kHz. This offers correlation at binary precision. The second option is greyscale SLMs which operate at 120 Hz. This is currently the only option to reach more than binary precision. The performance of the entire processor is determined by the
10 part with the lowest frequency. The main bottleneck with multiple bit precision is the operating frequency of the greyscale SLM. There is currently no easy solution to increase the frequency of greyscale SLMs.

Optical processing is more appropriately applied to cases where high-throughput relatively-complex operations are the priority, with less of an emphasis on numerical precision. More details about the Optalysys optical processor have been published in Macfaden et al. (2017) and Mazauric et al. (2017b). The results from the optical processing tests will be discussed in the
15 next section.

3.6 Comparison between processors in terms of runtime and energy consumption

A lot of the speedup achieved by running the spectral transform on accelerators is lost if the CPUs are idle during the computation on the accelerators. The cost of data transfer between CPU and accelerator needs to be considered but has not been included in the speedup numbers in this section. To take full advantage of high-bandwidth connected accelerators the entire
20 simulation has to be run on the same node, which requires to substantially reduce the memory footprint of the model.

For the CPUs and GPUs used in this paper the overall cost is dominated by the cost of the hardware and therefore by the number of sockets/devices required to reach the desired runtime. In addition to the number of devices we also compare the energy consumption. The large number of zero operations caused in the optimised GPU version by the padding of the matrices in the Legendre transform makes it impossible to do a fair comparison between CPU and GPU by comparing metrics based on
25 floating point operations including comparing roofline plots.

In the full operational model the 18km resolution ensemble member ($N_T = 639$) using 30 nodes on the Cray XC40 takes about 1.4s per time-step and the spectral transform component is about 15 percent (0.21s) of the overall runtime. Measurements with the dwarf on the Cray XC40 at 18km resolution ($N_T = 639$) resulted in 4.35s per timestep on a single node (4 MPI tasks, 18 threads per task), and 1.77s on 2 nodes. The energy consumption was measured at around 0.3Wh on the Cray XC40, which
30 compares to 0.026 Wh measured on 4 V100 GPUs on a DGX1 which take 0.12s per time-step. The energy measurement on the XC40 is based on proprietary power management counters. The measurement on the V100 GPUs uses the NVIDIA-smi monitoring tool. Tests in ESCAPE on the latest generation of Intel Skylake CPUs have shown 0.12s per timestep using 13 Intel® Xeon Skylake 6148 nodes (connected via a fat-tree EDR network) as shown in Fig. 9. This parallel CPU version has not seen the more radical changes which have been used in redesigning the algorithm for the parallel GPU and serial CPU version.




Figure 13. Log-log plot of the energy consumption vs. wall-clock time for the BiFFT dwarf and corresponding to the combination of one direct and one inverse transformation for 525 fields. Each data point is the result of averaging the outcome of two separate runs. Grey lines connect runs with the same number of OpenMP threads (1 resp. 4). Added are lines of constant power (light blue lines), including the power delivered by a node in the idle state (orange line). Indices next to each data point denote the number of MPI tasks. The black dot represents the estimate of the Optalysys optical processor when using a greyscale SLM. The performance of the optical processor at binary precision is much higher (not shown).

A comparison between CPUs and optical processor with greyscale SLM is shown in Fig. 13. The energy consumption of the optical processor is much lower than for the CPU. The runtime of the optical processor is larger due to the relatively slow performance of the greyscale SLM which is currently necessary to reach the precision required for NWP applications. The latter does not allow to exploit the superior energy-to-solution of optical processors in time-critical NWP applications; more details about the comparison between CPU and optical processor can be found in Van Bever et al. (2018). In general, for hybrid arrangements of CPU and accelerators, we conclude that optimising energy-to-solution requires algorithmic changes that facilitate the simultaneous use of all the available processors.

