

Optical Networks for Grid and Cloud Computing Applications

Chris Develder, *Member IEEE*, Marc De Leenheer, *Member IEEE*, Bart Dhoedt, *Member IEEE*, Mario Pickavet, *Senior Member IEEE*, Didier Colle, *Member IEEE*, Filip De Turck, *Member IEEE*, Piet Demeester, *Fellow IEEE*

Abstract — The evolution towards grid and cloud computing as observed for over a decennium illustrates the crucial role played by (optical) networks in supporting today’s applications. In this paper, we start from an overview of the challenging applications in both academic (further referred to as *scientific*), enterprise (*business*) and non-professional user (*consumer*) domains. They pose novel challenges, calling for efficient interworking of IT resources, for both processing and storage, as well as the network that interconnects them and provides access to their users. We outline those novel applications’ requirements, including sheer performance attributes (which will determine the quality as perceived by end users of the cloud applications), as well as the ability to adapt to changing demands (usually referred to as elasticity) and possible failures (i.e. resilience). In outlining the foundational concepts that provide the building blocks for grid/cloud solutions that meet the stringent application requirements we highlight, a prominent role is played by optical networking. The pieces of the solution studied in this respect span the optical transport layer as well as mechanisms located in higher layers (e.g. anycast routing, virtualization), and their interworking (e.g. through appropriate control plane extensions and middleware). Based on this study, we conclude by identifying challenges and research opportunities that can enable future-proof optical cloud systems (e.g. pushing the virtualization paradigms to optical networks).

Index Terms—Optical Networks, Cloud Computing, Grid Computing, Virtualization

I. INTRODUCTION

DURING the evolution of computing, characterized by Moore’s law, the role and scale of networks has been incessantly growing, as illustrated in Fig. 1. Just as the

Manuscript received July 1, 2011. Work described in this paper was partly funded by the European Commission through the 7th ICT-Framework Programme (FP7) projects Geysers (grant FP7-ICT-248657), Ofelia (grant FP7-ICT-258365) and SPARC (grant FP7-ICT-258457), as well as by Ghent University through GOA Optical Grids (grant 01G01506).

The authors are with IBCN, Dept. of Information Technology (INTEC), Ghent University – IBBT, Ghent, Belgium (e-mail: {firstname.lastname}@intec.ugent.be).

M. De Leenheer and C. Develder (in part) are supported as post-doctoral fellows of the Research Foundation – Flanders (FWO–VI). M. De Leenheer is currently a visiting post-doctoral researcher at the University of California - Davis, and is supported by a Fellowship of the Belgian American Educational Foundation (BAEF).

pioneering Mark I and ENIAC systems in the 1940s, the mainframe computers that were produced starting in the late 1950s through the 1960s were essentially stand-alone systems. Over time, access to such powerful machines was provided through terminals that remotely connected to a mainframe computer (such as IBM 360, VAX, PDP etc.) which performed the core processing for the multi-user group it thus served. Similarly, also personal computers conceived in the late 1970s (enabled by the microprocessor designed in the late 1960s), initially were prevalently used as stand-alone systems that could seemingly serve all the needs of their respective users. Gradually, and starting mainly in academic installations at universities, those PCs were tied up with each other thanks to Metcalfe’s invention of Ethernet in the 1970s. His visionary realization that this networking significantly increases the value of the whole set of end points is now incontestable, looking at the current Internet and its crucial role in almost any business.

In the 1990s, the increasing networking capabilities, and in particular the high capacities (needed for e.g. massive data volumes generated in particle physics experiments) and low latencies (e.g. enabling interactive visualization of that data) achievable using optical technologies led to the inception of a so-called computational grid [1]: in analogy with the power grid, this would allow users to obtain computing power on demand, irrespective of their location. As with previous evolutions (cf. PC, LAN, Internet), the grid concept was first implemented in academic circles. There, the idea rose to offer access to e.g. powerful computing facilities to remote users. The development of this concept gave birth to many worldwide grid infrastructure initiatives (cf. Open Science Grid [2], the Enabling Grid for E-Science (EGEE) [3] now continued as the European Grid Infrastructure (EGI), and TeraGrid [4]).

Inspired by the success of the grid paradigm in scientific circles, the cloud computing ideas arose in the 2000s, building on the seminal idea of “computation provided as a public utility” (as suggested back in 1961 by John McCarthy). The evolution to a wide-spread adoption of such “utility computing” could be a logical next step where functionality is gradually pushed further into the network (see Fig. 1). That evolution initially manifested in client-server based architectures in local networks, and continued with remotely hosted parts in web-based solutions. Whereas initially the network interactions were mainly of a point-to-point nature,

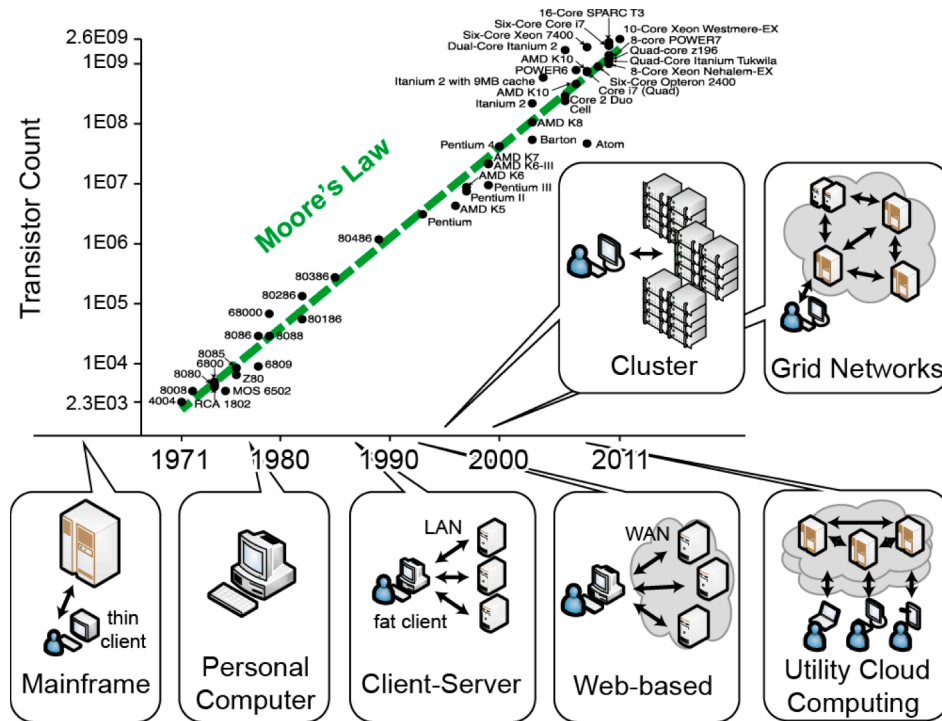


Fig. 1. Computing power has been growing exponentially, which is reflected in Moore's Law describing the doubling of the number of transistors on a single integrated circuit approximately every two years. To make efficient use of these resources, distributed computing paradigms have been devised, eventually enabling offloading of applications to the cloud and realizing the idea of utility computing. Clearly, this evolution builds heavily on advances in networking.

they now rely on interworking of many distributed components as observed in grid computing, and present-day cloud-connected systems.

In this paper, we outline the characteristics of cloud and grid applications, their requirements and the fundamental concepts that lay at the basis of solutions to address them. We particularly advocate an evolution towards “optical clouds”, stressing the important role that especially optical networking technology can play in realizing next generation cloud solutions. Note that we mainly will discuss wide-area networks, and will not delve into e.g. optical interconnects within data centers [5]. Since clearly optical networks are part of a bigger grid/cloud eco-system (see Fig. 2), we will also elaborate some non-network-specific concepts.

The remainder of this paper is structured as follows: we start in Section II by sketching the applications that gave rise to grids and clouds, and highlight their differences as well as similarities. In Section III we derive the requirements for technologies to support them. In Section IV, we explain the characteristics of grids and clouds, referring to the typical applications they serve. The foundational components of solutions that (optical) grids/clouds build on, to meet the challenging requirements of the novel applications they serve, are identified in Section V. Before concluding, Section VI gives an overview of remaining challenges as well as a selection of research projects that address (some of) them.

II. NOVEL APPLICATIONS

Easy access to powerful, and often distributed, hardware and software resources has been of key importance for science

(e.g. the groundbreaking EGEE project [3] aimed at efficiently distributing large amounts of experimental data obtained in the Large Hadron Collider (LHC) at CERN to thousands of physicists) as well as business (e.g. the hosting facilities offered by Amazon.com). As the resource offerings continue to grow and are also coming in reach of non-professional users (referred to as “consumers” hereafter), we are witnessing the emergence of a wide variety of novel applications also targeting this audience. A common observation in all these domains is the increasing reliance on networking (cf. Fig. 1), and hence a distributed infrastructure as sketched in Fig. 2. In this section, an overview is presented of applications envisaged to run on large distributed infrastructures, assisting the requirement elicitation process discussed subsequently in Section III.

A. Scientific applications

The need for resources exceeding the capacities offered at a single research center has been a major motivation for constructing high-performance distributed infrastructures. These e-Science applications still present an important fraction of the workload, and are focused on either scientific computing or on distributing/collecting data.

Scientific computing: Scientific simulations (e.g. weather forecasting, computational fluid dynamics) are the core applications in this category. Originally, these simulations were organized as a bag-of-tasks, i.e. loosely coupled tasks running in parallel on the system without much need for communication or interdependencies (a typical example being Monte Carlo simulations [6]). However, a clear evolution

towards coupled tasks is being perceived: the work is organized as a set of interdependent tasks (dependencies are specified through a directed acyclic graph), resulting in a workflow to be executed on the infrastructure. In addition, tasks running in parallel can require communication to notify intermediate results.

Data-centric computing: Distributing or collecting experimental data is the main focus in this category. Data obtained through expensive equipment (e.g. high-energy physics, astronomy) is distributed to large scientific communities for further analysis and interpretation. Systems supporting this should of course offer guarantees for the availability of this expensive data as well as high bandwidth network connectivity to transport these data sets efficiently. A second, more recent, application type consists of data collection in sensor networks. Typically, data originating from a possibly large collection of sensors (typically a few 10s up to a few 1000s) is aggregated and stored persistently. Resilient and efficient data access is of paramount importance for this type of applications. First generation sensor networks were typically dedicated to collecting parameter values at moderate frequencies (e.g. environmental parameters each 15 minutes), leading to only moderate network and computing requirements. However, nowadays sensors can send multimedia data across the network requiring intensive back-end processing for analysis purposes [7] (e.g. video stream analysis to detect abnormal behavior in surveillance applications).

B. Business applications

In terms of applications in typical business settings, we observe roughly the following classes: transactional systems, collaborative tools, multimedia applications, and data mining.

Transactional systems: Businesses in various sectors have relied on transactional systems (e.g. web shop applications, stock exchange transaction management) for a few decades. This transaction processing is now typically outsourced, to allow businesses to focus on their core business and to realize cost reductions due to the economies of scale. Note that in this domain, machine-to-machine (M2M) applications in e.g. stock exchange applications are foreseen to lead to interactions on ever shorter timescales (see Fig. 3), implying that even optical networks’ capabilities and computation speed will be pushed to the edge [8].

