Geoscientific Model Development

Open Access

Experiences with distributed computing for meteorological applications: grid computing and cloud computing

F. Oesterle[1,*]**, S. Ostermann**[2]**, R. Prodan**[2]**, and G. J. Mayr**[1]

[1]Institute of Atmospheric and Cryospheric Science, University of Innsbruck, Innrain 52, 6020 Innsbruck, Austria
[2]Institute of Computer Science, University of Innsbruck, Innsbruck, Austria
[*]née Schüller

Correspondence to: F. Oesterle (felix.oesterle@uibk.ac.at)

Abstract. Experiences with three practical meteorological applications with different characteristics are used to highlight the core computer science aspects and applicability of distributed computing to meteorology. Through presenting cloud and grid computing this paper shows use case scenarios fitting a wide range of meteorological applications from operational to research studies. The paper concludes that distributed computing complements and extends existing high performance computing concepts and allows for simple, powerful and cost-effective access to computing capacity.

1 Introduction

Meteorology has an ever growing need for substantial amounts of computing power, be it for sophisticated numerical models of the atmosphere itself, modelling systems and workflows, coupled ocean and atmospheric models or the accompanying activities visualisation or dissemination. In addition to the increased need for computing power, more data are being produced, transferred and stored, which increases the problem. Consequently, concepts and methods to supply the compute power and data handling capacity also have to evolve.

Until the beginning of this century, high performance clusters, local consortia and/or buying cycles on commercial clusters were the main methods to acquire sufficient capacity. Starting in the mid-1990s, the concept of grid computing, in which geographical and institutional boundaries only play a minor role, became a powerful tool for scientists. Foster and Kesselman (2003) published the first and most cited def-

inition of the grid: "A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities". In the following years the definition changed to viewing the grid not as a computing paradigm, but as an infrastructure that brings together different resources in order to provide computing support for various applications, emphasising the social aspect (Foster and Kesselman, 2004; Bote-Lorenzo et al., 2004). Grid initiatives can be classified as *compute grids*, i.e. solely concentrated on raw computing power, or *data grids* concentrating on storage/exchange of data.

Many initiatives in the atmospheric sciences have utilised compute grids. One of the first climatological applications to use a compute grid is the Fast Ocean Atmospheric Model (FOAM) (Nefedova et al., 2006). They performed ensemble simulations of a coupled climate model on the Teragrid, a US-based grid project sponsored by the National Science Foundation. More recently, Fernández-Quiruelas et al. (2011) provided an example with the Community Atmospheric Model (CAM) for a climatological sensitivity study investigating the connection of sea surface temperature and precipitation in the El Niño area. Todorova et al. (2010) presents three Bulgarian projects investigating air pollution and climate change impacts. WRF4SG utilises grid computing with the Weather Research and Forecast Model (WRF) (Blanco et al., 2013) for various applications in weather forecasting and extreme weather case studies. TIGGE, the THORPEX Interactive Grand Global Ensemble, partly uses grid computing to generate and share atmospheric data between various partner (Bougeault et al., 2010). The Earth

system grid ESGF (Earth System Grid Federation) is a US–European data grid project concentrating on storage and dissemination of climate simulation data (Williams et al., 2009).

Cloud computing is slightly newer than grid computing. Resources are also pooled, but this time usually within one organisational unit, mostly within commercial companies. Similar to grids, applications range from services based on demand to simply cutting ongoing costs or determining expected capacity needs.

The most important characteristics of clouds are condensed into one of the most recent definitions by Mell and Grance (2011): "Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources, (e.g. networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction". Further definitions can be found in Hamdaqa and Tahvildari (2012), Vaquero and Rodero-Merino (2008) or Vaquero and Rodero-Merino (2008). One of the few papers to apply cloud technology to meteorological research is Evangelinos and Hill (2008), who conducted a feasibility study for cloud computing with a coupled atmosphere–ocean model.

In this paper, we discuss advantages and disadvantages of both infrastructures for atmospheric research, show the supporting software ASKALON, and present three examples of meteorological applications, which we have developed for different kinds of distributed computing: projects *MeteoAG* and *MeteoAG2* for a compute grid, and *RainCloud* for cloud computing. We look at issues and benefits mainly from our perspective as users of distributed computing. Please note we describe our experiences but do not show a direct, quantitative comparison, as we did not have the resources to run experiments on both infrastructures with identical applications.

2 Aspects of distributed computing in meteorology

2.1 Grid and cloud computing

Our experiences in grid computing come from the projects MeteoAG and MeteoAG2 within the national effort Austrian-Grid (AGrid), including partners and supercomputer centres from all over Austria (Volkert, 2004). AGrid phase 1 started in 2005 and concentrated on research of basic grid technology and application. Phase 2, started in 2008, continued to build on research of phase 1 and additionally tried to make AGrid self-sustaining. The research aim of this project was not to develop conventional parallel applications that can be executed on individual grid machines, but rather to unleash the power of the grid for single distributed program runs. To simplify this task, all grid sites are required to run a similar Linux operating system. At the height of the project AGrid consisted of nine clusters distributed over five locations in Austria including various smaller sites with ad hoc desktop

PC networks. The progress of the project, its challenges and solutions were documented in several technical reports and other publications (Bosa and Schreiner, 2009).

For cloud computing, plenty of providers offer services, e.g. Rackspace or Google Compute Engine. Our cloud computing project *RainCloud* uses Amazon Web Services (AWS), simply because it is the most well known and widely used. AWS offers different services for computing, different levels of data storage and data transfer, as well as tools for monitoring and planning. The services most interesting for meteorological computing purposes are Amazon Elastic Compute cloud (EC2) for computing and Amazon Simple Storage Service (S3) for data storage. For computing, so-called *instances*, (i.e. virtual computers) are defined according to their compute power relative to a reference CPU, available memory, storage and network performance.

Figure 1 shows the basic structure of cloud computing on the right side and AGrid as a grid example on the left side. In both cases an additional layer, so-called middleware, is applied between the compute resources and the end user. The Middleware layer handles all necessary scheduling, transfer of data and set up of cloud nodes. Our middleware is ASKALON (Ostermann et al., 2008), which is described in more detail in Sect. 2.2.