3.7 Sustainability of code optimisation techniques

The optimisation results presented in Sect. 3.6 are just an example for one of the dwarfs. Similar improvements have been achieved for the advection scheme MPDATA (Douriez et al., 2018) and the radiation scheme ACRANE2 (Poulsen and Berg, 2017). These optimisations are hardware specific and will be difficult to maintain on upcoming new architectures. As a strategy towards performance portability and sustainability we have worked in the ESCAPE project on domain specific languages through the GridTools framework. The main goal of the GridTools library is to provide a solution for weather and climate models to run one code on many different architectures (portability) and achieve good performance (performance portability).

However, the main operational product of GridTools so far focused on solutions for lat-lon grid models like COSMO. The work developed in the ESCAPE project aimed at extending the DSL support for irregular grids, the efficient generation of backends for multiple architectures and evaluating DSL useability.

The DSL developments have been used to implement a portable version of the MPDATA dwarf. The DSL version hides such details as the nested loops and the OpenACC directives used to specify properties of the GPU kernel and data layouts of the Fortran arrays. Furthermore, the DSL allows to compose several of these operators together, which is used by the library to apply advanced performance optimisations like loop fusion or software managed caches. Comparing the Fortran OpenACC kernel with the DSL version gives us a speedup of $2.1\times$ for the DSL version. This speedup could also be achieved by hand-tuned optimisation. The DSL prevents the repeated manual effort of tuning the code for multiple architectures. At the same time the DSL allows optimisations to be performed which would otherwise make the code unreadable. More details about this work including code examples on how to use the new backend to GridTools can be found in Osuna (2018).

In addition to GridTools, ESCAPE explored domain specific languages via the CLAW DSL (Clement et al., 2018) for the cloud microphysics dwarf. In particular the use of the Single Column Abstraction (SCA), where physical parameterisations are defined solely in terms of a single horizontal column, enables domain scientists to define governing equations purely in terms of vertical dependencies without needing to account for parallelisation issues. The CLAW DSL then inserts loops over the data-parallel horizontal dimension specific to the hardware architecture and programming model (OpenMP, OpenACC) via source-to-source translation, allowing multiple architectures to be targeted from a single source code. A GPU implementation of the CLOUDSC dwarf generated by automated source translation tools has been used to generate similar performance results to the ones presented by Xiao et al. (2017) on K80 GPUs, using the OpenACC backend of the CLAW DSL.

20 4 Conclusions and outlook

The ESCAPE project has introduced the concept of Weather & Climate dwarfs as fundamental domain-specific building blocks into the weather and climate community. Dwarfs' categorisation of computational and communication patterns has been extremely useful in further breaking down the complexity of weather prediction and climate models, and advancing their adaptation to future hardware. Prototype implementations of the dwarfs have been used to work on code optimisation and benchmarking new architectures. Dwarfs include measures for verifying scientific correctness of the implementations, documentation and input from domain scientists. Our dwarfs are small enough to be fairly easy to learn, and at the same time they represent a significant part of the performance of the whole weather prediction model.

The paper gives an overview of the dwarfs identified in the ESCAPE project, while illustrating in fair detail the optimisation cycle within ESCAPE applied to the spectral transform dwarf that encapsulates historically perhaps the most fundamental algorithmic motif in operational global weather models (Williamson, 2007).

To avoid code duplication we used the data structure and mesh handling framework provided by Atlas. Atlas has been extended in ESCAPE to support limited area grids and DSL with GridTools. The participating weather models and most prototype implementations are based on Fortran and all our optimisations can be incorporated into Fortran code including

the use of CUDA functions on GPUs by calling C functions from Fortran. Selected code optimisations are now routinely used in ECMWF's services. Besides optimisations of the existing code, improved algorithms have been developed which are specifically targeted at improving performance on large scale systems, which include a multigrid preconditioner for the elliptic solver to reduce iteration counts in iterative solvers, a HEVI time-integration scheme with significantly improved stability and alternative finite difference methods on the sphere in the context of reducing global communications across large processor counts, and include alternative solution procedures for spectral transforms at fixed energy cost, with FFTs and spherical harmonics realised on optical processors.