Efficient collaboration: To facilitate efficient interaction between people at different geographic locations, virtual meeting systems and collaborative frameworks can come to the rescue. Such virtual meeting systems offer interaction facilities beyond traditional video conferencing applications, including processing streams from different participants to render a single image where all participants (or their avatars) appear to meet in a virtual room. Additional features include automatic detection of points of interest (e.g. automatic zooming to a presenter), gesture and mood detection as well as stream adaptation to client side capabilities or the network. Collaborative frameworks allow working on shared objects (e.g. text document or presentation) or engineering designs

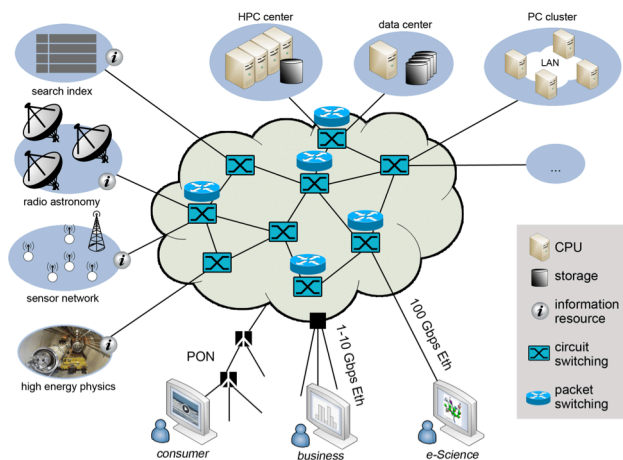


Fig. 2. Different types of users tap into a resource rich distributed system. This system offers infrastructural resources (storage, computation and connectivity) as well as information (e.g. experimental data). High bandwidth connectivity, offered by different optical technologies, is needed to facilitate advanced application scenarios. As illustrated with some typical examples, the access technologies may vary depending on consumer-oriented, business or e-Science applications.

[9]. To produce different viewpoints on-demand of this shared object again requires dedicated remote rendering facilities, and depending on the type of collaboration, also individual manipulations on this shared object can require substantial computational resources.

Multimedia: The multimedia sector is heavily relying on distributed infrastructures, both for storing/retrieving high-quality multimedia content (e.g. digital storage of ingested raw material for further editing, consultation of multimedia archive to identify a sequence satisfying certain topical criteria) as well as for processing this material (e.g. transcoding raw material to a suitable stream format). These systems are obviously very demanding in terms of offering robust storage capacities and network performance (in view of the data rates needed in this context) [10].

Data mining: As electronic data is becoming available in almost every application domain, the need to understand the underlying structure of this data and to identify trends in large data volumes is being recognized (e.g. to build usage profiles of expensive infrastructures, understand customer behavior, produce meaningful recommendations, build a search index for a large data volume etc.).

C. Consumer applications

During the 1990s, the emergence of the world wide web has seduced the non-professional users to make use of distributed infrastructures. Since then, a whole range of applications has emerged, offering very attractive features beyond simple web browsing.

Personal content: Manipulating (storage, retrieval and processing) personal content has become mainstream with the advent of Web 2.0 applications (e.g. Flickr, YouTube). These systems are extremely demanding in term of storage resources, especially in view of the number of users. A second challenge concerns the unpredictability of this user community: certain content items can suddenly become extremely popular

(leading to flash crowds), putting high and unexpected load on parts of the system, including the network.

Gaming: Gaming oriented applications have stringent requirements on processing capacities and latency (action games typically requiring less than 70ms between the action and the visual effect). Serious games (i.e. applications using gaming technology in more professional settings) share these characteristics, and depending on the application at hand latency requirements are stringent (e.g. flight simulation immersive applications) or relaxed (e.g. virtual tourism applications allowing to visit a city remotely). Efforts are ongoing to offer gaming applications through a centralized facility (e.g. GamesAtLarge [11]).

In this same gaming context, massive multi-player on-line role playing games (MMORPG) can serve over 1 million users simultaneously (distributed over different game instances, called “realms”). These games present the same virtual world to end-users, who are represented by avatars. Users navigate and perform strategic actions in this virtual world, getting more powerful capabilities as the game evolves. A particular challenge here consists of serving these large audiences, also characterized by a certain level of unpredictable behavior (e.g. users moving collectively to a certain part of the virtual world, causing excessive load on one particular server).

Augmented Reality: In augmented reality applications, real world information is supplemented with synthetically generated content. Typically, a user is equipped with a head mounted display and a camera, capturing the scene. This scene is analyzed in real time, and relevant meta-data is displayed in the display. This scene analysis should happen in real-time, and requires considerable computational resources.

Interactive TV: As a last representative in this application category, we mention advanced interactive television settings. In such settings, a 360° scene is captured in high quality, and the end-user can select the viewpoint from which he prefers to inspect this scene (e.g. view a goal in a football match from the reverse angle) [12]. Again, considerable processing is needed (to combine the information of different cameras in order to present the desired view) in combination with substantial network bandwidths (typically, UHD camera systems are used to capture the scene, and at least HD quality is required for streaming the content to the consumer).

III. REQUIREMENTS

Based on a scan of aforementioned application scenarios, we have identified the following requirements for (optical) grid/cloud solutions, as summarized qualitatively in Table I:

- **On-demand set-up:** As indicated before, a crucial characteristic of cloud applications is that they rely on on-demand instantiation of the required network and IT resources. Thus, this calls for dynamic set-up mechanisms to quickly allocate the necessary storage, computational and network resources and make the necessary configurations. Note that set-up and tear-down of resource allocations should be user-friendly (i.e. highly automatized) and responsive (see higher). This

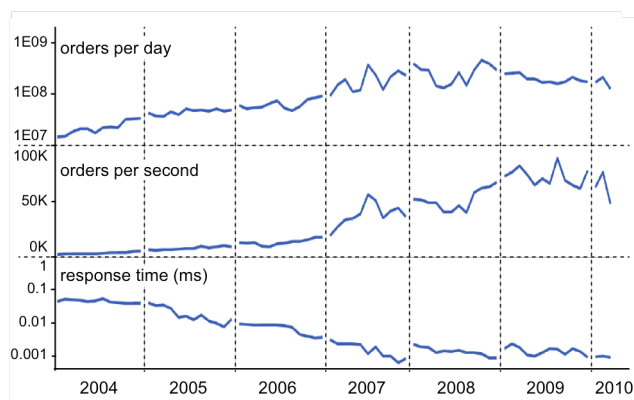


Fig. 3. The financial industry poses extreme challenges, striving towards zero latency. To eliminate any delay, real-time market data is distributed uncompressed, which drives the application needs towards terabit networking. Competitiveness of the market is no longer measured in milliseconds (ms), but in microseconds (μ s) [8].

requirement is common to all domains (therefore it was not listed as an entry in Table I), although clearly holding times and arrival rates of set-up requests can vary (e.g. typically fewer, but longer lasting application instances in scientific than in consumer applications).

- **Resource volume and granularity:** Almost all aforementioned applications pose challenges in terms of the sheer amount of resources they require. Here, resources may be for (i) data storage, (ii) processing capacity (CPU) and (iii) network bandwidth. The amount of resources (data volumes, processing, bandwidth) may greatly vary among applications (e-Science vs. business and residential use cases, see Table I). Hence, different paradigms may need to coexist to enable resource allocations with appropriate granularity, and possibly share common underlying hardware infrastructure.
- **Response time & latency:** Fast set-up of hardware/software resources may prove challenging for interactive/real-time services (e.g. multimedia editing). Especially for business or consumer applications (as opposed to batches of independent tasks in some e-Science scenarios), deadlines might be rather strict. An extreme example is presented in stock exchange transaction systems, illustrated in Fig. 3. There competitiveness of the market is no longer measured in milliseconds (ms), but in microseconds (μ s) and hence any delay in the trading path needs to be eliminated [8]. The ratio between holding times and set-up times for the service will greatly influence the efficiency of resource utilization. In addition, the responsiveness requirements of the application during runtime may also prove challenging (cf. interactivity) and impose constraints on the network. Note that a related issue is packet loss: usually, for real-time applications delay is more crucial than loss, and a certain degree of loss is even tolerable (e.g. for stock exchange applications, delaying market data wreaks havoc on trading applications and is significantly worse than discard). The reverse often holds for e.g. data-centric applications: losing the data may

TABLE I
REQUIREMENTS FOR NOVEL APPLICATIONS

| Application | Storage* | CPU* | Bandwidth* | Response time & latency [†] | Scale [‡] | Elasticity [^] | Multiple tasks [§] |
|---|----------|------|------------|--------------------------------------|--------------------|-------------------------|-----------------------------|
| <i>Scientific</i> | | | | | | | |
| Scientific computing (e.g. numerical simulations) | + | +++ | o / + / ++ | -- | 1 | + (<i>workflow</i>) | ++ |
| Data-centric computing (e.g. LHC) | +++ | ++ | ++ | -- | <10 | + | ++ |
| Sensor applications | ++(+) | + | +(+) | - | 1000s | - | ++ (<i>aggregation</i>) |
| <i>Business</i> | | | | | | | |
| Virtual meeting | - | + | + | + | 10s | - | - |
| Collaborative frameworks | o | + | + | o | 10s | + | - |
| Multimedia processing & editing | + | ++ | +++ | - / o | 1 | ++ | + (<i>workflow</i>) |
| Multimedia storage & retrieval | ++ | + | ++ | - / o | 1 | ++ | + (<i>workflow</i>) |
| Data mining | + | ++ | +++ | - | N/A | - | ++ |
| Transactional systems | + | + | + | - | N/A | - | ++ |
| <i>Consumer</i> | | | | | | | |
| Multimedia storage, editing & processing | +(+) | (+) | ++ | - | 1 | ++ | (+) |
| Action games | - | +(+) | +(+) | ++ | 10s | + | - |
| MMORPG | - | + | o | o | 1000s | - | ++ |
| Interactive TV | - | +(+) | ++ | + | 1 | + | - |
| Augmented reality | - | + | ++ | ++ | 10s | + | (+) |
| Virtual Tourism | - | + | +(+) | + | 10s | + | - |

*: Qualitative measures: - Low, o Neutral, + High

[†]: Real-time requirements: + Hard real-time (<100ms), o Soft real-time (<1s), - Non-real-time (order of seconds or more)

[‡]: Number of entities for a single application instance

[^]: Requirement for amount of resources of a single application instance: + Fluctuating, - Constant

[§]: Distributed nature of a single application instance: + Separate tasks, - Rather monolithic

involve repeating an expensive experiment.

- *Scalability*: The grid/cloud platform will need to scale well with not just the complexity of the application (in terms of resource needs), but also with the volume of applications it needs to host, as well over widely dispersed geographical areas. E.g., while consumer application requirements may be modest, the sheer number of users, and hence amount of application instances, may be gigantic (as opposed to e.g. a relatively small specialist research community in an e-Science field). The performance indicators to assess this scalability typically are resource utilization and response time, in addition to possibly application-specific metrics (delay, throughput, etc.).
- *Elasticity*: the needs of a particular application may vary over time. Thus, the amount of resources tied to a particular application instance may clearly need to vary over time. This means that in addition to on-demand set-up, the resource needs may vary during the lifetime of the application. Additionally, the number of users requiring a particular application may also vary over time (e.g. mainly during daytime for business, more in the evening for consumers, and rather continuously for scientific applications).
- *Multiple tasks*: Applications may differ greatly in terms of their internal structure. Depending on their nature as well as the platform they are implemented on, they may be composed of multiple sequential and/or parallel processes. Application instances could be monolithic, or rather composed of multiple tasks with varying degree of

interdependency (cf. bag-of-tasks vs workflows). The interdependency between those processes will greatly affect flexibility that could be exploited to schedule them in both the time and space domain. Thus, various resource allocation mechanisms should be supported to exploit that varying degree of flexibility. This is a challenge, since efficiently making the decision when and where to run what job/task [13] becomes an NP-complete problem for sets of interdependent tasks [14].