In the following sections, we list advantages and disadvantages of grid and cloud concepts, which affected our research most (see Table 1 for a brief overview). Criteria are extracted from literature, most notably Foster et al. (2008) containing a general comparison with all vital issues, Mell and Grance (2011), Hamdaqa and Tahvildari (2012) and Foster and Kesselman (2003). The discussed issue of security of sensitive and valuable data did not apply to our research and operational setting. However, for big and advanced operational weather forecasting this might be an issue due to its monetary value. Because the hardware and network is completely out of the end user's control, possible security breaches are harder or even impossible to detect. If security is a concern, detailed discussions can be found in Cody et al. (2008) for grid computing, and Catteddu (2010) and Feng et al. (2011) for cloud computing.

2.1.1 Advantages and disadvantages grid

+ *Handle massive amounts of data*. The full atmospheric model in MeteoAG generated large amounts of data. Through grid tools like *gridftp* (Allcock et al., 2002) we were able to efficiently transfer and store all simulation data.

+ *Access to high performance computing (HPC) which suits parallel applications, (e.g. Message Passing Interface, MPI)*. The model used in MeteoAG, as many other meteorological models, is a massive parallel application parallelised with MPI. On single systems they run efficiently; however, across different HPC clusters latencies
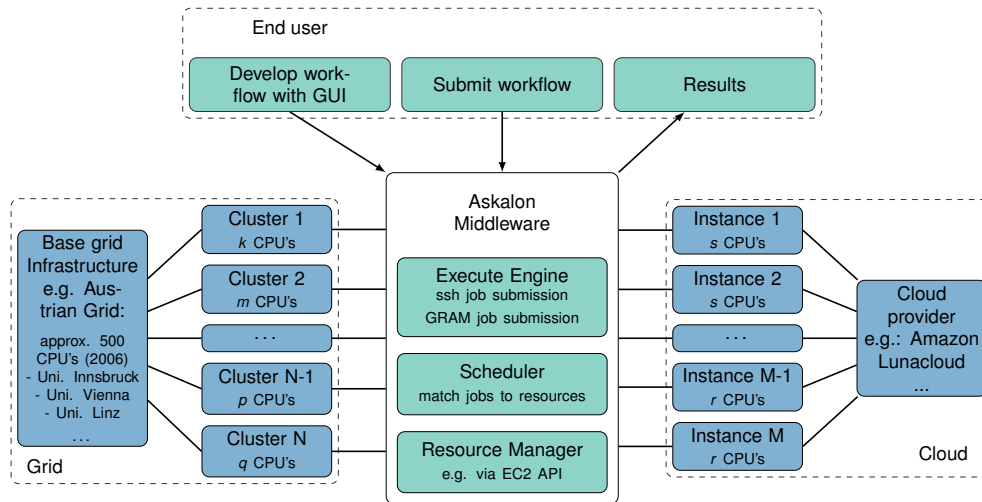
Figure 1. Schematic set-up of our computing environment for grid (left) and cloud (right) computing. End users interact with the ASKALON middleware via a Graphical User Interface (GUI). The number of CPUs per cluster provided by the base grid varies, whereas the instance types of cloud providers can be chosen. Execute engine, scheduler and resource manager interact to effectively use the available resources and react to changes in the provided computing infrastructure.

Table 1. Overview of advantages/disadvantages of grids and clouds affecting our applications most. For a detailed discussion see Sect. 2.

grids	clouds
+ Massive amounts of data	+ Cost
+ Access to parallel computing enabled high performance computing (HPC)	+ Full control of software set-up
	+ Simple on-demand access
− Inhomogeneous hardware architectures	− Slow data transfer
− Complicated set-up and inflexible handling	− Not suitable for MPI computing
− Special compilation of source code needed	

become too high. A middleware can leverage the advantage of access to multiple machines and run applications on suitable machines and appropriately distribute parts of workflows in parallel.

– *Different hardware architectures.* During tests in MeteoAG, we discovered problems due to different hardware architectures (Schüller et al., 2007). We tested different systems with exactly the same set-up and software and got consistently different results. In our case this affected our complex full model, but not our simple model. The exact cause is unclear, but most likely a combination of programming, the libraries and set-up down to the hardware level.

– *Difficult to set up and maintain as well as inflexible handling.* For us, the process of getting necessary updates, patches or special libraries needed in meteorology onto all grid sites was complex and lengthy or sometimes even impossible due to operating system limitations.

– *Special compilation of source code.* To get the most out of the available resources, the executables in MeteoAG

needed to be compiled for each architecture, with possible side effects. Even in a tightly managed project like AGrid, we had to supply three different executables for the meteorological model, with changes only during compilation, not in the model code itself.

Other typical characteristics are not as important for us. The "limited amount of resources" never influenced us as they were always vast enough to not hinder our models. The "need to bring your own hardware/connections" is also a small hindrance, since this is usually negotiable or the grid project might have different levels of partnership.

2.1.2 Advantages and disadvantages cloud computing

+ *Cost.* Costs can easily be determined and planned. More about costs can be found in Sect. 4.

+ *Full control of software environment, including operating system (OS) with root access.* This proved to be one of the biggest advantages for our workflows. It is easy to install software, special libraries or modify any component of the system. Cloud providers usually offer most

standard operating systems as images/Amazon Machine Image (AMI), but tuned images can also be saved permanently and made publicly available (with additional storage costs).

+ *Simple on-demand self-service.* For applications with varying requirements for compute resources or with repeated but short needs for compute power, simple on-demand self-service is an important characteristic. As long as funds are available the required amount of compute power can be purchased. Our workflow was never forced to wait for instances to be available. Usually our standard on-demand Linux instances were up and running within 5–10 s (Amazon's documentation states a maximum of 10 min).

− *Slow data transfer and hardly any support for MPI computing.* Data transfer to and from cloud instances is slow as well as higher network latency between the instances. Only a subset of instance types are suitable for MPI computing. This limitation makes cloud computing unsuitable for large-scale complex atmospheric models.

"Missing information about underlying hardware" has no impact on our workflow, as we are not trying to optimise a single model execution. "No common standard between clouds" and the possibility of "a cloud provider going out of business" is also unimportant for us. Our software relies on common protocols like ssh and adaptation to a new cloud provider could be done easily by adjusting the script requesting the instances.

2.2 Middleware ASKALON

To make it as simple as possible for a (meteorological) end user to use distributed computing resources, we make use of a so-called middleware system. ASKALON, an existing middleware from the Distributed and Parallel Systems group in Innsbruck, provides integrated environments to support the development and execution of scientific workflows on dynamic grid and cloud environments (Ostermann et al., 2008).