Code optimisations performed in the ESCAPE project targeted Intel CPUs, Intel Xeon Phi processors, NVIDIA GPUs and Optalysys optical processors. In NWP applications, the bottleneck in terms of performance is usually the memory bandwidth between processor and main memory. Having fast interconnects can provide a significant speedup. Having all of the processing units used by the simulation connected with such a fast interconnect is still a challenge. Using accelerators only for a small part of the code reduces their benefit if the associated CPUs are idle while the accelerators perform their computations. Hence, a large part of the code has to be moved to the accelerator (cf. Fuhrer et al., 2018; Schalkwijk et al., 2015) or computations on the CPU are overlapped with computations on the accelerator (cf. Dziekan et al., 2019).

Most of the work done in the ESCAPE project in terms of code optimisation focused on hardware specific optimisations. This makes the optimisations difficult to maintain on upcoming new hardware architectures. As an approach towards performance portability ESCAPE explored a DSL prototype for advection with MPDATA by using the GridTools framework, and a prototype of the cloud microphysics dwarf by using the CLAW DSL. DSLs are a promising tool to enable good performance on multiple architectures while maintaining a single code base. However, designing a domain specific language that is user friendly and at the same time close to hand tuned performance on each architecture is challenging. When assessing the efficacy of diverse numerical methods - e.g., finite-volume vs. spectral spatial discretisation, or Eulerian vs. Lagrangian time stepping - it is important to account for all computational costs. Spectral transforms with semi-implicit, semi-Lagrangian time-integration has always been considered as being poorly suited for large supercomputers due to the high volume of communication. Our work indicates that thanks to much larger time-steps the overall communication cost is not necessarily worse than for other more local methods. Again overlapping different parts of the model such that useful computation can be done while data is communicated is a way forward and needs to be high priority in future research.

When moving towards very high resolution global simulations of $O(1\text{km})$ or less and considering exascale computations on a variety of emerging HPC architectures, there is a continued interest and need in pursuing fundamentally different algorithmic approaches that simply do not communicate beyond a certain halo size while retaining all other favourable properties, such as removing time-step size restrictions as in the semi-Lagrangian advection with semi-implicit time-stepping. Such approaches, as well as the identification of dwarfs in sea-ice and ocean model components of the Earth System, crucial for the performance of coupled applications, are further investigated in ESCAPE-2. These algorithms will utilise DSLs for performance portability.

Code availability. The data structure framework Atlas is available at <https://github.com/ecmwf/atlas> under an Apache License 2.0. The Grid-Tools framework used as a domain specific language approach for MPDATA is available at <https://github.com/GridTools/gridtools> under a BSD-3-Clause license. The CLAW DSL used for the cloud microphysics dwarf is available at <https://github.com/claw-project/claw-compiler> under a BSD-2-Clause license. Model codes developed at ECMWF are the intellectual property of ECMWF and its member states, and therefore the IFS code and the IFS-FVM code are not publicly available. Access to a reduced version of the IFS code may be obtained from ECMWF under an OpenIFS license (see <http://www.ecmwf.int/en/research/projects/openifs> for further information, last access: 28 May 2019). The ALARO and ALADIN codes, along with all their related intellectual property rights, are owned by the members of the ALADIN consortium and are shared with the members of the HIRLAM consortium in the frame of a cooperation agreement. This agreement allows each member of either consortium to license the shared ALADIN-HIRLAM codes to academic institutions of their home country for noncommercial research. Access to the codes of the ALADIN System can be obtained by contacting one of the member institutes or by submitting a request in the contact link below the page of the ALADIN website (<http://www.umr-cnrm.fr/aladin/>, last access: 28 May 2019) and the access will be subject to signing a standardised ALADIN-HIRLAM license agreement. The COSMO-EULAG model, along with all of its related intellectual property rights, is owned by the members of the COSMO consortium under the cooperation agreement. This agreement allows each member of the COSMO consortium to license the COSMO-EULAG code (<http://www.cosmo-model.org/content/consortium/licencing.htm>, last access: 28 May 2019) without fee to academic institutions of their home country for noncommercial research. The code of the GRASS model is intellectual property of Météo-France and is not publicly available. The code of the ESCAPE dwarfs is intellectual property of ECMWF. A license for educational and non-commercial research can be obtained from ECMWF (see <http://www.hpc-escape.eu> for contact details, last access: 28 May 2019). For the GPU optimisation of the spectral transform dwarf we used the PGI compiler version 17.10, CUDA 9.0.176 and OpenMPI 2.1.3. For the GPU optimisation of the MPDATA dwarf we used the PGI compiler version 17.10, CUDA 9.2.88 and OpenMPI 2.1.3. CUDA includes the compiler nvcc used to compile the library function wrappers, the libraries themselves and the profiling tool nvprof which we used for profiling on the GPUs. For the CPU optimisation we used Intel compilers and libraries version 2018.1.163. This includes the compilers icc and ifort and the libraries mkl and mpi. The work on the optical processor Optalysys used the MATLAB software version 2017b.

