InterGrid: A Case for Internetworking of Islands of Grids

Marcos Dias de Assunção and Rajkumar Buyya

Grid Computing and **D**istributed **S**ystems (GRIDS) Laboratory
Department of Computer Science and Software Engineering
The University of Melbourne, Australia
Email: {marcosd, raj}@csse.unimelb.edu.au
http://www.gridbus.org

Abstract: Over the last few years, several Grids have been set up to share resources such as computers, data, and instruments to enable collaborative research. These Grids follow models and heterogeneous policies restricted by the requirements of e-Science applications for which they have been created, which has resulted in islands of Grids. A structure that allows Grids to grow in a scalable manner by allowing peering between Grids is still not possible. There has been a profound interest in providing interoperability between different Grids. In this paper, we propose a model that (a) promotes interlinking of islands of Grids through peering arrangements to enable inter-Grid resource sharing; (b) provides a scalable structure for Grids that allow them to interconnect with one another and grow in a sustainable way; (c) creates a global cyberinfrastructure to support e-Science and e-Business applications. This work identifies the need for a business ecosystem, called the InterGrid, which promotes the internetworking of islands of Grids and allows Grids to grow in a similar way as the Internet. We examine current global systems that grow by enabling networks of networks and propose an architecture for the InterGrid. The architecture is composed of InterGrid Gateways responsible for managing peering arrangements between Grids. We discuss the main components of the architecture and present a research agenda to enable the InterGrid vision.

Keywords: Grid computing, cyberinfrastructure, internetworking, InterGrid, decentralized resource management.

1. Introduction

The growing popularity of Internet-based communication, computing, storage and software technologies has led to the emergence of the Grid computing paradigm that allows secure and coordinated sharing of globally distributed resources. Grid computing supports a range of e-Science and e-Business applications [1-3]. Its ultimate goal is the creation of a cyberinfrastructure that allows scientists and practitioners to cope with the scale and complexity of both current and next-generation scientific challenges [4-6]. Toward this, various national programs have initiated e-Science projects to enable resource sharing and collaboration among scientists. Such endeavors generally follow a restricted organizational model based on the idea of Virtual Organizations (VOs) [7]. These models and dispersed Grid initiatives have resulted in islands of Grids without resource sharing between them.

Interlinking of islands of Grids is needed to provide a global Grid-based cyberinfrastructure [8]. Nevertheless, much beyond the need for interoperability at middleware level, interlinking of Grids requires peering arrangements, advanced and automated mechanisms for inter-Grid resource allocation, reservation, accounting, and scheduling. In addition, Grids need to adopt mechanisms that enable administrative separation, by allowing networks of networks, similar to many network-based systems such as the Internet, the Web, and numerous social and biological systems [9-12]. Currently, however, Grids follow organizational models that stop them from peering with one another. In other words, the structure of current Grids does not follow principles such as the peering between Internet Service Providers (ISPs) present in the Internet. Therefore, likewise the peering of ISPs, we need to investigate policies for interlinking of Grids and how the peering arrangements between Grids will be made.

There are currently several Grid facilities across the world and efforts to promote interoperability between them. Examples include Enabling Grids for E-science in Europe (EGEE) [13] and Open Science Grid [14]. There exist also several national and international initiatives such as TeraGrid [15, 16], APAC (the Australian Grid) [5], K*Grid in Korea [17], NAREGI in Japan [18], OurGrid in Brazil [19], among others [20, 21] that aim to provide national Grid facilities. Grid middleware, such as Globus [22], provides a means for uniform access to resources in these Grids, but do not provide much in means to interlink them. Furthermore, many Grid systems require users to support the same types of middleware to interoperate. Although efforts have taken place towards supporting interoperability between different Grid middleware systems (such as between Globus and UNICORE [23], and OSG and EGEE), much of this work continues to focus on translating between messaging formats and data models instead of dealing with policy issues. Efforts at the Open Grid Forum, such as "Grid Interoperability Now" (GIN) are also focusing on interoperability between Grids. However, the current work on interoperability does not focus on peering arrangements between Grids. Interoperability and common protocols are important, but not enough to promote interlinking of islands of Grids. For example, a set of common communication protocols underlies the Internet, but when ISPs peer with one another they consider their policies, economic issues, social and economic impact when they peer with one another. It is therefore important to identify the key issues of current Grid technologies that do not allow them to evolve to such an InterGrid level. The questions that be answered are:

- What are the architectural issues that prevent current Grid architectures to scale to the InterGrid?
- What kind of structure should the InterGrid have to promote internetworking of Grids?
- What kind of ecosystem, peering arrangements and policies do we need to have?
- What are the coordination mechanisms that we need to put in place to enable the InterGrid?
- What are the incentives for end users, laboratories, organizations, service providers, and Grid facilities in general to engage in such a network of Grids?
- Should we consider macroeconomics-based, bilateral, and multilateral agreements when linking such Grids?