- *Geographical scale & awareness*: In terms of spatial flexibility, the grid/cloud system will need to be aware of certain constraints. E.g. some applications may rely on resources that are only available on particular locations (specialized equipment). Also for interactive applications, awareness of the location of resources can be benefited from to maximize user satisfaction. Note that such constraints will also affect topological design and resource dimensioning for the infrastructure.
- *Resiliency*: As in traditional services, reliability can be a prime concern [15]. In the considered grid/cloud applications, this applies to both the network and IT resources for storage and computation. Whereas various approaches have been defined in each of these domains (network vs IT), their interdependency needs to be kept in mind to successfully deploy resilient grids/clouds. Also, the most efficient way to provide resilience may depend on bandwidth, response time, etc. requirements of the application. (E.g. if the application relies on mainly short computational tasks to be executed, a fairly simple reissuing of the tasks on alternate resources may

be acceptable.)

Note that to address these requirements, the constraints for the underlying network technology may be dependent on network segment (access / data center / backbone) [16]. Thus, different approaches may be adopted in different network segments. In that case, we need them to seamlessly coexist and contribute to an effective end-to-end solution. For that, clearly a holistic view of all resource kinds (storage, computation, network) will be most beneficial to assure efficient resource utilization. Apart from this unification of network and IT resources, most important for grid/cloud platforms to deal with scalability, elasticity, etc. is to appropriately address spatial and temporal dynamics stemming from the elasticity and geographical awareness (to meet e.g. latency constraints) requirements. These indicate that, from a network perspective, the traffic patterns may be quite different from those in today’s common applications. Just as the classical Poisson model (that appropriately modeled telephone call arrivals) was questioned by the observed long-range dependency (LRD) and self-similarity in Internet traffic [17, 18] (even though correct assessment needs caution [19]), we may expect new models to accurately characterize the nature of grid/cloud applications. Such models could capture both e.g. rather predictable diurnal patterns correlated with human activity [20], but also effects such as flash crowds attracted by new, highly appealing (and hence popular) content/applications [21].

IV. GRIDS AND CLOUDS

The applications outlined in Section III imply a wide variety of resources that are interconnected in a distributed computing environment as sketched in Fig. 2. Information is generated in massive amounts by experimental facilities such as the LHC, as well as observational infrastructure (e.g. radio astronomy, or sensor networks). This data needs to be reliably stored and further distributed to various data centers, to be further analyzed. As pointed out before, also business- or consumer-oriented applications are increasingly challenging in terms of amount of data exchange and processing. Thus, clouds and grids imply the need for efficient interworking of networking, storage and computing resources.

Grids originated mainly from the needs of e-Science applications (cf. Open Science Grid [2], the Enabling Grid for E-Science or EGEE [3], and TeraGrid [4]), esp. high performance computing (HPC). The seminal work of Foster et al. now lists use cases beyond that, venturing into more business-oriented applications (e.g. fault diagnosis in jet engines, biomedical imaging) or even multiplayer video gaming (see Part III of [1]). The core characteristics of a grid are concisely summarized in this three point checklist [22]: *a Grid 1) coordinates resources that are not subject to centralized control, 2) uses standard, open, general-purpose protocols and interfaces, and 3) delivers non-trivial qualities of service (QoS)*. The requirement for coordination originates from the nature of grid applications (e.g. see the science applications in Table I): they are primarily job oriented, implying a need for e.g. workflow-based task coordination

taking into account inter-task communication and dependencies [23], and sometimes require the ability to make resource reservations in advance. In the academic environment they were conceived in, open interfaces are a necessity for a path to global sharing of e.g. research facilities across multiple administrative domains. The need for QoS for the network mainly pertains to bandwidth (cf. large data volumes), since grids usually do not natively support interactive applications [24] (cf. non-real-time nature of many e-Science applications, Table I) and hence latency and delay tend to be less of an issue. Grids do offer time sharing of resources, but generally do not offer explicit (dynamic) partitioning of the hardware infrastructure into virtual resources (which is a foundation of the cloud paradigm) [25].

Build on the basic ideas of grids (e.g. similar coordination across resources may be required in cloud mashups), clouds manifest themselves in more commercially oriented applications (as opposed to the public funded research oriented grids), which often involve loosely coupled tasks, and are typically interactive [24]. In terms of infrastructure, they typically run in large data centers (as opposed to HPC infrastructure for many grid applications). The essential characteristics of clouds become apparent when studying cloud definitions [23]: user friendliness, virtualization, scalability, pay-per-use model, and SLAs. The user friendliness stems from the wider target audience (i.e. business and consumers, vs academics for e.g. e-science oriented grids). This mainly refers to the easy access to, and the deployment and configuration of the resources used (which can range from hardware to applications, see further), which is typically internet-based, leveraging the service-oriented Web 2.0 paradigms (e.g., [26]). Scalability is also one of the main drivers for cloud deployments, exploiting a pay-as-you-grow approach, which appeals to businesses. The success of cloud computing clearly is based on this scalability (which can also involve automatic adaptation), where cloud providers benefit from economies of scale and statistical multiplexing: using virtualization allows them to operate cost-effectively, avoiding peak load (over)provisioning. This virtualization is a key difference that clouds bring to the table, compared to grids. Virtualization allows them to share resources in a safe way, facilitating to respect the performance SLAs agreed with their users. (Compared to traditional grid applications, in particular the response time and latency requirements of business/consumer applications can be more stringent, see Table I.) Virtualization furthermore enables migration to other servers, both for performance and resilience against failures. Also, monitoring in clouds is quite challenging (partly because the user is not at liberty to install and run his/her own monitoring infrastructure), whereas grids apply a different trust model where users, via identity delegation, can access and browse resources at various sites that contain resources. In grids, these resources are typically not that much abstracted nor virtualized. (However, in view of the clear advantages of virtualization, also grids are evolving in that direction [27]; see below, Section V). Further in-depth analysis of grid vs. cloud computing is discussed in [24].

In terms of architectures, the many attempts (e.g. [28, 29, 30]) to classify various cloud paradigms seem to converge to a layered “everything as a service” (XaaS) taxonomy, comprising the following layers (see Fig. 4):

- *Software as a Service (SaaS)*: This layer comprises all applications that run on the cloud and provide a direct service to the customer/user. This layer can be further subdivided according to the application level offered. On top, we have the actual *Applications* which are basically the final service offered to an end user, such as Google Docs or Microsoft Office Live. Clearly, they can be composed (following e.g. a service-oriented approach) of lower level services: [28] categorizes them further into *Basic* (e.g. Google Maps) and *Composite Application Services*.
- *Platform as a Service (PaaS)*: Users of PaaS are provided with an application or development platform, which allows them to e.g. create SaaS applications/services. The PaaS layer can be further decomposed into *Programming Environments* and *Execution Environments* (for instance Django running on top of Google App Engine). The former provides programming-language-level environment (i.e. well-defined APIs), whereas the latter offers the run-time execution environment that can take care of e.g. automatic scaling and load balancing.
- *Infrastructure as a Service (IaaS)*: This lowest level provides the underlying resources, i.e. storage, computing and networking, which PaaS/SaaS rely on. (Note that SaaS applications do not necessarily rely on the intermediate PaaS layer.) The “resources” can refer to physical resources (i.e. servers), but these often are virtualized. Hence we distinguish both virtual and physical resources (where the former abstract/partition the latter). These resources can be further abstracted into what [28] calls “basic infrastructure services” providing higher level functionality than that offered by a typical OS (e.g. Google File System as storage service). Offering database functionality is an example of “higher infrastructure services” (e.g. Amazon’s SimpleDB or Google’s BigTable).

We note that [29] also mentions Hardware as a Service (HaaS), referring to providers that offer server infrastructure and take care of operation, management and upgrades of the hardware. In our view this can be categorized as the physical resource sub-layer of IaaS. (Note that [30] treats IaaS and HaaS as synonyms.) On top of SaaS, [28] also introduces an extra human-as-a-service (HuaaS) layer, but in our view, these can be seen as particular application services (thus residing in the SaaS layer) which rely on interaction and actual data processing by multiple collaborating people. Also, from an architectural perspective, the intermediate layers (esp. PaaS, but also some supporting functions for e.g. monitoring and management of IaaS components, as well as basic application services in SaaS) can be seen as “cloud middleware”, which [30] categorizes in *User Level* and *Core Middleware*.

Looking at de facto standard grid technology such as the

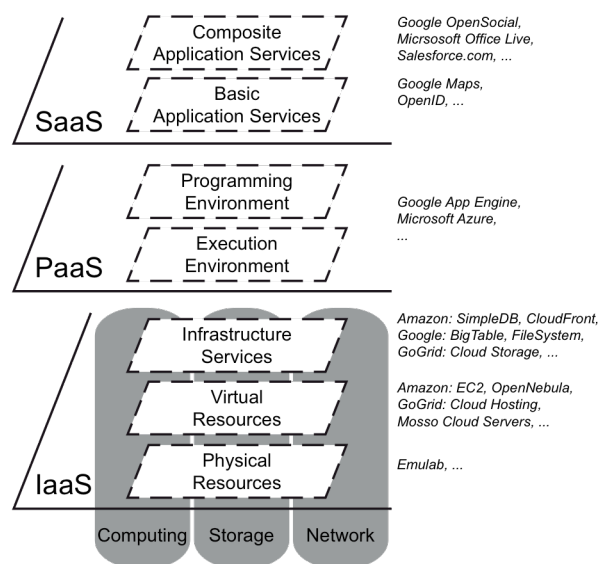


Fig. 4. Cloud computing paradigms can be classified as forms of “anything as a service” (XaaS): (i) Software-as-a-Service (SaaS), (ii) Platform-as-a-Service (PaaS) and (iii) Infrastructure-as-a-Service (IaaS). The physical resources which cater for these services include servers for computing and storage, as well as network equipment.

Globus toolkit, we note that the grid software today mostly seems to be situated on this “middleware” layer, which one could (arguably) position on the PaaS level. The evolution to web-service based access to grids [31] could be seen as more SaaS-like grid offerings. Thus, from a conceptual point of view, it seems that grids are converging to the same (if not, very similar) layered architecture. While clouds (business/consumer oriented) can be seen as an evolution of grids (more biased towards science and HPC), we hence believe their concepts and architectures may converge (see further, in Section VI).

Note that in the taxonomies/classification of cloud paradigms, the network is often neglected. Youssef et al. [29] coin the term “Communication as a Service (CaaS)”, as one of the three types of infrastructure (next to computational resources and data storage), offering dynamic provisioning of virtual overlays for traffic isolation or dedicated bandwidth with QoS guarantees etc. They mainly refer to interfaces for the creation of on-demand communication services or channels (e.g. using SIP and IMS-like technologies). Lower layer virtualization is not addressed there. Similarly, [28] indicates network as a “basic infrastructure service”, stating OpenFlow [32] as an example. Nevertheless, the concept of network virtualization, and especially its unified integration into a grid/cloud-like paradigm is very much an open research problem (as discussed further, Section VI).

V. SOLUTIONS

In this section, a number of technologies are introduced that help fulfilling the requirements as discussed in Section III. First we focus on optical networking technologies, detailing

the relevant transmission and switching techniques in various levels of the network. Our attention then shifts to routing paradigms and how these enable desirable features such as geographical awareness and scalability. As discussed previously, virtualization is of vital concern to cloud computing, and thus we introduce the concept and demonstrate its effectiveness in meeting scalability and elasticity demands. Finally, we discuss current and future control and management solutions and relate their features back to the original requirements.