To account for the heterogeneity and the loosely coupled nature of resources from grid and cloud providers, ASKALON has adopted a workflow paradigm (Taylor et al., 2007) based on loosely coupled coordination of atomic activities. Distributed applications are split in reasonably small execution parts, which can be executed in parallel on distributed systems, allowing the runtime system to optimise resource usage, file transfers, load balancing, reliability, scalability and handle failed parts. To overcome problems resulting from unexpected job crashes and network interruptions, ASKALON is able to handle most of the common failures. Jobs and file transfers are resubmitted on failure and jobs might also be rescheduled to a different resource if transfers or jobs failed more than 5 times on a resource (Plankensteiner et al., 2009a). These features still exist in the cloud version

but play a less important role as resources showed to be more reliable in the cloud case.

Figure 1 shows the design of the ASKALON system. Workflows can be generated in a scientist-friendly Graphical User Interface (GUI) and submitted for execution by a service. This allows for long lasting workflows without the need for the user to be online throughout the whole execution period.

Three main components handle the execution of the workflow:

1. *Scheduler.* Activities are mapped to physical (or virtualised) resources for their execution with the end user deciding which pool of resources are used. A wide set of scheduling algorithms is available, e.g. Heterogeneous Earliest Finish Time (HEFT) (Zhao and Sakellariou, 2003) or Dynamical Critical Path - Clouds (DCP-C) (Ostermann and Prodan, 2012). HEFT, for example, takes as input tasks, a set of resources, the times to execute each task on each resource and the times to communicate results between each job on each pair of resources. Each task is assigned a priority and then distributed onto the resources accordingly. For the best possible scheduling, a training phase is needed to get a function that relates the problem size to the processing time. Advanced techniques in prediction and machine learning are used to achieve this goal (Nadeem and Fahringer, 2009; Nadeem et al., 2007).

2. *Resource manager.* cloud resources are known to "scale by credit card" and theoretically an infinite amount of resources is available. The resource manager has the task to provision the right amount of resources at the right moment to allow the execute engine to run the workflow as the scheduler decided. Cost constraints must be strictly adhered to as budgets are in practice limited. More on costs can be found in Sect. 4.

3. *Execute engine.* Submission of jobs and transfer of data to the compute resources is done with a suitable protocol, e.g. ssh or Globus resource allocation manager (GRAM) in a Globus environment.

– *System reliability.* An important feature which is distributed over several components of ASKALON is the capability to handle faults in distributed systems. Resources or network connections might fail any time and mechanisms as described in Plankensteiner et al. (2009a) are integrated in the execution engine Qin et al. (2007) allowing workflows to finish even when parts of the system fail.

3 Applications in meteorology

In the following subsections, we detail the three applications we developed for usage with distributed computing.

All projects investigate orographic precipitation over complex terrain. The most important distributed computing characteristics of the projects are shown in Table 3.

3.1 MeteoAG

MeteoAG started as part of the AGrid computing initiative. Using ASKALON we created a workflow to run a full numerical atmospheric model and visualisation on a grid infrastructure (Schüller, 2008; Schüller et al., 2007; Schüller and Qin, 2006). The model is the non-hydrostatic Regional Atmospheric Modeling System (RAMS; version 6), a fully MPI parallelised Fortran-based code (Cotton et al., 2003). The National Center for Atmospheric Research (NCAR) graphics library is used for visualisation. Due to all AGrid sites running a similar Linux OS, no special code adaptations to grid computing were needed.

We simulated real cases as well as idealised test cases in the AGrid environment. Most often these were parameter studies testing sensitivities to certain input parameters with many slightly different runs. The investigated area in the realistic simulations covered Europe and a target area over western Austria. Several nested domains are used with a horizontal resolution of the innermost domain of 500 m and 60 vertical levels (approx. 7.5 million grid points). Figure 2 shows the workflow deployed to the AGrid. Starting with many simulations with a shorter simulation time, it was then decided which runs to extend further. Only runs where heavy precipitation occurs above a certain threshold were chosen. Post-processing done on the compute grid includes extraction of variables and preliminary visualisation, but the main visualisation was done on a local machine.

The workflow characteristics relevant for distributed computing are a few (20–50) model instances but highly CPU intensive as well as lots of interprocess communications. Results of this workflow require a substantial amount of data transfer between the different grid sites and the end user (O(200 Gb)).

Upon investigation of our first runs it was necessary to provide different executables for specific architectures (32 bit, 64 bit, 64 bit Intel) to get optimum speed. We ran into a problem while executing the full model on different architectures. Using the exact same static executable with the same input parameters and set-up led to consistently different results across different clusters (Schüller et al., 2007). For real case simulations, these errors are negligible compared to errors in the model itself. But for idealised simulations, e.g. investigation of turbulence with an atmosphere initially at rest, where tiny perturbations play a major role, this might lead to serious problems. We were not able to determine the cause of these differences. It seems to be a problem of the complex code of the full model and its interaction with the underlying libraries. While we can only speculate on the exact cause, we strongly advise using a simple and quick test such as simu-

lating an atmosphere at rest or linear orographic precipitation to test for such differences.

3.2 MeteoAG2

MeteoAG2 is the continuation of MeteoAG and also part of AGrid (Plankensteiner et al., 2009b). Based on the experience from the MeteoAG experiments, we hypothesise that it would be much more effective to deploy an application consisting of serial CPU jobs. ASKALON is optimised for submission of single core parts of a workflow, which avoids internal parallelism and communication of activities and allows for the best control over the execution within ASKALON. Thus MeteoAG2 uses a simpler meteorological model, the linear model (LM) of orographic precipitation (Smith and Barstad, 2004). The model computes only very simple linear equations of orographic precipitation, is not parallelised, and has short runtime, O(10 s), even with high resolutions (500 m) over large domains. LM is written in Fortran. ASKALON is again used for workflow execution and Matlab routines for visualisation.

With this workflow, rainfall over the Alps was investigated by taking input from the European Centre for Medium-Range Weather Forecasts (ECMWF) model, splitting the Alps into subdomains (see Fig. 3a) and running the model within each subdomain with variations in the input parameters. The last step combines the results from all subdomains and visualises them. Using grid computing allowed us to run many O(50 000) simulations in a relatively short amount of time O(h). This compares to about 50 typical, albeit a lot more complex runs in current operational meteorological set-ups.