Author contributions. Andreas Müller performed energy measurements for the spectral transform dwarf, implemented the optimised spectral transform for post-processing, worked on the implementation of the MPDATA dwarf, supported maintenance for all dwarf prototypes and assembled the paper. Willem Deconinck is the main developer of the data structure framework Atlas and contributed the work on improving and extending Atlas. He also was involved in creating, supporting and maintaining all of the dwarf prototypes. Christian Kühnlein and Piotr Smolarkiewicz are the main developers of IFS-FVM and contributed to the corresponding elliptic solver and MPDATA dwarfs. Gianmarco Mengaldo was involved in the creation of the dwarfs and provided improvements to early drafts of the manuscript. Michael Lange and Valentin Clement provided DSL work for the cloud microphysics dwarf. Nils Wedi and Peter Bauer supervised the entire ESCAPE project as coordinators. They were also involved in writing and improving the manuscript. Nils Wedi performed the measurements for the cost profiles in Figure 3. Michail Diamantakis implemented the semi-Lagrangian dwarf prototype. Sarah-Jane Lock supported the work on HEVI time integration methods. Mats Hamrud, Sami Saarinen and George Mozdzyński were involved in creating initial versions of the dwarf prototypes. Sami Saarinen also contributed the code for energy measurements on the Cray XC40 supercomputer. Daniel Thiemert provided project management for the ESCAPE project. Michael Glinton and Pierre Bénard implemented the GRASS model. Fabrice Voitus and Charles Colavolpe worked on HEVI time-stepping methods. Philippe Marguinaud and Yongjun Zheng performed simulations of communication

costs. Joris Van Bever, Daan Degrauwe, Geert Smet and Piet Termonia provided energy and runtime measurements for ALARO, analysed the components of the RMI ensemble prediction system and implemented limited area support in Atlas. Kristian P. Nielsen, Bent H. Sass, Jacob W. Poulsen and Per Berg provided work on the radiation dwarf. Carlos Osuna and Oliver Fuhrer provided work on GridTools and were involved in adding GPU support to Atlas. Michael Baldauf contributed work on HEVI DG. Mike Gillard and Joanna Szmelter worked on improving the preconditioning for the elliptic solver. Enda O'Brien, Alastair McKinstry, Oisín Robinson, Parijat Shukla and Michael Lysaght provided the LAITRI dwarf and created initial CPU optimisations and OpenACC ports for all of the dwarf prototypes. Michał Kulczewski, Milosz Ciznicki, Wojciech Piątek, Sebastian Ciesielski, Marek Błażewicz, Krzysztof Kurowski, Marcin Procyk, Paweł Spychala and Bartosz Bosak worked on performance modelling using the roofline model which allowed a better understanding of the experimental measurements. Zbigniew Piotrowski and Andrzej Wyszogrodzki compared different time-integration methods, supported optimisation of IFS-FVM and provided energy measurements of COSMO-EULAG. Erwan Raffin, Cyril Mazauric, David Guibert, Louis Douriez and Xavier Vigouroux provided the CPU optimisations for the spectral transform and MPDATA dwarf prototypes. Alan Gray and Peter Messmer contributed the GPU optimisations for the spectral transform and MPDATA dwarf prototypes. Alexander J. Macfaden and Nick New contributed the work on using optical processors for the spectral transform.