A. Optical Technology

Optical networking can play a key role in the realization of grid and cloud computing systems. In the following, we discuss the capabilities of optical transmission and switching techniques, and its application in access and local networks.

Optical transmission is generally accepted as the most cost-effective way to realize high-bandwidth connections in the long-haul network [33, 34]. The technology's ability to transfer huge data volumes with low latency has made optical networks the de facto standard to connect data centers that provide computing and storage services in grid and cloud computing networks. Furthermore, techniques to offer resilient network operation despite common failures such as fiber cuts have been researched in depth and are widely deployed today [35].

Frequency division multiplexing permits optical end-to-end connections over a specific frequency or wavelength, a technique known as (Dense) Wavelength Division Multiplexing or (D)WDM. Commercially available line rates offered by a single wavelength include 10, 40 or 100 Gbps, while channels are typically spaced 50 or 100 GHz apart [36]. All-optical end-to-end connections have very low latencies that are, in general, only limited by the communication distance and physical impairments, since no processing takes place during transit. These can also achieve tremendous scalability because of the very high data rates offered. However, as indicated in Table I, not all users (in particular consumers) require high data rates; forcing low speed connections on high data rate wavelengths ultimately wastes bandwidth and reduces resource utilization.

A first step towards better matching user requirements and offered bandwidth consists in careful network planning such that multiple wavelength granularities can coexist; these are referred to as mixed line rate (MLR) optical networks [37]. Note though that bandwidth may still be stranded if end-to-end traffic does not suffice to fill an entire wavelength.

An elegant solution to this problem is through re-engineering of the static and coarse channel spacing into a more flexible and finer spectrum grid [34]. Next to its original objective of improving spectrum efficiency, the smaller channel width also reduces the mismatch between offered bandwidth and user demands. Sub-wavelength granularities are offered through transmission over a limited number of low data rate channels, while a large number of such channels may also be combined to offer true scalability in the form of super

wavelength capacity. As such, elastic access to optical bandwidth is achieved by on-demand growing and shrinking of the allocated resources [38, 39]. This obviously could be very effective for the spatial and temporal nature of cloud traffic as discussed in Section III. Moreover, the work in [40] proposed to use optical OFDM for cloud computing environments, with a specific focus on its capabilities regarding virtualization of optical networking infrastructure. Although very promising, elastic optical networking is still a relatively new research field. As such, the main technical challenges to successfully adopt the technology in grid and cloud computing networks, are not yet fully understood.

Apart from transmission technology, *optical switching* techniques also attempt to bridge the gap between optical bandwidth and user demands. More specifically, WDM corresponds to an Optical Circuit Switching (OCS) solution, where bandwidth is reserved exclusively for communication between source and destination. As discussed before, this may lead to inefficient use of bandwidth if the traffic demand does not match the full capacity of such a wavelength. Another drawback is that the reservation process takes a non-negligible amount of time to complete, because of a two-way signaling operation to allocate bandwidth and configure switching devices. Despite these issues, successful examples of WDM-based Grids include the CineGrid [41] and TransLight projects [42], which utilize optical circuits to build grid computing networks on a global scale. Likewise, optical networking technology is an essential component to interconnect datacenters on a global scale as is done in, for instance, Amazon's EC2 and Google's cloud offerings. Further examples of practical implementations can be found in Section V.D.

Whenever jobs transfer a sufficient volume of data, the overhead of the reservation process will remain low (or may be reduced by considering Fast Optical Circuit Switching (FOCS) [43]). However, the problem is compounded when the actual data transfer time becomes relatively small compared to the reservation time. Table I indicates that this is the case for some business and most residential or consumer applications, particularly when compared to e-Science jobs. In response, some researchers have considered Optical Burst Switching (OBS) when less stringent data transfer and processing requirements are needed [44, 45]. The idea is to map each job onto a data burst and to reserve optical bandwidth only for the actual duration of the data transfer [46]. As of yet, OBS has proven difficult to implement due to the strict requirements it poses to the optical switching and electronic processing devices, and suffers from poor performance (i.e., low bandwidth utilization and relatively high blocking) under certain circumstances.

Ultimately, Optical Packet Switching (OPS) can offer the finest granularity to bandwidth, but current implementations are severely limited because optical memory is still a major challenge [47]. Furthermore, OPS does not offer power savings compared to electronic packet switching [48]. Either way, both OBS and OPS hold the potential for fine-grained




Fig. 5. Since grid/cloud users usually do not care where their tasks are executed, they can be executed at a different site under failure conditions. Such exploitation of *relocation* to protect against single link failures (1L) can lead to significant savings in bandwidth (13% to 20% for a particular case study on a 28-node EU network, with 3 server sites). This however comes at a price of extra server capacity (9% to 14%). Yet, to protect against server site failures as well, backup servers need to be provided anyhow. Protection against both single link and complete data center failures (1LS) imposes extra network capacity, but relocation allows to provision less servers than in case of 1+1 data center protection (+52% to +60% in the case study, vs. +100% for 1+1) [67].

bandwidth granularities and low latencies, as well as excellent elasticity features. Their success depends however on the maturity, energy efficiency and cost-effectiveness of their implementations. Recently, Optical Flow Switching (OFS) [49] has been proposed as an end-to-end transport service supporting end-to-end lightpaths that are set-up for 100s of milliseconds or longer (for comparisons to OCS, OBS see [50]): this seems quite appropriate for any grid/cloud application with substantial data exchanges (see bandwidth requirements in Table I).

Similar to MLR networks, some researchers have combined multiple switching granularities in a single networking architecture [51, 52, 53]. In particular, an optical cloud solution was demonstrated in [54].

On the other hand, Table I also contains a number of business and e-Science scenarios where users require very large bandwidths, in the order of several tens or hundreds of Gbps. In these cases, multiple wavelengths can be combined into a waveband that can be switched as a single entity. Moreover, all wavelengths in a fiber can also be switched jointly, to establish the ultimate scalability in terms of data volume [55].

Whereas we previously stated that optical technology is mainly used in the long haul, core network, the technology has also been adopted in *access networks* [56, 57]. For instance, passive optical networks (PONs) are a cost-effective way to offer high-speed network connectivity to residential or business users [58]. Combined with their (almost) symmetrical upload and download speeds, time-division multiplexed PONs potentially allow OPS or OBS to be deployed, while wavelength-division multiplexed PONs can be based on FOCS or even OCS.

Besides access networks, optical technology can also play an important role in *local networks*. Most notably, as the scale and requirements of data centers grow, optics is increasingly being introduced in what is traditionally an electronically switched environment. Careful network design can lead to highly scalable data center internal networks and reduce cost, energy consumption and complexity in topology [59, 60]. We

do not consider this in more detail in this paper, but instead refer the interested reader to work introduced in [5].

In summary, advances in optical technology clearly address the requirements for *high data volumes*, and *elasticity* (if allocated bandwidth can be dynamically adapted). Making efficient use of these capabilities with appropriate control and management approaches (see Section V.D) should ensure *scalability*.

B. Routing

Routing and path computation algorithms are undeniably of fundamental concern in communication networks. In the following, we argue that these algorithms need to be carefully designed in a grid or cloud computing environment. To this end, we focus on anycast and multicast routing algorithms in general and in optical networks in particular. It will become clear how anycast routing helps to meet the requirements for *scalability* (cf. load balancing) and *resilience* (e.g. relocating jobs to an alternate server, see Fig. 5), and could exploit *geographical awareness* (e.g. sending to closest candidate destination of anycast-set). We also present advanced optimization strategies for anycast-routed service networks, and conclude with the issues inherent to the planning and dimensioning of an anycast-routed network.

Traditionally, data networks employ unicast routing algorithms (e.g. shortest path routing) for transferring data from source to a given destination. In cloud and grid networks however, each user-generated task can be serviced at multiple locations in the network. Moreover, the exact service location and network route is of less importance to the end user. Instead, his main interest lies in successful execution of the task, while observing the quality attributes as specified by the SLA. A fundamental concept to realize such service-oriented networks is the *anycast* routing principle [61]. Anycast routing specifically enables users to transmit data for processing, storage or service delivery, without assigning an explicit destination. By simply using an anycast address, service providers can offer a generic interface to end users for a wide

range of services and applications. The challenging task of finding suitable network and IT resources for a given task requires knowledge of state information on the network and IT resources, including e.g. current traffic loads. Consequently, desirable scalability features such as load balancing or congestion control can be implemented, where network and IT resource constraints are jointly taken into account [62, 63].

Next to anycast, we also consider *multicast* routing and its application in grid or cloud networks. It can obviously be very useful to, for instance, distribute an identical data set to multiple computing nodes. Alternatively, computing and streaming of data to multiple end users is of concern in case of, for example, video transcoding and scene rendering in gaming environments.

Optical networks in particular may benefit from the use of anycast or multicast routing algorithms. For instance, consumer-oriented grid applications can be supported by an optical grid architecture based on anycast routing and OBS [45]. Another example is an optical network interconnecting multiple data centers on a global scale. Energy consumption is of great concern in this case, and as such sites may be switched on or off depending on demand; anycast routing then assists in finding an appropriate data center.

A variety of anycast-based optimizations may be considered to further improve cloud network performance. For instance, not only the location of execution of the service can impact performance, but also the location of the data being processed, or the proximity to the end user. Thus, the execution of computational tasks should be directed to the most appropriate locations while minimizing communication overhead [24]. Indeed, job turn-around times can be reduced significantly in this way [64], since moving data repeatedly to distant CPUs increasingly could become a bottleneck. In effect, this introduces geographical awareness to cloud networks, an important requirement as listed in Section III. Note that to exploit such data locality knowledge in advanced schedulers, such locations must be advertised¹. However, in current cloud offerings for storage, this information is rarely exposed (although e.g. Amazon EC2 allows choosing virtual resource instances in a particular geographic region).

Technically speaking however it is not straightforward to deploy anycast or multicast routing algorithms in a grid or cloud computing environment. Indeed, such environments are composed of large numbers of IT and networking resources, whose states are inherently dynamic. Efficient techniques to collect and manage this state information, and to calculate optimal routing decisions, are still an active research topic given the challenges to deploy them on a large scale (e.g., the European Grid Infrastructure, a continuation of EGEE [3], offers around 250,000 CPU cores at any point in time)

Finally, one should note that *dimensioning* anycast-routed

networks forms a considerable challenge; indeed traditional network design methods are difficult to use since we are not given the complete so-called traffic matrices stating required bandwidths between given source and destination pairs, but rather only the origin of the demand, while the destination can be freely chosen [65]. However, substantial bandwidth and energy savings can be achieved by intelligent use of the anycast routing principle. For instance, we can exploit the anycast principle to reroute jobs to an alternate location under failure conditions [66, 67]. Illustrative case studies show that this may incur bandwidth savings around 20% compared to shared path protection towards the original destination, as illustrated in Fig. 5 [67]. Similarly, by exploiting such relocation to protect against both link and data center failures, we need only ca. 50% extra servers (vs. 100% if we were to provide 1+1 data center protection).

C. Virtualization

A fundamental driver for the cloud computing paradigm is the concept of *virtualization*, which is the process of either *partitioning* a single physical resource in multiple virtual resources (1:N), or *aggregation* of multiple physical resources in a single virtual resource (N:1). As such, a common physical infrastructure (which may be an IT resource, a network resource, or any combination thereof) can be shared among multiple users. A key characteristic is that these users do not interfere with each other as they operate within their own virtual infrastructure. Furthermore, *abstraction* is necessary so both physical and virtual resources can be described in a common language to identify their attributes and capabilities in a generic way. Note that virtualization techniques in general aid in offering fine granularity and elasticity in terms of resources that are available to the user. In the following, we consecutively discuss virtualization techniques on the network level, for IT resources, and the combination of both network and IT resources. We also cover benefits and potential disadvantages of using virtualization, both in general terms and its effect on performance parameters in particular.