The workflow deployed to the grid (Fig. 3b) is simple with only two main activities: preparing all the input parameters for all subdomains and then the parallel execution of all runs. One of the drawbacks of MeteoAG2 is the very strict set-up that was necessary due to the state of ASKALON at that time, e.g. no robust if-construct yet, and the direct use of model executables without wrappers. The workflow could not easily be changed to suit different research needs, e.g. change to different input parameters for LM or to using a different model.

3.3 RainCloud

In switching to cloud computing, RainCloud uses an extended version of the same simple model of orographic precipitation as MeteoAG2. The main extension to LM is the ability to simulate different layers, while still retaining its fast execution time (Barstad and Schüller, 2011). The software stack again includes ASKALON, the Fortran-based LM, python scripts and Matplotlib for visualisation.

The inclusion of if-constructs in ASKALON and a different approach to the scripting of activities, (e.g. wrapping the model executables in python scripts and calling these) allows RainCloud to be used in different set-ups. We are now able

Table 2. Prices and specifications for Amazon EC2 on-demand instances mentioned in this paper, running Linux OS in region *EU-west* as of November 2014. m1.xlarge, m1.medium and m2.4xlarge are previous generations which were used in our experiments. Storage is included in the instance, additional storage is available for purchase. One elastic compute unit (ECU) provides the equivalent CPU capacity of a 1.0–1.2 GHz 2007 Opteron or 2007 Xeon processor.

Instance family	Instance type	vCPU	ECU	Memory (GiB)	Storage (GB)	Cost USD h^{-1}
General purpose	m3.medium	1	3	3.75	1×4 SSD	0.077
	m3.2xlarge	8	26	30	2×80 SSD	0.616
Compute optim.	c3.8xlarge	32	108	60	2×320 SSD	1.912
Storage optim.	hs1.8xlarge	16	35	117	24×2048	4.900
Micro instances	t1.micro	1	Var	0.615	EBS only	0.014
General purpose	m1.xlarge	4	8	15	4×420	0.520
	m1.medium	1	2	3.75	1×410	0.130
Memory optim.	m2.4xlarge	8	26	68.4	2×840	1.840
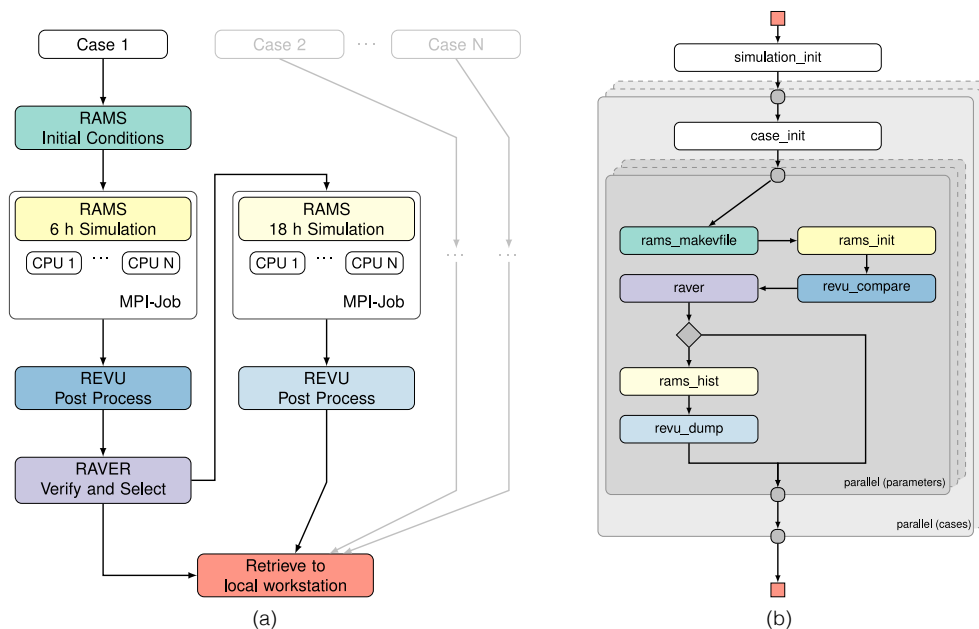
Figure 2. Workflow of MeteoAG using the Regional Atmospheric Modelling System (RAMS) and supporting software REVU (extracts variables) and RAVER (analyses variables). Each case represents a different weather event. **(a)** Meteorological representation with indication which activities are parallelised using Message Passing Interface (MPI). **(b)** Workflow representation of the activities as used by ASKALON middleware. In addition to the different cases, selected variables are varied within each case. Same colours between the subfigures.

to run the workflow in three flavours without any changes: idealised, semi-idealised and realistic simulations as well as different settings, operational and research. Figure 4b depicts the workflow run on cloud computing. Only the first two activities, *PrepareLM* and *LinearModel* have to be run, the others are optional. This workflow fits a lot of meteorological applications as it has the following building blocks:

– preparation of the simulations (*PrepareLM*);

– execution of a meteorological model (*LinearModel*);

– post-processing of each individual run, e.g. for producing derived variables (*PostProcessSingle*);

– post-processing of all runs (*PostprocessFinal*).

All activities are wrapped in Python scripts. As long as the input and output between these activities are named the same, everything within the activity can be changed. We use archives for transfer between the activities, again allowing different files to be packed into these archives.