Acknowledgements. The authors thank three anonymous referees for their thorough reviews and many detailed suggestions that helped to improve the presentation. The ESCAPE project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671627, see also <http://www.hpc-escape.eu/>. This work has been also supported by HPC resources provided by Poznan Supercomputing and Networking Center and corresponding research activities under the MAESTRO grant number DEC-2013/08/A/ST6/00296 in National Science Centre (NCN). An award of computer time was provided by the INCITE program. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE office of Science User Facility supported under contract DE-AC05-00OR22725.

References

- Asanović, K., Bodik, R., Catanzaro, B. C., Gebis, J. J., Husbands, P., Keutzer, K., Patterson, D. A., Plishker, W. L., Shalf, J., Williams, S. W., and Yelick, K. A.: The landscape of parallel computing research: a view from Berkeley, Tech. Rep. UCB/EECS-2006-183, EECS Department, University of California, Berkeley, <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf>, 2006.
- 5 Asanović, K., Wawrzynek, J., Wessel, D., Yelick, K., Bodik, R., Demmel, J., Keaveny, T., Keutzer, K., Kubiawicz, J., Morgan, N., Patterson, D., and Sen, K.: A view of the parallel computing landscape, *Comm. ACM*, 52, 56, <https://doi.org/10.1145/1562764.1562783>, 2009.
- Bauer, P., Thorpe, A., and Brunet, G.: The quiet revolution of numerical weather prediction, *Nature*, 525, 47–55, <https://doi.org/10.1038/nature14956>, 2015.
- Bénard, P. and Gintion, M.: Circumventing the pole problem of reduced lat-lon grids with local schemes. Part I: analysis and model formulation, *Q. J. Royal Meteorol. Soc.*, 145, 1377–1391, <https://doi.org/10.1002/qj.3509>, 2019.
- 10 Clement, V., Ferrachat, S., Fuhrer, O., Lapillonne, X., Osuna, C. E., Pincus, R., Rood, J., and Sawyer, W.: The CLAW DSL, in: Proceedings of the Platform for Advanced Scientific Computing Conference - PASC '18, ACM Press, <https://doi.org/10.1145/3218176.3218226>, 2018.
- Colavolpe, C., Voitus, F., and Bénard, P.: RK-IMEX HEVI schemes for fully compressible atmospheric models with advection: analyses and numerical testing, *Q. J. Royal Meteorol. Soc.*, 143, 1336–1350, <https://doi.org/10.1002/qj.3008>, 2017.
- 15 Colella, P.: Defining software requirements for scientific computing, DARPA HPCS Presentation, 2004.
- Deconinck, W.: Development of Atlas, a flexible data structure framework, Tech. rep., ECMWF, <http://arxiv.org/abs/1908.06091>, 2017a.
- Deconinck, W.: Public release of Atlas under an open source license, which is accelerator enabled and has improved interoperability features, Tech. rep., ECMWF, <http://arxiv.org/abs/1908.07038>, 2017b.
- Deconinck, W., Bauer, P., Diamantakis, M., Hamrud, M., Kühnlein, C., Maciel, P., Mengaldo, G., Quintino, T., Raoult, B., Smolarkiewicz, P. K., and Wedi, N. P.: Atlas - a library for numerical weather prediction and climate modelling, *Comput. Phys. Commun.*, 220, 188–204, <https://doi.org/10.1016/j.cpc.2017.07.006>, 2017.
- 20 Douriez, L., Gray, A., Guibert, D., Messmer, P., and Raffin, E.: Performance report and optimized implementations of weather & climate dwarfs on multi-node systems, Tech. rep., ECMWF, <http://arxiv.org/abs/1908.06097>, 2018.
- Dziekan, P., Waruszewski, M., and Pawlowska, H.: University of Warsaw Lagrangian Cloud Model (UWLCM) 1.0: a modern large-eddy simulation tool for warm cloud modeling with Lagrangian microphysics, *Geosci. Model Dev.*, pp. 2587–2606, <https://doi.org/10.5194/gmd-12-2587-2019>, 2019.
- 25 Feng, W., Lin, H., Scogland, T., and Zhang, J.: OpenCL and the 13 dwarfs: a work in progress, in: Proceedings of the third joint WOSP/SIPEW International Conference on Performance Engineering - ICPE '12, ACM Press, <https://doi.org/10.1145/2188286.2188341>, 2012.
- 30 Flamm, K.: Measuring Moore’s law: evidence from price, cost, and quality indexes, National Biuro of Economic Research, Working Paper No. 24553, April 2018, <https://doi.org/10.3386/w24553>, 2018.
- Fuhrer, O., Chadha, T., Hoefler, T., Kwasniewski, G., Lapillonne, X., Leutwyler, D., Lüthi, D., Osuna, C., Schär, C., Schulthess, T. C., and Vogt, H.: Near-global climate simulation at 1km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0, *Geosci. Model Dev.*, 11, 1665–1681, <https://doi.org/10.5194/gmd-11-1665-2018>, 2018.
- 35 Gintion, M. R. and Bénard, P.: Circumventing the pole problem of reduced lat-lon grids with local schemes. Part II: validation experiments, *Q. J. Royal Meteorol. Soc.*, 145, 1392–1405, <https://doi.org/10.1002/qj.3495>, 2019.