Network virtualization [68, 69], and more specifically on optical links, has seen numerous studies and implementations in the form of Virtual Private Networks (VPNs), which connect a number of known end-points over a dedicated communications infrastructure. For instance, optical Layer 1 VPNs [70] are provisioned to establish connectivity between data centers, as such creating a high performance public or private cloud network. Such L1VPNs can be rapidly allocated or even reconfigured, to cater for fast changes in server loads. Finally, various bandwidth granularities are supported by L1VPNs, to ensure the provisioned capacity follows bandwidth demands.

Virtualization of networking resources has also been considered an essential technology to realize the future Internet [71], since it allows researchers to *slice* off parts of an operational network and use it as an experimental testbed for novel architectures and protocols. This implies that private grid or cloud infrastructure can be fully customized according

¹ Another incentive for cloud providers to disclose such geographic information may stem from regulatory pressures that force enterprises to be careful where exactly their data is stored/processed.

to the owner's requirements. Consequently, one may choose to deploy non-standard protocols for reasons of, for instance, scalability.

Finally, whereas VPNs create isolated logical networks on a common physical substrate, recent work introduces virtualization in most if not all network elements, such as the switching fabric, the routing and forwarding engine, and the control plane [32].

IT resources are made up of multiple components such as a processing unit, storage devices and working memory. *Virtualization of computer systems* results in a virtual machine (VM) that offers all the capabilities of the host resource. These VMs can be instantiated and configured on-demand and introduce a relatively limited overhead. Furthermore, partitioning and aggregating of, for instance, storage resources, leads to the desirable properties of granularity and scalability, respectively. It should be clear that IT resource virtualization adheres to the majority of requirements as listed in Section III. This can be partly explained by the technology's widespread adoption and proven commercial success as exemplified by, among others, Xen [72], KVM [73] or VMWare. Further details concerning virtualization of IT resources can be found in [74, 75].

Only recently, *combined virtualization* of both networking and IT resources has gained widespread attention, mainly due to the popularity of the grid and cloud computing concepts. In particular, the Generalized Architecture for Dynamic Infrastructure Services (GEYSERS) is a European FP7 project that designs and implements a novel architecture for seamless and coordinated provisioning of both optical and IT resources, and develops the necessary tools and software to realize this objective [27]. The idea is to introduce a Logical Infrastructure Composition Layer (LICL) that manages the physical infrastructure consisting of both network and IT resources, and exposes these as virtual resources in a generic way. These, in turn, are combined to form virtual infrastructures that operate independently from each other, and each deploys its own control plane solution as desired. Additional features include dynamic up/down-grading of these infrastructures, as well as guaranteed end-to-end service delivery over diverse resources and complex different technologies. One of the project's main technical challenges is to demonstrate that the proposed architecture can gracefully scale to support large physical infrastructures as well as handle a significant number of virtual network instances.

In parallel, several higher-level cloud management toolkits have been proposed to handle aspects of IT resource virtualization combined with advanced job scheduling, monitoring, storage and user management. In contrast to the aforementioned GEYSERS project however, their support for network virtualization is limited to at most connectivity through L2 tunneling. Modern examples of the latter are OpenNebula or Eucalyptus besides several others [76, 77, 78]. These software solutions allow transforming a network of cluster nodes to cooperate in managed cloud network. A very

thin layer of management software must be installed on the machines, providing a standardized API to the previously mentioned platforms. The software toolkits keep track of available resources, adequate scheduling between different cloud service requests and allows status monitoring of the provisioned services and resources.

In general, virtualization can offer a number of *qualitative advantages* over more traditional models, including stricter isolation between users, more flexible enforcement of security policies and higher levels of trust [79]. However, one should not assume these advantages to be implied by virtualization, as careful design remains essential to successfully operate these services. In particular, the study in [80] demonstrates the trade-off inherent to optical WDM network virtualization, and specifically the effect of isolating virtual networks on network dimensions and the control plane scalability.

Revisiting the cloud and grid requirements, virtualization mainly caters for *elasticity* and *scalability*, and addresses diversity of applications in terms of *granularity* of their resource needs. The flexibility of on-demand resource provisioning of virtualized resources, due to the less stringent dependence on the availability of a particular physical resource, also enables extra *resilience* opportunities.

D. Control and Management

A major challenge in successfully operating optical cloud networks is to have an integrated control and management plane for both network and IT resources. The cloud management middleware solutions as discussed in the previous section can effectively control and virtualize the cloud's resources. However they are currently not ready to natively support optical networking technology. Typically, these toolkits rely on pre-configured L2 connectivity (Ethernet in most cases) to create virtual networks for the cloud services they offer. Traffic within different virtual networks is typically isolated using either MAC address filters or VLANs. Prior studies have shown that these Ethernet-based techniques are not really secure, nor scalable (limit of 2^{12} VLANs) [81].

These challenges are by and large addressed in the aforementioned GEYSERS project. Indeed, the LICL toolset performs abstraction of optical networking equipment and IT resources, supporting a truly multi-technology and multi-vendor environment [27]. Virtual infrastructures can then be built with arbitrary granularity, while LICL optimizes, among others, utilization and energy efficiency of physical resources. Furthermore, the project's extensions to the GMPLS control plane introduce novel constraint-based routing algorithms that incorporate network, IT and energy-related parameters. As such, applications can operate within a converged architecture, as, for instance, SLAs can be negotiated and enforced vertically, implying tight cooperation between the application, control plane, LICL and physical layer.

Alternatively, OpenFlow [32] has been proposed as a sort of *network operation system*, whereby the forwarding tables in switches and routers can be freely and fully programmed. OpenFlow's approach towards programmable networks

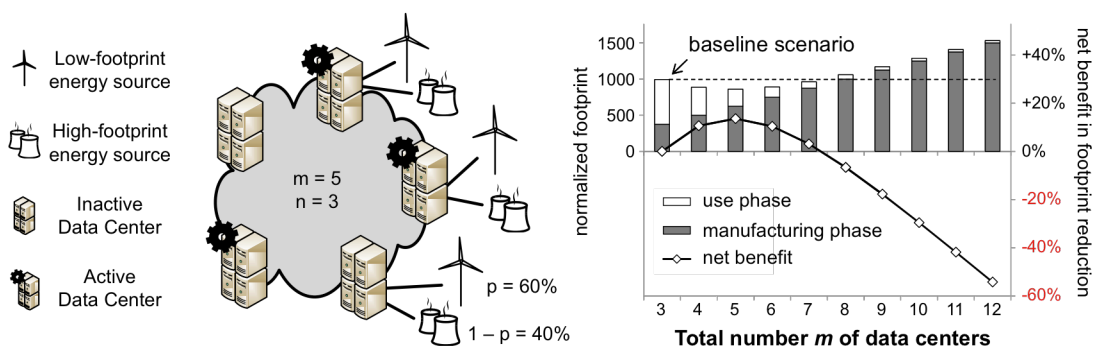


Fig. 6. Data centers may be powered by a combination of “green”, low-footprint energy sources at one point in time, and conventional high-footprint sources at another, depending on availability of the green source (e.g. with a probability of $p=60\%$). To minimize footprint, dynamic migration towards another center may be considered (powering off, or putting to sleep, the vacated servers) if the green source is not available locally, but it is at the remote site. A case study detailed in [97] showed that the net benefit in terms of footprint reduction, taking into account also manufacturing, depends on the number of locations (sample result shown is for a constant demand requiring $n=3$ servers, and a varying number m of locations between which we can migrate).

(achieved by separating the control from the data path) allows virtualization of Ethernet switches and IP routers. By appropriately programming the routers/switches, the network can be logically partitioned in so-called slices (to be seen as separate virtual networks), where each slice can use its own set of (experimental) protocols and policies. As of now, OpenFlow has gained considerable support from commercial vendors, although intelligence regarding IT resources has not been considered yet. Note that a ForCES enables a similar separation between the virtual network's control/operation and the management of the physical network infrastructure [82]. Yet another alternative to virtual network management is proposed in [83], where routers are virtualized and allowed to move within the network.

Whereas cloud projects have not yet fully embraced optical technology, the grid community very well realized the opportunities that optical networks offer — given their high bandwidth and low latency (which is especially relevant in interactive applications, e.g. using 3D and/or high-resolution visualizations). Various projects target so-called *lambda grid* scenarios, which are commonly OCS-based grid networks that offer user-controlled bandwidth provisioning in an on-demand fashion. To incorporate that dynamical optical network functionality in a grid context, two fundamental approaches can be followed [44]: either the grid and optical layer are managed separately in an overlay manner, or an integrated control plane is deployed, e.g. by extending the optical control plane for grid resource provisioning. Clearly, the latter approach allows for joint optimization of both optical network and IT resources. The optical network community has been working extensively on the control plane extensions towards grid scenarios: for an elaborate discussion of these issues, and relevant projects and organizations (such as the Open Grid Forum, OGF), we refer to [84].

Of particular interest is the PHOSPHORUS project [85], which addresses some of the key technical challenges to enable on-demand, end-to-end network services across multiple, independent, high-performance transport domains. One solution is based on the Generalized Multi-Protocol Label

Switching (GMPLS) protocol suite, which is frequently deployed to bridge the gap between optical transport technology and the IP layer. (Another solution uses a service plane approach called Harmony [86].) GMPLS includes signaling, routing and path-computation functionality for a wide range of circuit-based and packet-based technologies across multiple layers. In the PHOSPHORUS project, GMPLS was extended to include attributes and states of Grid end resources, along with algorithms to use this information for anycast-based routing and improved resiliency (in terms of both network and IT resource failures). Apart from offering very fine traffic granularity, the project was the first to build an optical grid solution that adhered to all requirements as listed in Sec. III.

The Dynamic Resource Allocation in GMPLS Optical Networks (DRAGON) project in USA focused on the dynamic provisioning of network resources, including the establishment of traffic-engineering paths using a distributed control plane. The main objective was to support e-Science applications with deterministic network resources to link computational clusters, storage arrays, visualization facilities etc. [87]. TransLight, a global project involving government-funded research labs and universities in EU, Japan and North America, designed and implemented a system of optical networks to move massive amounts of data directly under e-scientists' control [42]. The OptIPuter idea builds on the availability of dedicated wavelength channels between distributed clusters to link them together in a virtual metacomputer [88].

VI. CHALLENGES AND OPPORTUNITIES

As already pointed out, major challenges are related to the integration of cloud and grid technologies (which mainly focus on the IT resources) with optical networking. A second set of associated research opportunities stem from the fact that optical clouds could play a significant role in addressing concerns on greenhouse gas emissions. Finally, also specific grid/cloud related issues need to be addressed. We summarize each of those in the following.

Firstly, for the optical network one of the main challenges will be to provide virtualization mechanisms that fit the needs of grid/cloud applications. This can imply both sharing of the same physical resources (and thus exploit statistical multiplexing (cf. OPS, OBS) and/or time sharing (cf. fast circuit switching), while still offering isolation), as well as partitioning them (cf. slicing). Thorough assessment of the costs/benefits of such virtualization approaches in terms of e.g. resource utilization penalties/gains is mainly an open question.