The operational set-up produces spatially detailed, daily probabilistic precipitation forecasts for the Avalanche Service Tyrol (Lawinenwarndienst Tirol) to help forecast avalanche danger. Figure 4a shows the schematic of our operational workflow. Starting with data from the ECMWF, we forecast and visualise precipitation probabilities over Tyrol with a spatial resolution of 500 m. Additionally, research type
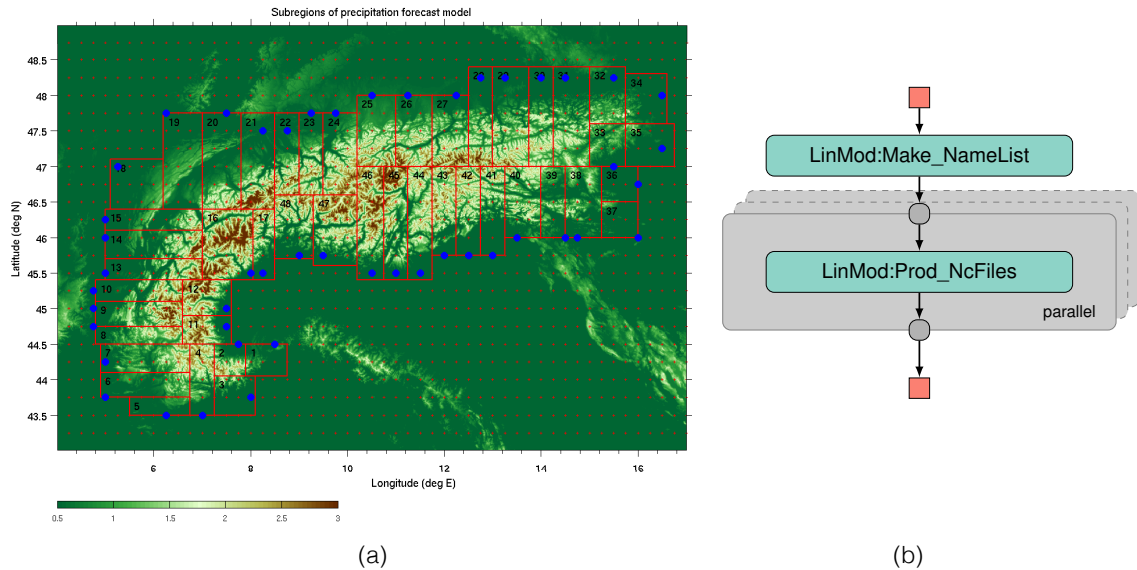
Figure 3. Set-up and workflow of MeteoAG2 using the linear model (LM) of orographic precipitation. **(a)** Grid set-up of experiments in MeteoAG2 with dots representing grid points of the European Center of Medium Range Weather Forecast (ECMWF) used to drive the LM. Topography height in kilometres a.m.s.l. **(b)** Workflow representation of the activities as used by ASKALON. Activity MakeNML prepares all input sequentially. ProdNCfile is the main activity with the linear model run in parallel on the grid. Panel **(a)** courtesy of Plankensteiner et al. (2009b).
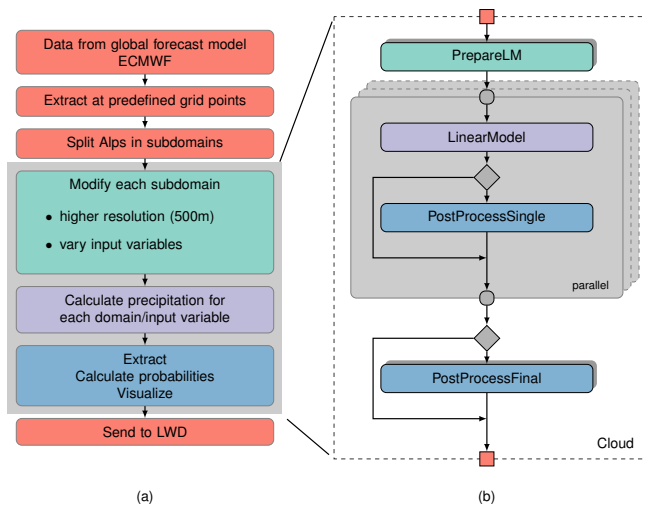
Figure 4. Workflow of RainClouds operational setting for the Avalanche Warning Service Tyrol (LWD) using the double layer linear model (LM) of orographic precipitation. Input data from the European Center for Medium Range Weather Forecast (ECMWF). **(a)** Meteorological flow chart with parts not executed on cloud (in red). **(b)** Workflow with activities as used by ASKALON. Same colours between subfigures.

experiments are used to test, explore and run experiments with new developments in LM through parameter studies.

Our workflow set-ups vary substantially in required computation power as well as data size. The operational job is run daily during winter, whereas research types are run in bursts. Data usage *within* the cloud can be substantial O(500 Gb) with all flavours, but with big differences of data transfer from the cloud back to the local machine. Operational results are small, of the order of O(100 Mb), while research results can amount to O(100 Gb), influencing the overall runtime and costs due to the additional data transfer time.

4 Costs, performance and usage scenarios

4.1 Costs

To define the exact costs for a *dedicated server* system or the participation in a *grid* initiative is not trivial, and often even unknown to the provider; we contacted several of them, but due to complicated budgeting methodologies the final costs are not obvious. Greenberg and Hamilton (2008) discussed costs for operating a server environment for data services from a provider perspective, including servers, infrastructure, power requirements and networking. However, the authors did not include the cost of human resources for, e.g., system administration. Patel and Shah (2005) included human resources and establish a cost model for set-up and maintenance of a data centre. Grids may have different and negotiable levels of access and participation, with varying associated costs to the user. Some initiatives, e.g. PRACE (Guest et al., 2012), offer free access to grid resources after a proposal/review process.

Table 3. Overview of our projects and their workflow characteristics.

Project	MeteoAG	MeteoAG2	RainCloud
Type	grid	grid	cloud
Meteorological model	RAMS (Regional Atmospheric Modeling System)	single layer linear model of orographic precipitation	double layer linear model of orographic precipitation
Model type	complex full numerical model parallelised with Message Passing Interface (MPI)	simplified model	double layer simplified model
Parallel runs	20–50	approx. 50 000	> 5000 operational, > 10 000 research
Runtime	several days	several hours	1–2 h operational/ <1 h research
Data transfer	O(200GB)	O(1GB)	O(MB) – O(1GB)
Workflow flexibility	strict	strict	flexible
Applications	parameter studies, case studies	downscaling	parameter studies, downscaling, probabilistic forecasts, model testing
Intent	research	research	operational, research
Frequency	on demand	on demand	operational: daily; research: on demand
Programming	shell scripts, Fortran, NCAR Graphics, MPI	shell scripts, Fortran, Matlab	python, Fortran

Cloud computing on the other hand offers simpler and transparent costs. Pricing varies depending on the provider, capability of a resource and also on the geographical region. Prices (as of November 2014) of AWS on-demand compute instances for Linux OS can be found in Table 2 and range from USD 0.014 up to $\sim 5\,h^{-1}$ (region Ireland). Cheaper instance pricing is available through spot instances where one bids on spare resources. These resources might get cancelled if demand rises, but are a valid option for interruption-tolerant workflows or for developing a workflow.