- Johnston, B. and Milthorpe, J.: Dwarfs on accelerators: enhancing OpenCL benchmarking for heterogeneous computing architectures, in: Proceedings of the 47th International Conference on Parallel Processing Companion - ICPP '18, ACM Press, <https://doi.org/10.1145/3229710.3229729>, 2018.
- 5 Kaltofen, E. L.: The “seven dwarfs” of symbolic computation, in: Numerical and Symbolic Scientific Computing, edited by Langer, U. and Paule, P., pp. 95–104, Springer, https://doi.org/10.1007/978-3-7091-0794-2_5, 2011.
- Katzav, J. and Parker, W. S.: The future of climate modeling, *Clim. Change*, 132, 475–487, <https://doi.org/10.1007/s10584-015-1435-x>, 2015.
- Krommydas, K., Feng, W., Antonopoulos, C. D., and Bellas, N.: OpenDwarfs: characterization of dwarf-based benchmarks on fixed and reconfigurable architectures, *J. Signal Process. Syst.*, 85, 373–392, <https://doi.org/10.1007/s11265-015-1051-z>, 2015.
- 10 Kühnlein, C., Deconinck, W., Klein, R., Malardel, S., Piotrowski, Z. P., Smolarkiewicz, P. K., Szmelter, J., and Wedi, N. P.: FVM 1.0: a nonhydrostatic finite-volume dynamical core formulation for IFS, *Geosci. Model Dev.*, 12, 651–676, <https://doi.org/10.5194/gmd-12-651-2019>, 2019.
- Lawrence, B. N., Rezny, M., Budich, R., Bauer, P., Behrens, J., Carter, M., Deconinck, W., Ford, R., Maynard, C., Mullerworth, S., Osuna, C., Porter, A., Serradell, K., Valcke, S., Wedi, N., and Wilson, S.: Crossing the chasm: how to develop weather and climate models for next generation computers?, *Geosci. Model Dev.*, 11, 1799–1821, <https://doi.org/10.5194/gmd-11-1799-2018>, 2018.
- 15 Macfaden, A. J., Gordon, G. S. D., and Wilkinson, T. D.: An optical Fourier transform coprocessor with direct phase determination, *Sci. Rep.*, 7, <https://doi.org/10.1038/s41598-017-13733-1>, 2017.
- Mazauric, C., Raffin, E., and Guibert, D.: Recommendations and specifications for data scope analysis tools, Tech. rep., ECMWF, <http://arxiv.org/abs/1908.06095>, 2017a.
- Mazauric, C., Raffin, E., Vigouroux, X., Guibert, D., Macfaden, A., Poulsen, J., Berg, P., Gray, A., and Messmer, P.: Performance report and optimized implementation of weather & climate dwarfs on GPU, MIC and Optalysys optical processor, Tech. rep., ECMWF, <http://arxiv.org/abs/1908.06096>, 2017b.
- 20 Mengaldo, G.: Batch 1: definition of several weather & climate dwarfs, Tech. rep., ECMWF, <http://arxiv.org/abs/1908.06089>, 2016.
- Mengaldo, G., Wyszogrodzki, A., Diamantakis, M., Lock, S.-J., Giraldo, F., and Wedi, N. P.: Current and emerging time-integration strategies in global numerical weather and climate prediction, *Arch. Comput. Methods Eng.*, pp. 1–22, <https://doi.org/10.1007/s11831-018-9261-8>, 25 2018.
- Messer, O. E. B., D’Azevedo, E., Hill, J., Joubert, W., Berrill, M., and Zimmer, C.: MiniApps derived from production HPC applications using multiple programming models, *Int. J. High Perform. Comput. Appl.*, 32, 582–593, <https://doi.org/10.1177/1094342016668241>, 2016.
- Michalakes, J., Govett, M., Benson, R., Black, T., Juang, H., Reinecke, A., and Skamarock, B.: AVEC Report: NGGPS Level-1 Benchmarks and software evaluation, Tech. rep., NOAA, Boulder, US, <https://repository.library.noaa.gov/view/noaa/18654>, 2015.
- 30 Mozdzynski, G., Hamrud, M., and Wedi, N. P.: A partitioned global address space implementation of the European Centre for Medium Range Weather Forecasts Integrated Forecasting System, *Int. J. High Perform. Comput. Appl.*, 29, 261–273, <https://doi.org/10.1177/1094342015576773>, 2015.
- Müller, A., Kopera, M. A., Marras, S., Wilcox, L. C., Isaac, T., and Giraldo, F. X.: Strong scaling for numerical weather prediction at petascale with the atmospheric model NUMA, *Int. J. High Perform. Comput. Appl.*, 2, 411–426, <https://doi.org/10.1177/1094342018763966>, 2018.
- 35 Müller, A., Gillard, M., Nielsen, K. P., and Piotrowski, Z.: Batch 2: definition of novel weather & climate dwarfs, Tech. rep., ECMWF, <http://arxiv.org/abs/1908.07040>, 2017.