Clearly, on-demand provisioning will remain important in optical networks in a grid/clouds context, but it should be offered in an integrated way, to allow coordinated IT and network resource provisioning. Also, some grid/cloud applications will require joint provisioning of a complete network topology (see e.g. [89]), as opposed to traditional point-to-point connections. Moreover, as previously indicated, not all end points will be specified (cf. anycast). Fundamentally, this leads to new network design and optimization questions that need to be addressed (e.g. resilient dimensioning of optical networks for anycast applications [67]), as well as challenging provisioning and scheduling problems (e.g. multi-site data aggregation [90, 91], exploiting flexibility in the time domain [92], etc.). In addition, control and management plane technologies should be put in place that realize interworking/integration of IT and network worlds. While this work is being addressed in ongoing research projects, exhaustive assessment of e.g. scalability of the approaches and robustness against failures today is still largely unaddressed.

Secondly, in view of the growing importance of greenhouse gas emission reduction, also the energy efficiency of networking solutions in general, and optical cloud solutions in particular needs careful consideration. Optical networks provide some key advantages in terms of power consumption when compared to electronic technologies [93]. Evaluating the cloud computing paradigm in terms of power consumption requires a careful balancing of network, storage and processing energy [94]. Elasticity and virtualization already promise accurate matching of on-demand provisioned resources to the instantaneous demand, where powering off currently unused servers can cut energy consumption [95]. More revolutionary ideas, such as locating data centers close to renewable energy sources and moving cloud jobs from one data center to another depending on the instantaneous renewable energy production (sometimes dubbed as “follow the sun / follow the wind”) are also being investigated [96], as illustrated in Fig. 6. Such relocations to more environmentally friendly locations clearly could increase the dynamicity of traffic patterns, which emphasizes the need for a both flexible and responsive network. Note that to get the complete picture, besides the emphasis on power consumption during the use phase of network and IT equipment, also full life cycle analyses of the involved equipment are required for a more accurate assessment of greenhouse gas emission impact. The example of Fig. 6 shows the net benefit of providing data

centers at multiple locations (m in total), for the illustrative case where $n=3$ centers are required at any given moment: while adding more locations, the probability of powering $n=3$ of them with green energy increases, yet the equipment manufacturing footprint eventually outweighs the gain [97].

Lastly, also in the grid/cloud paradigm itself, important questions remain unanswered. Observing the growing interest of e-Science users in cloud(-like) solutions to serve their application needs (e.g. [98, 29] and references in [30]), we believe the integration of grid and cloud paradigms poses a promising challenge to address. While both grids and clouds can be expected to continue to co-exist (serving e.g. different application domains, such as e-science grids vs business and consumer oriented clouds), they can share many architectural concepts (and eventually soft- and hardware components). Such (partial) convergence involves introducing the virtualization concept in grid scenarios, including associated user-friendly interfaces exploiting them [31], while keeping the high performance characteristics, especially in terms of networking [23]. Indeed, particularly the network seems to be the bottleneck when running distributed e-Science applications in current commercial clouds [99, 100, 101]. This will involve enhanced support for (network) QoS, monitoring, federation of different organizations: issues that have been addressed previously (partly) in grids, but may need to be revisited [23], especially to incorporate virtualization. Federation needs to address e.g. trust/security issues related to sharing of large volumes of data and computational resources in an untrusted multitenancy cloud [28]. While cloud solutions today seem to perform quite satisfactory, getting insight in application performance requires a fine balance of business application monitoring, enterprise server management, virtual machine monitoring, and hardware maintenance, and will be a significant challenge [24]. How to use such monitoring information to efficiently address elasticity and keep meeting SLAs by automatic up/downscaling both IT and networks is a next (mainly unexplored) step. The multi-dimensionality of the grid/cloud setup, which allows not just for up/downscaling but also e.g. replicating or migrating/relocating processes to alternate locations leads to intriguing questions on what alternative will be most effective. Also, questions arise at which layer (IaaS/PaaS/SaaS) to address aforementioned scalability/elasticity, resilience, etc.

Looking at the plethora of existing cloud systems and offerings, it is clear that integration and interoperability of all the services and applications probably remains one of the biggest challenges (including e.g. a unified resource description language covering both network and IT resources). Even though some de facto standards seem to be emerging (e.g. Amazon’s interfaces that are also supported by open source solutions; increasing popularity of Xen; etc.), standardization bodies and fora (such as OGF, VXDL forum, etc.) will be crucial in setting the scene and establish an open, competitive environment. In such a new environment, the traditional roles of IT and network providers might shift (which may impact network structure, and even traffic

patterns, cf. the observed evolution in [16]). Hence, opportunities are plenty to develop new business models and for new players to enter the scene in this exciting arena (even though the development of a profitable revenue model could be quite challenging). However, lack of standardization (esp. addressing network and IT convergence) will likely hamper such open competition, thus delaying innovation and potentially blocking the evolution towards optical grids/clouds.

VII. CONCLUSION

In the evolution of computing paradigms, to grids and more recently clouds, the role of the network becomes increasingly important. We have outlined the novel applications that gave rise to this evolution, and identified their requirements. Even though academic applications (e-Science in particular) are pioneering inventors and adopters of new technologies and paradigms, we clearly illustrated that also business- (cf. financial market applications striving for zero latency and towards terabit networking) and consumer-oriented applications are increasingly demanding. We summarized a taxonomy of grid and cloud systems that offer an answer in meeting those stringent demands.

The foundational concepts and technologies that can realize those systems have been identified. An important role will need to be played by optical networking technology, where multigranular switching concepts such as mixed line rate (MLR), flexible (grid-less) switching, can help to address the requirements for flexible bandwidth. From a routing perspective, many novel applications imply extensive distribution of data (e.g. for distributing experimental measurements or sensor data) which can benefit from optical multicasting. The grid/cloud paradigm also gives rise to anycast-routing: the end-user often does not greatly care where his processes are running, thus introducing a degree of freedom (hence optimization) in deciding where to serve which requests (and efficiently routing the involved data to/from there). Cloud computing heavily relies on virtualization of IT resources, thus offering logical partitioning and possibly aggregation to efficiently serve time-varying volumes of application requests.

Moving towards “optical clouds” by pushing the virtualization paradigms to optical networks would enable full grid/cloud convergence and hence realize a future proof platform offering flexible, scalable IT and network resources. Routing concepts such as anycast, and coordinated reservation of a (virtual) network topology, jointly with IT resources calls for innovative allocation algorithms (as well as monitoring tools and e.g. autonomic management [102, 103, 104] to help implement them) and related network design solutions. From a control and management perspective, we note that convergence of IT-oriented cloud toolkits and on-demand provisioning technologies such as GMPLS could be a way forward to realize “optical clouds”, for which we identified main challenges.

ACKNOWLEDGMENT

The authors are grateful to prof. Biswanath Mukherjee (UC Davis, CA, USA) for guidance in setting the theme of this paper. The authors would also like to thank their colleagues Wouter Tavernier for providing input to Section III.D and Sofie Demeyer for assistance in preparing Fig. 2. The part of Fig. 1 on Moore’s Law is based on the corresponding Wikimedia Commons figure (available on http://commons.wikimedia.org/wiki/File:Transistor_Count_and_Moore%27s_Law_-_2011.svg).

REFERENCES

- [1] I. Foster and C. Kesselman, *The Grid: Blueprint for a New Computing Infrastructure (2nd ed.)*. Elsevier, 2004.
- [2] R. Pordes, D. Petravick, B. Kramer, D. Olson, M. Livny, A. Roy, P. Avery, K. Blackburn, T. Wenaus, F. Würthwein, I. Foster, R. Gardner, M. Wilde, A. Blatecky, J. McGee, and R. Quick, “The open science grid,” *J. Physics Conf. Series*, vol. 78, no. 1, pp. 12–57, 2007.
- [3] F. Gagliardi, B. Jones, F. Grey, M.-E. Bégin, and M. Heikkurinen, “Building an infrastructure for scientific grid computing: status and goals of the egee project,” *Philos. Transact. A Math Phys. Eng. Sci.*, vol. 363, no. 1833, pp. 1729–1742, 15 Aug. 2005.
- [4] D. Reed, “Grids, the TeraGrid and beyond,” *IEEE Computer*, vol. 36, no. 1, pp. 62–68, Jan. 2003.
- [5] M. Glick, A. Krishnamoorthy, and C. Schow, “Optics in the data center: Introduction to the feature issue,” *IEEE/OSA J. Opt. Commun. Netw.*, vol. 3, no. 8, p. OD1, Aug. 2011.
- [6] K. Binder and D. W. Heermann, *Monte Carlo Simulation in Statistical Physics: An Introduction*. Springer, 2010.
- [7] L. Luo, Q. Cao, C. Huang, T. Abdelzaher, J. Stankovic, and M. Ward, “EnviroMic: Towards cooperative storage and retrieval in audio sensor networks,” in *Proc. 27th Int. Conf. Distributed Computing Systems (ICDCS 2007)*, Toronto, Canada, 25–29 Jun. 2007, p. 34.
- [8] A. Bach, “The financial industry’s race to zero latency and terabit networking,” in *Proc. Optical Fiber Commun./National Fiber Optic Engineers Conf. (OFC/NFOEC 2011)*, Los Angeles, CA, USA, 8–10 Mar. 2011, keynote presentation. [Online]. Available: http://www.ofcnoec.org/osa.ofc/media/Default/PowerPoint/2011/Service-Provider-Summit-Keynote_Andrew-Bach.pptx
- [9] J. Erickson, S. Spence, M. Rhodes, D. Banks, J. Rutherford, E. Simpson, G. Belrose, and R. Perry, “Content-centered collaboration spaces in the cloud,” *IEEE Internet Comput.*, vol. 13, no. 5, pp. 34–42, Sep.–Oct. 2009.
- [10] W. Zhu, C. Luo, J. Wang, and S. Li, “Multimedia cloud computing,” *IEEE Signal Process. Mag.*, vol. 28, no. 3, pp. 59–69, May 2011.
- [11] I. Nave, H. David, A. Shani, Y. Tzruya, A. Laikari, P. Eisert, and P. Fechteler, “Games@large graphics streaming architecture,” in *Proc. IEEE Int. Symp. Consumer Electronics (ISCE 2008)*, Vilamoura, Portugal, 14–16 Apr. 2008, pp. 1–4.
- [12] E. Cooke and N. O’Connor, “Scalable virtual viewpoint image synthesis for multiple camera environments,” in *Proc. 9th Int. Conf. Information Visualisation (IV 2005)*, London, UK, 6–8 Jul. 2005, pp. 393–397.
- [13] V. Hamscher, U. Schwiegelshohn, A. Streit, and R. Yahyapour, “Evaluation of job-scheduling strategies for grid computing,” in *Proc. 1st IEEE/ACM Int. Workshop Grid Comput. (GRID 2000)*, ser. LNCS, vol. 1971, Bangalore, India, 17–20 Dec. 2000, pp. 191–202.
- [14] J. Yu and R. Buyya, “A taxonomy of scientific workflow systems for grid computing,” *ACM SIGMOD Rec.*, vol. 34, no. 3, pp. 44–49, Sep. 2005.
- [15] J.-P. Vasseur, M. Pickavet, and P. Demeester, *Network recovery - Protection and restoration of optical, SONET- SDH, IP, and MPLS*, D. Clarck, Ed. Morgan Kaufmann Publishers, 2004.
- [16] C. F. Lam, H. Liu, B. Koley, X. Zhao, V. Kamalov, and V. Gill, “Fiber optic communication technologies: What’s needed for datacenter network operations,” *IEEE Commun. Mag.*, vol. 48, no. 7, pp. 32–39, Jul. 2010.
- [17] V. Paxson and S. Floyd, “Wide area traffic: the failure of Poisson