Figure 5 shows the difference between spot and on-demand pricing for 25 test runs of our operational workflow (circle and x; right y axis). All runs use 32 cores but a different number of instances, i.e. only one c3.8xlarge (32 cores) instance, but 32 m1.medium (1 core) instances. Runtime only includes the actual workflow, not the spin-up needed to prepare the instances. It usually takes 5–10 s for an instance to become available and another 2–5 min to set up the system and install necessary libraries and software. Spot and on-demand only differ in the pricing scheme not in the computational resources themselves. With spot pricing we achieved savings between 65 and 90 %, however with an additional start-up latency of 2–3 min (compared to 5–10 s).

To give an idea, a very simplified cost comparison can be done with the purchasing costs of dedicated hardware, excluding costs for system administration, cooling or power. The operational part of RainCloud runs on 32 cores for approximately 3 h per day for 6 months of the year, i.e. 550 h per year.

- A dedicated 32 core server with 64 GB RAM costs around USD 5500 (various brands, excluding tax, Austria, November 2014).

- A comparable on-demand AWS instance (c3.x8large; 32 cores, 60 GB RAM) could run for USD $\sim 2800\,h$ at $1.91\,h^{-1}$ pricing.

Assuming no instance price variance, our operational workflow could be run on AWS for approximately 5 years, the usual depreciation time for hardware. This suggests that AWS is the cheaper alternative for RainCloud, since hardware is only one part of the total cost of ownership of a dedicated system.

4.2 Performance

For our operational RainCloud workflow, Fig. 5 shows the effect of different instance types on the runtime. First, a clear
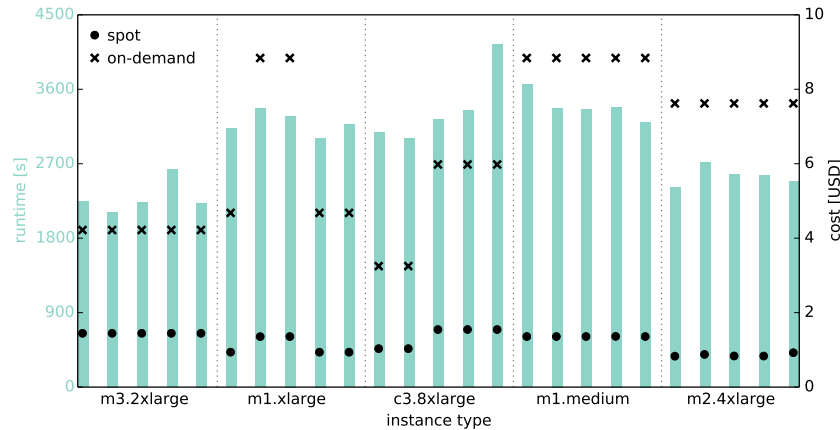
Figure 5. Bars show overall runtime of one operational run on various Amazon EC2 instance types, each with a total of 32 cores (left y axis). Each bar represents one workflow invocation with the corresponding instance type. Dots show costs for on-demand instances (x) and spot instances (circle; right y axis). Only the execution part is shown, spin-up time, i.e. preparation and installation (2–5 min) is not included. See Table 3 for exact specifications. All experiments were run during March 2014 with the exact same set-up.

difference between the instance types is evident, with the longest running taking nearly twice as long as the shortest one. Second, even within one instance type, runtime varies by 10–20 percent. Serial execution on a 1 core desktop PC takes about 12 h, i.e. a speedup of ~ 18 (a runtime of ~ 0.66 h as seen in Fig. 5). Based on these experiments our daily operational workflow uses four m3.2xlarge instances.

To put this into relation, Schüller et al. (2007) showed a speedup for MeteoAG of multiple cores vs. 1 core for a short running test set-up of ~ 5, with higher speedups possible for a full complex workflow run. For MeteoAG2, Plankensteiner et al. (2009b) showed a speedup of ~ 120 when executing that workflow on several grid machines compared to the execution on a single desktop PC. However, as these are different workflows, no comparison between the type of computing resources can be made from these performance measures.

4.3 Usage scenarios

Different usage scenarios are commonly found in meteorology. For choosing the right type of computing system, several issues need to be taken into account. Only above a certain workflow scale is it worth the effort to move away from a local machine. grids usually have a steep learning curve, clouds offer simple (web) interfaces and local clusters are somewhere in the middle. To make the most out of cloud computing (and to some extent out of grid computing), it is best to have a workflow which can be split into small, independent components.

In an "operational scenario with frequent invocations", either clouds and grids might be suitable depending on the amount of data transferred and the complexity of the model. Time critical data dissemination of forecast products can be sped up with (data) grids. "Operational scenarios with in-

frequent invocations" might benefit from using grid or even cloud computing, avoiding the need for a local cluster. Examples are recalculation/reanalysis of seasonal/climate simulations or updating of model output statistics (MOS) equations. One important consideration for operational workflows is the scheduling latency, i.e. the time between submitting a job and its actual execution. Berger et al. (2009) and Lingrand and Montagnat (2009) show median latencies of 100 s for Enabling Grids for E-Science in Europe (EGEE) grid, but with frequent outliers upwards to 30 min and more (Rain-Cloud 10–120 s).

For a "research scenario with bursts of high activity with many small tasks", cloud computing fits perfectly. The costs are fully controllable and only little set-up is required. Examples of such use cases include parameter studies with simple models or computation of MOS. If a lot of data transfer is needed, grid computing is the better alternative. "Research applications with big, long running, data intensive simulations" such as high-resolution complex models are best run on grids or local clusters.

5 Conclusions

We successfully deployed meteorological applications on distributed computing infrastructure of both grids and clouds. Our meteorological applications range from a complex atmospheric limited-area model to a simplified model of orographic precipitation. Adhering to some limitations/considerations, distributed computing can cater to both.

A consideration to be taken into account for both concepts is security. With grids, it is relatively easy to determine users and potential access to data as all resources and locations are known. With clouds, this is nearly impossible/impractical to do this and potential breaches are hard to detect.

If the grid is seen as an agglomeration of individual supercomputers, complex parallelised models are simple to deploy and efficient to use in a research setting. The compute power is usually substantially larger than what a single institution could afford. However, in an operational setting the immediate availability of resources might not be a given. This is an issue that needs to be addressed in advance. For data storage and transfer, e.g. dissemination of forecasts, grids are a powerful tool.