- Neumann, P., Düben, P., Adamidis, P., Bauer, P., Brück, M., Kornbluh, L., Klocke, D., Stevens, B., Wedi, N., and Biercamp, J.: Assessing the scales in numerical weather and climate predictions: will exascale be the rescue?, *Philos. Trans. Royal Soc. A: Mathematical, Physical and Engineering Sciences*, 377, <https://doi.org/10.1098/rsta.2018.0148>, 2019.
- Osuna, C.: Report on the performance portability demonstrated for the relevant weather & climate dwarfs, Tech. rep., ECMWF, <http://arxiv.org/abs/1908.06094>, 2018.
- 5 Palmer, T.: Climate forecasting: Build high-resolution global climate models, *Nature*, 515, 338–339, <https://doi.org/10.1038/515338a>, 2014.
- Phillips, S. C., Engen, V., and Papay, J.: Snow white clouds and the seven dwarfs, in: 2011 IEEE Third International Conference on Cloud Computing Technology and Science, IEEE, <https://doi.org/10.1109/cloudcom.2011.114>, 2011.
- Poulsen, J. W. and Berg, P.: Tuning the implementation of the radiation scheme ACRANEB2, Tech. rep., DMI report 17-22, http://www.dmi.dk/fileadmin/user_upload/Rapporter/TR/2017/SR17-22.pdf, 2017.
- 10 Robinson, O., McKinsty, A., and Lysaght, M.: Optimization of IFS subroutine LAITRI on Intel Knights Landing, Tech. rep., PRACE White Papers, <https://doi.org/10.5281/zenodo.832025>, 2016.
- Schalkwijk, J., Jonker, H. J. J., Siebesma, A. P., and Meijgaard, E. V.: Weather forecasting using GPU-based Large-Eddy Simulations, *Bull. Am. Meteorol. Soc.*, 96, 715–723, <https://doi.org/10.1175/bams-d-14-00114.1>, 2015.
- 15 Schulthess, T. C., Bauer, P., Wedi, N., Fuhrer, O., Hoefler, T., and Schär, C.: Reflecting on the goal and baseline for exascale computing: a roadmap based on weather and climate simulations, *Comput. Sci. Eng.*, 21, 30–41, <https://doi.org/10.1109/mcse.2018.2888788>, 2019.
- Shukla, J., Palmer, T. N., Hagedorn, R., Hoskins, B., Kinter, J., Marotzke, J., Miller, M., and Slingo, J.: Toward a new generation of world climate research and computing facilities, *Bull. Am. Meteorol. Soc.*, 91, 1407–1412, <https://doi.org/10.1175/2010bams2900.1>, 2010.