- modeling,” *IEEE/ACM Trans. Netw.*, vol. 3, no. 3, pp. 226–244, Jun. 1995.
- [18] G. Terdik and T. Gyires, “Lévy flights and fractal modeling of internet traffic,” *IEEE/ACM Trans. Netw.*, vol. 17, pp. 120–129, Feb. 2009.
- [19] P. Abry, P. Borgnat, F. Ricciato, A. Scherrer, and D. Veitch, “Revisiting an old friend: on the observability of the relation between long range dependence and heavy tail,” *Telecommun. Syst.*, vol. 43, pp. 147–165, Apr. 2010.
- [20] A. Sang and S.-q. Li, “A predictability analysis of network traffic,” *Computer Networks*, vol. 39, no. 4, pp. 329–345, Jul. 2002.
- [21] J. Jung, B. Krishnamurthy, and M. Rabinovich, “Flash crowds and denial of service attacks: Characterization and implications for CDNs and web sites,” in *11th Int. Conf. World Wide Web (WWW 2002)*, 2002, pp. 293–304.
- [22] I. Foster, “What is the grid? a three point checklist,” *GRID today*, vol. 1, no. 6, pp. 32–36, 2002. [Online]. Available: <http://www-fp.mcs.anl.gov/~foster/Articles/WhatIsTheGrid.pdf>
- [23] L. M. Vaquero, L. Rodero-Merino, J. Caceres, and M. Lindner, “A break in the clouds: towards a cloud definition,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 1, pp. 50–55, Jan. 2009.
- [24] I. Foster, Y. Zhao, I. Raicu, and S. Lu, “Cloud computing and grid computing 360-degree compared,” in *Proc. Grid Computing Environments Workshop (GCE 2008)*, Austin, TX, USA, 16 Nov. 2008, pp. 1–10.
- [25] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, “Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the 5th utility,” *Futur. Gener. Comp. Syst.*, vol. 25, no. 6, pp. 599–616, 2009.
- [26] E. Ciepiela, D. Harezlak, J. Kocot, T. Bartynski, M. Kasztelnik, P. Nowakowski, T. Guba, M. Malawski, and M. Bubak, “Exploratory programming in the virtual laboratory,” in *Proc. Int. Multiconf. Computer Science and Information Technology (IMCSIT 2010)*, Wisla, Poland, 18–20 Oct. 2010, pp. 621–628.
- [27] P. Vicat-Blanc, S. Soudan, S. Figuerola, J. A. Garcia, Espin, J. Ferrer, E. Lopez, J. Buysse, X. Chen, A. Jukan, G. Landi, N. Ciulli, A. Tzanakaki, M. Brogle, L. van Laarhoven, B. Belter, F. Anhalt, E. Escalona, R. Nejabati, D. Simeonidou, Y. Demchenko, C. Ngo, C. de Laat, M. Biancani, M. Roth, P. Donadio, J. Jimenez, J. Kowalczyk, and A. Gumaste, *The Future Internet*, ser. LCNS. Springer, 2011, vol. 6656/2011, ch. Bringing Optical Networks to the Cloud: an Architecture for a Sustainable future Internet, pp. 307–320.
- [28] A. Lenk, M. Klems, J. Nimis, S. Tai, and T. Sandholm, “What’s inside the cloud? an architectural map of the cloud landscape,” in *Proc. 2009 ICSE Workshop on Software Engineering Challenges of Cloud Computing (CLOUD 2009)*, Vancouver, Canada, 2009, pp. 23–31.
- [29] L. Youseff, M. Butrico, and D. Da Silva, “Toward a unified ontology of cloud computing,” in *Proc. Grid Computing Environments Workshop (GCE 2008)*, Austin, TX, USA, 16 Nov. 2008, pp. 1–10.
- [30] C. Vecchiola, S. Pandey, and R. Buyya, “High-performance cloud computing: A view of scientific applications,” in *Proc. 10th Int. Symp. Pervasive Systems, Algorithms, and Networks (I-SPAN 2009)*, Kaohsiung, Taiwan, 14–16 Dec. 2009, pp. 4–16.
- [31] N. Wilkins-Diehr, D. Gannon, G. Klimeck, S. Oster, and S. Pamidighantam, “TeraGrid science gateways and their impact on science,” *IEEE Computer*, vol. 41, no. 11, pp. 32–41, Nov. 2008.
- [32] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, “OpenFlow: enabling innovation in campus networks,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Apr. 2008.
- [33] J. Berthold, A. a. M. Saleh, L. Blair, and J. M. Simmons, “Optical networking: Past, present, and future,” *IEEE/OSA J. Lightwave Technol.*, vol. 26, no. 9, pp. 1104–1118, May 2008.
- [34] E. Ip and T. Wang, “100g and beyond: Transmission technologies for evolving optical networks and relevant physical-layer issues,” *Proc. IEEE*, vol. 100, no. 1, Jan. 2012.
- [35] S. Ramamurthy, L. Sahasrabudde, and B. Mukherjee, “Survivable WDM mesh networks,” *IEEE/OSA J. Lightwave Technol.*, vol. 21, no. 4, pp. 870–883, Apr. 2003.
- [36] “Spectral grids for WDM applications: DWDM frequency grids,” ITU-T, Tech. Rep. G.694.1, June 2002.
- [37] A. Nag, M. Tornatore, and B. Mukherjee, Fellow, Ieee, “Optical network design with mixed line rates and multiple modulation formats,” *IEEE/OSA J. Lightwave Technol.*, vol. 28, no. 4, pp. 466–475, Feb. 2010.
- [38] M. Jinno, H. Takara, B. Koziicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, “Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies,” *IEEE Commun. Mag.*, vol. 47, no. 11, pp. 66–73, Nov. 2009.
- [39] A. N. Patel, P. N. Ji, J. P. Jue, and T. Wang, “Routing, wavelength assignment, and spectrum allocation in transparent flexible optical WDM (FWDM) networks,” in *Proc. Photonics in Switching (PS 2010)*. Optical Society of America, paper PDPWG1, 2010.
- [40] M. Jinno and Y. Tsukishima, “Virtualized optical network (VON) for agile cloud computing environment,” in *Proc. Optical Fiber Commun./National Fiber Optic Engineers Conf. (OFC/NFOEC 2009)*. San Diego, CA: Optical Society of America, 22–26 Mar. 2009, p. OMG1.
- [41] C. de Laat and L. Herr, “Ultra high definition media over optical networks (CINEGRID),” in *Proc. Optical Fiber Commun./National Fiber Optic Engineers Conf. (OFC/NFOEC 2009)*, San Diego, CA, 22–26 Mar. 2009, p. OWK1.
- [42] T. DeFanti, C. de Laat, J. Mambretti, K. Neggers, and B. St. Arnaud, “Translight: a global-scale lambda grid for e-science,” *Commun. ACM*, vol. 46, no. 11, pp. 34–41, Nov. 2003.
- [43] K.-I. Sato and H. Hasegawa, “Optical networking technologies that will create future bandwidth-abundant networks (invited),” *IEEE/OSA J. Opt. Commun. Netw.*, vol. 1, no. 2, pp. A81–A93, Jun. 2009.
- [44] D. Simeonidou, R. Nejabati, G. Zervas, D. Klonidis, A. Tzanakaki, and M. J. O’Mahony, “Dynamic optical-network architectures and technologies for existing and emerging grid services,” *IEEE/OSA J. Lightwave Technol.*, vol. 23, no. 10, pp. 3347–3357, Oct. 2005.
- [45] M. De Leenheer, P. Thysebaert, B. Volckaert, F. D. Turck, B. Dhoedt, P. Demeester, D. Simeonidou, R. Nejabati, G. Zervas, D. Klonidis, and M. J. O’Mahony, “A view on enabling-consumer oriented grids through optical burst switching,” *IEEE Commun. Mag.*, vol. 44, no. 3, pp. 124–131, Mar. 2006.
- [46] Y. Chen, C. Qiao, and X. Yu, “Optical burst switching: A new area in optical networking research,” *IEEE Network*, vol. 18, no. 3, pp. 16–23, 2004.
- [47] M. J. O. Mahony, D. Simeonidou, D. K. Hunter, and A. Tzanakaki, “The application of optical packet switching in future communication networks,” *IEEE Commun. Mag.*, vol. 39, no. 3, pp. 128–135, 2001.
- [48] R. S. Tucker, “Optical packet switching: A reality check,” *Opt. Switch. Netw.*, vol. 5, no. 1, pp. 2–9, 2008.
- [49] V. Chan, “Optical flow switching networks,” *Proc. IEEE*, vol. 100, no. 1, Jan. 2012.
- [50] G. Weichenberg, V. Chan, and M. Medard, “Design and analysis of optical flow-switched networks,” *IEEE/OSA J. Opt. Commun. Netw.*, vol. 1, no. 3, pp. B81–B97, Aug. 2009.
- [51] G. S. Zervas, M. De Leenheer, L. Sadeghioon, D. Klonidis, Y. Qin, R. Nejabati, D. Simeonidou, C. Develder, B. Dhoedt, and P. Demeester, “Multi-granular optical cross-connect: Design, analysis, and demonstration,” *IEEE/OSA J. Opt. Commun. Netw.*, vol. 1, no. 1, pp. 69–84, Jun. 2009.
- [52] M. De Leenheer, C. Develder, J. Buysse, B. Dhoedt, and P. Demeester, “Performance analysis and dimensioning of multi-granular optical networks,” *Opt. Switch. and Netw.*, vol. 6, no. 2, pp. 88–98, Apr. 2009.
- [53] H. L. Vu, A. Zalesky, E. Wong, Z. Rosberg, S. Bilgrami, M. Zukerman, and R. Tucker, “Scalable performance evaluation of a hybrid optical switch,” *IEEE/OSA J. Lightwave Technol.*, vol. 23, no. 10, pp. 2961–2973, oct. 2005.
- [54] G. Zervas, V. Martini, Y. Qin, E. Escalona, R. Nejabati, D. Simeonidou, F. Baroncelli, B. Martini, K. Torkmen, and P. Castoldi, “Service-oriented multigranular optical network architecture for clouds,” *IEEE/OSA J. Opt. Commun. Netw.*, vol. 2, no. 10, pp. 883–891, Oct. 2010.
- [55] X. Cao, V. Anand, and C. Qiao, “Waveband switching in optical networks,” *IEEE Commun. Mag.*, vol. 41, no. 4, pp. 105–112, 2003.
- [56] M. De Andrade, G. Kramer, R. Roy, and P. Chowdhury, “Evolution of optical access networks: Architectures and capacity upgrades,” *Proc. IEEE*, vol. 100, no. 1, Jan. 2012.
- [57] L. Kazovsky, S.-W. Wong, T. Ayhan, K. M. Albeyoglu, M. Ribeiro, and