Taking grid as a structure, workflows involving MPI are not simple to exploit. As with clouds, it is much more effective to deploy an application consisting of serial jobs with as little interprocess communication as possible.

Heterogeneity of the underlying hardware cannot be ignored for grid computing as quality tests showed (Schüller et al., 2007). Differences arising solely based on the used hardware might influence very sensitive applications. However, this is application-specific and needs to be tested for each set-up.

The set-up and access to cloud infrastructure is a lot simpler and involves less effort than participation in a grid project. Grids require hardware and more complex software to access, whereas access to clouds is usually kept as simple as possible.

Cloud (commercial) computing is very effective and cost saving tool for certain meteorological applications. Individual projects with high-burst needs or an operational setting with a simple model are two examples. Elasticity, i.e. access to a larger scale of resources, is one of the biggest advantages of clouds. Undetermined or volatile needs can be easily catered for. One option is to use clouds to baseline workflow requirements and then build and move to a correctly sized in-house cluster/set-up based on this prototyping.

Disadvantages of clouds include above-mentioned security issues, but one of the biggest problems for meteorological applications is data transfer. Transfer to and from the cloud and within the cloud infrastructure is considerably slower than for a dedicated cluster set-up or grids. Recently new instance types for massively parallel computing have been emerging, (e.g. Amazon), but high computation applications with only modest data needs are best suited for most clouds.

Private clouds remove some of the disadvantages of public clouds, security and data transfer are the most notable ones. However, using private clouds also removes the advantage of not needing hardware and system administration. We used a small private cloud to develop our workflow before going full scale on Amazon AWS with our operational set-up.

In a meteorological research setting with specialised software, clouds offer a flexible system with full control over operating system, installed software and libraries. Grids on the other hand are managed on individual grid sites and are more strict and less flexible. The same is true for customer service. clouds offer one contact for all problems and offer (paid) premium support as opposed to having to contact each system administration for every grid site.

In conclusion, both concepts are an alternative or a supplement to self-hosted high-performance computing infrastructure. We have laid out guidelines with which to decide whether one's own application is suitable to either or both alternatives.

Edited by: S. Unterstrasser

References

Allcock, B., Bester, J., Bresnahan, J., Chervenak, A. L., Foster, I. T., Kesselman, C., Meder, S., Nefedova, V., Quesnel, D., and Tuecke, S.: Data management and transfer in high-performance computational Grid environments, Parallel Comput., 28, 749–771, 2002.

Barstad, I. and Schüller, F.: An extension of Smith's linear theory of orographic precipitation: introduction of vertical layers, J. Atmos. Sci., 68, 2695–2709, 2011.

Berger, M., Zangerl, T., and Fahringer, T.: Analysis of overhead and waiting time in the EGEE production Grid, in: Proceedings of the Cracow Grid Workshop, 2008, 287–294, 2009.

Berriman, G. B., Deelman, E., Juve, G., Rynge, M., and Vöckler, J.-S.: The application of cloud computing to scientific workflows: a study of cost and performance, Philos. T. R. Soc. A., 371, 20120066, doi:10.1098/rsta.2012.0066, 2013.

Blanco, C., Cofino, A. S., and Fernandez-Quiruelas, V.: WRF4SG: a scientific gateway for climate experiment workflows, Geophys. Res. Abstr., EGU2013-11535, EGU General Assembly 2013, Vienna, Austria, 2013.

Bosa, K. and Schreiner, W.: A supercomputing API for the Grid, in: Proceedings of 3rd Austrian Grid Symposium 2009, edited by: Volkert, J., Fahringer, T., Kranzlmuller, D., Kobler, R., and Schreiner, W., Austrian Grid, Austrian Computer Society (OCG), 38–52, available at: http://www.austriangrid.at/index.php?id=symposium (last access: 17 June 2015), 2009.

Bote-Lorenzo, M. L., Dimitriadis, Y. A., and Sanchez, E. G. A.: Grid Characteristics and Uses: A Grid Definition, Vol. 2970, Springer, Berlin, Heidelberg, 2004.

Bougeault, P., Toth, Z., Bishop, C., Brown, B., Burridge, D., Chen, D. H., Ebert, B., Fuentes, M., Hamill, T. M., Mylne, K., Nicolau, J., Paccagnella, T., Park, Y.-Y., Parsons, D., Raoult, B., Schuster, D., Dias, P. S., Swinbank, R., Takeuchi, Y., Tennant, W., Wilson, L., and Worley, S.: The THORPEX interactive grand global ensemble, B. Am. Meteorol. Soc., 91, 1059–1072, 2010.

Catteddu, D.: Cloud computing: benefits, risks and recommendations for information security, in: Web Application Security SE – 9, edited by: Serrão, C., Aguilera Díaz, V., and Cerullo, F.,

Vol. 72 of Communications in Computer and Information Science, Springer, Berlin, Heidelberg, p. 17, 2010.

Cody, E., Sharman, R., Rao, R. H., and Upadhyaya, S.: Security in grid computing: a review and synthesis, Decis. Support Syst., 44, 749–764, 2008.

Cotton, W. R., Pielke Sr., R. A., Walko, R. L., Liston, G. E., Tremback, C. J., Jiang, H., McAnelly, R. L., Harrington, J. Y., Nicholls, M. E., Carrio, G. G., and McFadden, J. P.: RAMS 2001: Current status and future directions, Meteorol. Atmos. Phys., 82, 5–29, 2003.

Deelman, E., Singh, G., Livny, M., Berriman, B., and Good, J.: The cost of doing science on the cloud: the montage example, in: Proceedings of the 2008 ACM/IEEE conference on Supercomputing, p. 50, 2008.

Evangelinos, C. and Hill, C.: Cloud computing for parallel scientific HPC applications: feasibility of running coupled atmosphere-ocean climate models on Amazon's EC2, Ratio, 2, 2–34, 2008.

Feng, D.-G., Zhang, M., Zhang, Y., and Xu, Z.: Study on cloud computing security, J. Softw., 22, 71–83, 2011.

Fernández-Quiruelas, V., Fernández, J., Cofiño, A., Fita, L., and Gutiérrez, J.: Benefits and requirements of grid computing for climate applications. An example with the community atmospheric model, Environ. Modell. Softw., 26, 1057–1069, 2011.