- Van Bever, J., McFaden, A., Piotrowski, Z., and Degrauwe, D.: Report on energy-efficiency evaluation of several NWP model configurations, Tech. rep., ECMWF, <http://arxiv.org/abs/1908.06115>, 2018.
- 20 Wallemacq, P. and House, R.: Economic losses, poverty and disasters 1998-2017, Tech. rep., <https://www.unisdr.org/we/inform/publications/61119>, 2018.
- Wedi, N. P., Hamrud, M., and Mozdzyński, G.: A fast spherical harmonics transform for global NWP and climate models, *Mon. Weather Rev.*, 141, 3450–3461, <https://doi.org/10.1175/mwr-d-13-00016.1>, 2013.
- 25 Wedi, N. P., Bauer, P., Deconinck, W., Diamantakis, M., Hamrud, M., Kühnlein, C., Malardel, S., Mogensen, K., Mozdzyński, G., and Smolarkiewicz, P.: The modelling infrastructure of the Integrated Forecasting System: recent advances and future challenges, Tech. Rep. 760, Eur. Cent. For Medium-Range Weather Forecasts, Reading, UK, <https://doi.org/10.21957/thtpwp67e>, 2015.
- Wehner, M. F., Oliker, L., Shalf, J., Donofrio, D., Drummond, L. A., Heikes, R., Kamil, S., Kono, C., Miller, N., Miura, H., Mohiyuddin, M., Randall, D., and Yang, W.-S.: Hardware/software co-design of global cloud system resolving models, *J. Adv. Model. Earth Sys.*, 3, <https://doi.org/10.1029/2011ms000073>, 2011.
- 30 Williamson, D. L.: The evolution of dynamical cores for global atmospheric models, *J. Meteor. Soc. Japan*, 85B, 241–269, <https://doi.org/10.2151/jmsj.85b.241>, 2007.
- World Economic Forum: The 2019 global risks report, <https://www.weforum.org/reports/the-global-risks-report-2019>, 2019.
- Wu, J., Wyckoff, P., and Panda, D.: High performance implementation of MPI derived datatype communication over InfiniBand, in: 18th International Parallel and Distributed Processing Symposium, 2004. Proceedings., IEEE, <https://doi.org/10.1109/ipdps.2004.1302917>, 2004.
- 35 Xiao, H., Diamantakis, M., and Saarinen, S.: An OpenACC GPU adaptation of the IFS cloud microphysics scheme, ECMWF Tech. Memo. No. 805, <https://doi.org/10.21957/g9mjjlgeq>, 2017.

Zheng, Y. and Marguinaud, P.: Simulation of the performance and scalability of MPI communications of atmospheric models running on exascale supercomputers, *Geosci. Model Dev.*, 11, 3409–3426, <https://doi.org/10.5194/gmd-11-3409-2018>, 2018.