- A. Shastri, “Hybrid optical-wireless access networks,” *Proc. IEEE*, vol. 100, no. 1, Jan. 2012.
- [58] F. Effenberger, D. Clearly, O. Haran, G. Kramer, R. D. Li, M. Oron, and T. Pfeiffer, “An introduction to PON technologies,” *IEEE Commun. Mag.*, vol. 45, no. 3, pp. S17–S25, Mar. 2007.
- [59] G. Wang, D. G. Andersen, M. Kaminsky, K. Papagiannaki, T. S. E. Ng, M. Kozuch, and M. Ryan, “c-Through: Part-time optics in data centers,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 4, pp. 327–338, 2010.
- [60] N. Farrington, G. Porter, S. Radhakrishnan, H. H. Bazzaz, V. Subramanya, Y. Fainman, G. Papen, and A. Vahdat, “Helios: A hybrid electrical / optical switch architecture for modular data centers,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 4, pp. 339–350, 2010.
- [61] C. Partridge, T. Mendez, and W. Milliken, “Host anycasting service,” Internet Engineering Task Force, United States, RFC 1546, Nov. 1993.
- [62] T. Stevens, M. De Leenheer, C. Develder, B. Dhoedt, K. Christodoulopoulos, P. Kokkinos, and E. Varvarigos, “Multi-cost job routing and scheduling in grid networks,” *Futur. Gener. Comp. Syst.*, vol. 25, no. 8, pp. 912–925, Sep. 2009.
- [63] W. Guo, W. Sun, Y. Jin, W. Hu, and C. Qiao, “Demonstration of joint resource scheduling in an optical network integrated computing environment,” *IEEE Commun. Mag.*, vol. 48, no. 5, pp. 76–83, May 2010.
- [64] I. Raicu, Y. Zhao, I. T. Foster, and A. Szalay, “Accelerating large-scale data exploration through data diffusion,” in *Proc. Int. Workshop on Data-aware Distributed Computing (DADC 2008)*, San Jose, CA, USA, 8 Jun. 2008, pp. 9–18.
- [65] C. Develder, B. Mukherjee, B. Dhoedt, and P. Demeester, “On dimensioning optical grids and the impact of scheduling,” *Photonic Netw. Commun.*, vol. 17, no. 3, pp. 255–265, Jun. 2009.
- [66] J. Buysse, M. De Leenheer, B. Dhoedt, and C. Develder, “Providing resiliency for optical grids by exploiting relocation: A dimensioning study based on ILP,” *Comput. Commun.*, vol. 34, no. 12, pp. 1389–1398, Aug. 2011.
- [67] C. Develder, J. Buysse, A. Shaikh, B. Jaumard, M. De Leenheer, and B. Dhoedt, “Survivable optical grid dimensioning: anycast routing with server and network failure protection,” in *Proc. IEEE Int. Conf. Commun. (ICC 2011)*, Kyoto, Japan, 5–9 Jun. 2011.
- [68] N. M. K. Chowdhury and R. Boutaba, “A survey of network virtualization,” *Comput. Netw.*, vol. 54, no. 5, pp. 862–876, Apr. 2010.
- [69] Y. Wang, Y. Jin, W. Guo, W. Sun, and W. Hu, “Virtualized optical network services across multiple domains for grid applications,” *IEEE Commun. Mag.*, vol. 49, no. 5, pp. 92–101, May 2011.
- [70] T. Takeda, D. Brungard, D. Papadimitriou, and H. Ould-Brahim, “Layer 1 virtual private networks: driving forces and realization by GMPLS,” *IEEE Commun. Mag.*, vol. 43, no. 7, pp. 60–67, Jul. 2005.
- [71] T. Anderson, L. Peterson, S. Shenker, and J. Turner, “Overcoming the internet impasse through virtualization,” *IEEE Computer*, vol. 38, no. 4, pp. 34–41, Apr. 2005.
- [72] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, “Xen and the art of virtualization,” in *Proc. 19th ACM Symp. Operating Sys. Principles (SOSP 2003)*, Bolton Landing, NY, USA, 19–22 Oct. 2003, pp. 164–177.
- [73] A. Kivity, Y. Kamay, D. Laor, U. Lublin, and A. Liguori, “kvm: the linux virtual machine monitor,” in *Proc. Linux Symposium*, vol. 1, Ottawa, Canada, 27–30 Jun. 2007, pp. 225–230.
- [74] J. Smith and R. Nair, “The architecture of virtual machines,” *IEEE Computer*, vol. 38, no. 5, pp. 32–38, May 2005.
- [75] R. J. T. Morris and B. J. Truskowski, “The evolution of storage systems,” *IBM Syst. J.*, vol. 42, no. 2, pp. 205–217, Apr. 2003.
- [76] P. Sempolinski and D. Thain, “A comparison and critique of Eucalyptus, OpenNebula and Nimbus,” in *Proc. 2nd IEEE Int. Conf. Cloud Computing Technology and Science*, Indianapolis, IN, USA, 30 Nov.–3 Dec. 2010.
- [77] [Online]. Available: <http://www.openstack.org>
- [78] C. Baun and M. Kunze, “The KOALA cloud management service: a modern approach for cloud infrastructure management,” in *Proc. 1st Int. Workshop on Cloud Computing Platforms (CloudCP 2011)*, Salzburg, Austria, 10 Apr. 2011, pp. 1:1–1:6.
- [79] D. Owens, “Securing elasticity in the cloud,” *Commun. ACM*, vol. 53, pp. 46–51, Jun. 2010.
- [80] M. De Leenheer, J. Buysse, C. Develder, and B. Mukherjee, “Isolation of virtual optical networks,” in *Submitted to Int. Conf. on Computing, Networking and Communications (ICNC 2012)*, Maui, Hawaii, USA, 30 Jan. – 2 Feb. 2012.
- [81] G. Chiruvolu, A. Ge, M. Ali, and J. Rouyer, “Issues and approaches on extending Ethernet beyond LANs,” *IEEE Commun. Mag.*, vol. 42, no. 3, pp. 80–86, Mar. 2004.
- [82] A. Doria, J. Salim Hadi, R. Haas, H. Khosravi, W. Wang, L. Dong, R. Gopal, and J. Halpern, “Forwarding and control element separation (ForCES) protocol specification,” IETF, Tech. Rep. RFC 5810, Mar. 2010.
- [83] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford, “Virtual routers on the move: live router migration as a network-management primitive,” in *Proc. ACM SIGCOMM 2008 Conf. Data Commun. (SIGCOMM 2008)*, Seattle, WA, USA, 14–22 Aug. 2008, pp. 231–242.
- [84] A. Jukan and G. Karmous-Edwards, “Optical control plane for the grid community,” *IEEE Commun. Surveys Tutorials*, vol. 9, no. 3, pp. 30–44, 3rd Quarter 2007.
- [85] G. Zervas, E. Escalona, R. Nejabati, D. Simeonidou, G. Carrozzo, N. Ciulli, B. Belter, A. Binczewski, M. Stroinski, A. Tzanakaki, and G. Markidis, “Phosphorus grid-enabled GMPLS control plane (G2MPLS): architectures, services, and interfaces,” *IEEE Commun. Mag.*, vol. 46, no. 6, pp. 128–137, Jun. 2008.
- [86] A. Willner, J. Ferrer Riera, J. A. Garcia-Espin, S. Figuerola, M. De Leenheer, and C. Develder, “An analytical model of the service provisioning time within the Harmony network service plane,” in *Proc. 2nd IEEE Int. Workshop on Management of Emerging Networks and Services (IEEE MENS 2010)*, Miami, FL, USA, 10 Dec. 2010.
- [87] T. Lehman, J. Sobieski, and B. Jabbari, “DRAGON: a framework for service provisioning in heterogeneous grid networks,” *IEEE Commun. Mag.*, vol. 44, no. 3, pp. 84–90, Mar. 2006.
- [88] L. L. Smarr, A. A. Chien, T. DeFanti, J. Leigh, and P. M. Papadopoulos, “The OptiPutter,” *Commun. ACM*, vol. 46, pp. 58–67, Nov. 2003.
- [89] X. Liu, C. Qiao, D. Yu, and T. Jiang, “Application-specific resource provisioning for wide-area distributed computing,” *IEEE Netw.*, vol. 24, no. 4, pp. 25–34, Jul.-Aug. 2010.
- [90] D. Andrei, M. Tornatore, D. Ghosal, C. U. Martel, and B. Mukherjee, “On-demand provisioning of data-aggregation sessions over WDM optical networks,” *IEEE J. Lightwave Technol.*, vol. 27, no. 12, pp. 1846–1855, 15 Jun. 2009.
- [91] P. Kokkinos, K. Christodoulopoulos, and E. Varvarigos, “Efficient data consolidation in grid networks and performance analysis,” *Future Gener. Comput. Syst.*, vol. 27, pp. 182–194, Feb. 2011.
- [92] E. M. Varvarigos, V. Sourlas, and K. Christodoulopoulos, “Routing and scheduling connections in networks that support advance reservations,” *Comput. Netw.*, vol. 52, no. 15, pp. 2988–3006, Oct. 2008.
- [93] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, “Energy efficiency in telecom optical networks,” *IEEE Commun. Surveys Tutorials*, vol. 12, no. 4, pp. 441–458, 2010.
- [94] J. Baliga, R. W. A. Ayre, K. Hinton, and R. S. Tucker, “Green cloud computing: Balancing energy in processing, storage and transport,” *Proc. IEEE*, vol. 99, no. 1, pp. 1–19, Jan. 2011.
- [95] A. Beloglazov and R. Buyya, “Energy efficient resource management in virtualized cloud data centers,” in *Proc. 10th IEEE/ACM Int. Symp. on Cluster Comput. and the Grid (CCGrid 2010)*, May 2010, pp. 826–831.
- [96] S. Figuerola, M. Lemay, V. Rejjs, M. Savoie, and B. St. Arnaud, “Converged optical network infrastructures in support of future internet and grid services using IaaS to reduce GHG emissions,” *IEEE/OSA J. Lightwave Technol.*, vol. 27, no. 12, pp. 1941–1946, Jun. 2009.
- [97] W. Van Heddeghem, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, “Distributed computing for carbon footprint reduction by exploiting low-footprint energy availability,” *Futur. Gener. Comp. Syst.*, vol. 28, no. 2, pp. 405–414, Feb. 2012.
- [98] K. Keahy, R. Figueiredo, J. Fortes, T. Freeman, and M. Tsugawa, “Science clouds: Early experiences in cloud computing for scientific applications,” in *Proc. Cloud Computing and Its Applications (CCA 2008)*, Chicago, IL, USA, 22–23 Oct. 2008.
- [99] J. Napper and P. Bientinesi, “Can cloud computing reach the top500?”

- in *Proc. Combined workshops on UnConventional high performance computing workshop plus memory access workshop (UCHPC-MAW 2009)*, Ischia, Italy, 2009, pp. 17–20.
- [100] S. Ostermann, A. Iosup, N. Yigitbasi, R. Prodan, T. Fahringer, and D. Epema, “A performance analysis of EC2 cloud computing services for scientific computing,” in *Proc. 1st Int. Conf. Cloud Computing (CloudComp 2009)*, 2009, pp. 115–131.
- [101] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, “Above the clouds: A berkeley view of cloud computing,” EECS Department, University of California, Berkeley, CA, USA, Tech. Rep. UCB/EECS-2009-28, Feb. 2009.
- [102] B. Jennings, R. Brennan, W. Donnelly, S. Foley, D. Lewis, D. O’Sullivan, J. Strassner, and S. van der Meer, “Challenges for federated, autonomic network management in the future internet,” in *Proc. 1st IFIP/IEEE Int. Workshop on Management of the Future Internet (ManFI 2009)*, Long Island, NY, USA, 5 Jun. 2009, pp. 87–92.
- [103] L. Baresi, A. D. Ferdinando, A. Manzalini, and F. Zambonelli, “The CASCADAS framework for autonomic communications,” in *Autonomic Communication*, A. V. Vasilakos, M. Parashar, S. Karnouskos, and W. Pedrycz, Eds. Springer US, 2009, pp. 147–168.
- [104] J. Strassner, J.-K. Hong, and S. van der Meer, “The design of an autonomic element for managing emerging networks and services,” in *Proc. Int. Conf. Ultra Modern Telecommun. (ICUMT 2009)*, Moscow, Russia, 18–20 Oct. 2009, pp. 1–8.