Foster, I. and Kesselman, C.: The Grid 2: Blueprint for a New Computing Infrastructure, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.

Foster, I., Zhao, Y., Raicu, I., and Lu, S.: Cloud Computing and Grid Computing 360-Degree Compared, 2008 Grid Computing Environments Workshop, 1–10, 2008.

Foster, I. T. and Kesselman, C.: The Grid: Blueprint for a New Computing Infrastructure, 2nd Edn., Morgan Kaufmann, Amsterdam, 2004.

Greenberg, A. and Hamilton, J.: The cost of a cloud: research problems in data center networks, ACM SIGCOMM Computer Communication Review, 39, 68–73, 2008.

Guest, M., Aloisio, G., and Kenway, R.: The scientific case for HPC in Europe 2012–2020, tech. report, PRACE, 2012.

Hamdaqa, M. and Tahvildari, L.: Cloud computing uncovered: a research landscape, Adv. Comput., 86, 41–85, 2012.

Lingrand, D. and Montagnat, J.: Analyzing the EGEE production grid workload: application to jobs submission optimization, Lect. Notes Comput. Sc., 5798, 37–58, 2009.

Mell, P. and Grance, T.: The NIST definition of cloud computing recommendations of the National Institute of Standards and Technology, Special Publication 800–145, NIST, Gaithersburg, available at: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf (last access: 9 February 2015), 2011.

Nadeem, F. and Fahringer, T.: Predicting the execution time of grid workflow applications through local learning, in: High Performance Computing Networking, Proceedings of the Conference on Storage and Analysis, 1, 1–12, doi:10.1145/1654059.1654093, 2009.

Nadeem, F., Prodan, R., and Fahringer, T.: Optimizing Performance of Automatic Training Phase for Application Performance Prediction in the Grid, in: High Performance Computing and Communications, Third International Conference, HPCC 2007, Houston, USA, 26–28 September 2007, 309–321, 2007.

Nefedova, V., Jacob, R., Foster, I., Liu, Z., Liu, Y., Deelman, E., Mehta, G., Su, M.-H., and Vahi, K.: Automating climate science: large ensemble simulations on the TeraGrid with the GriPhyN virtual data system, e-science, 0, 32, doi:10.1109/E-SCIENCE.2006.261116, 2006.

Ostermann, S. and Prodan, R.: Impact of variable priced cloud resources on scientific workflow scheduling, in: Euro-Par 2012 Parallel Processing, edited by: Kaklamanis, C., Papatheodorou, T., and Spirakis, P., Vol. 7484 of Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 350–362, 2012.

Ostermann, S., Plankensteiner, K., Prodan, R., Fahringer, T., and Iosup, A.: Workflow monitoring and analysis tool for ASKALON, in: Grid and Services Evolution, Barcelona, Spain, 73–86, 2008.

Patel, C. D. and Shah, A. J.: Cost Model for Planning, Development and Operation of a Data Center, Technical Report HP Laboratories Palo Alto, HPL-2005-107(R.1), 09 June 2005.

Plankensteiner, K., Prodan, R., and Fahringer, T.: A new fault tolerance heuristic for scientific workflows in highly distributed environments based on resubmission impact, in: eScience'09, 313–320, 2009a.

Plankensteiner, K., Vergeiner, J., Prodan, R., Mayr, G., and Fahringer, T.: Porting LinMod to predict precipitation in the Alps using ASKALON on the Austrian Grid, in: 3rd Austrian Grid Symposium, edited by: Volkert, J., Fahringer, T., Kranzlmüller, D., Kobler, R., and Schreiner, W., Vol. 269, Austrian Computer Society, 103–114, 2009b.

Qin, J., Wieczorek, M., Plankensteiner, K., and Fahringer, T.: Towards a light-weight workflow engine in the ASKALON Grid environment, in: Proceedings of the CoreGRID Symposium, Springer-Verlag, Rennes, France, 2007.

Schüller, F.: Grid Computing in Meteorology: Grid Computing with – and Standard Test Cases for – A Meteorological Limited Area Model, VDM, Saarbrücken, Germany, 2008.

Schüller, F. and Qin, J.: Towards a workflow model for meteorol[o]gical simulations on the Austrian Grid, Austrian Computer Society, 210, 179–190, 2006.

Schüller, F., Qin, J., Nadeem, F., Prodan, R., Fahringer, T., and Mayr, G.: Performance, Scalability and Quality of the Meteorological Grid Workflow MeteoAG, Austrian Computer Society, 221, 155–165, 2007.

Smith, R. B. and Barstad, I.: A linear theory of orographic precipitation, J. Atmos. Sci., 61, 1377–1391, 2004.

Taylor, I. J., Deelman, E., Gannon, D., and Shields, M.: Workflows for e-Science, Springer-Verlag, London Limited, 2007.

Todorova, A., Syrakov, D., Gadjhev, G., Georgiev, G., Ganev, K. G., Prodanova, M., Miloshev, N., Spiridonov, V., Bogatchev, A., and Slavov, K.: Grid computing for atmospheric composition studies in Bulgaria, Earth Sci. Inf., 3, 259–282, 2010.

Vaquero, L. and Rodero-Merino, L.: A break in the clouds: towards a cloud definition, ACM SIGCOMM Computer Communication Review, 39, 50–55, 2008.

Volkert, J.: The Austrian Grid Initiative – high level extensions to Grid middleware, in: PVM/MPI, edited by: Kranzlmüller, D., Kacsuk, P., and Dongarra, J. J., Vol. 3241 of Lecture Notes in Computer Science, Springer, p. 5, 2004.

Williams, D. N., Drach, R., Ananthakrishnan, R., Foster, I. T., Fraser, D., Siebenlist, F., Bernholdt, D. E., Chen, M., Schwidder, J., Bharathi, S., Chervenak, a. L., Schuler, R., Su, M., Brown, D., Cinquini, L., Fox, P., Garcia, J., Middleton, D. E.,

Strand, W. G., Wilhelmi, N., Hankin, S., Schweitzer, R., Jones, P., Shoshani, A., and Sim, A.: The Earth System Grid: enabling access to multimodel climate simulation data, B. Am. Meteorol. Soc., 90, 195–205, 2009.

Zhao, H. and Sakellariou, R.: An experimental investigation into the rank function of the heterogeneous earliest finish time scheduling algorithm, in: Euro-Par Conference, 189–194, 2003.