Therefore, there is a need for a business ecosystem and means that take into consideration barriers in order to support internetworking (interlinking or peering) between islands of Grids in a decentralized manner. Internetworking in this context refers to peering arrangements and mechanisms for inter-Grid resource allocation, automated resource reservation, interconnection of information services and accounting, and cross Grid scheduling, which we term as InterGrid. Internetworking of islands of Grids would lead to the creation of networks of Grids. This would allow participants to allocate resources from different islands of Grids in a seamless manner, and permit peering between Grids under different administrative policies and political boundaries.

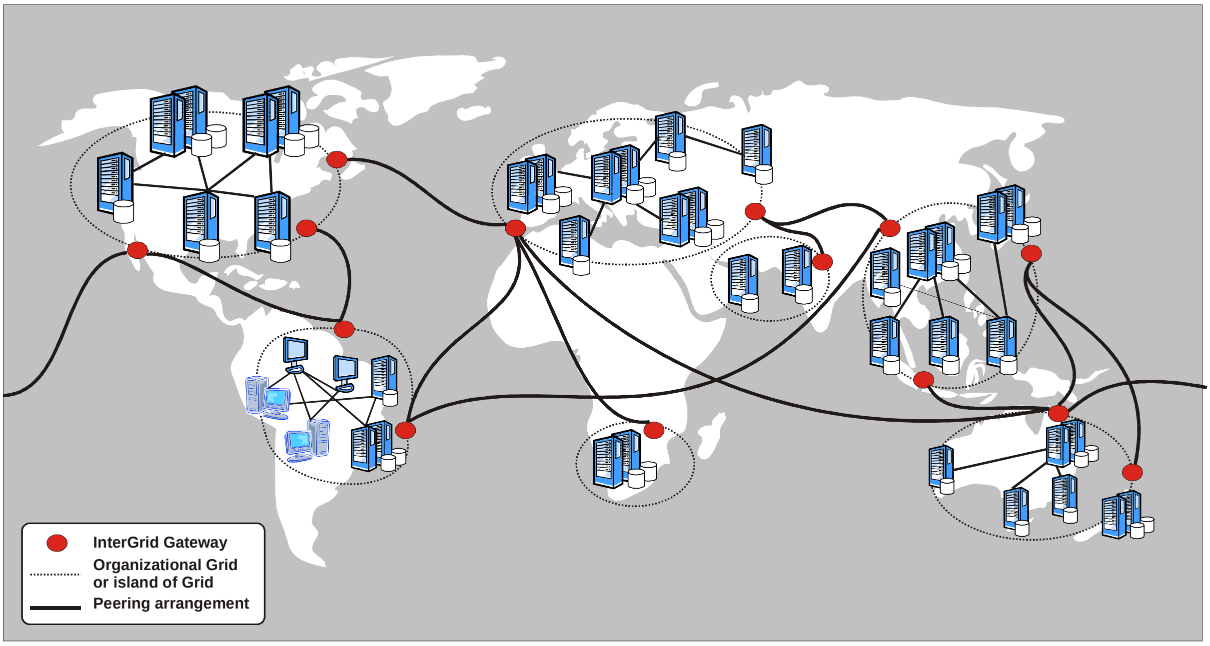

Figure 1. Abstract view of the InterGrid as an evolvable ecosystem that can grow from local Grids to a unique global Grid, while maintaining local autonomy.

In a similar fashion to the Internet, the InterGrid must be an evolvable system that can expand from organizational Grids to complex structures without major problems or scalability limitations. An abstract view of the InterGrid, presented in Figure 1, shows that it comprises of multiple islands of Grids and has a structure that grows from isolated Grids to a unique global ecosystem. Such islands of Grids interlink with one another through peering arrangements. It should not be difficult for an organization or a new service provider to join the InterGrid, without the need of complex and multilateral negotiations involving many organizations, or having to join multiple VOs.

To enable Grid internetworking, however, it is important to draw lessons from network-based systems, including the Internet, the Web and Content Delivery Networks (CDNs) [24]. It is also important to learn from these network-based systems by analyzing their patterns and characteristics, find out what can be applied to Grid computing. It is important to analyze how they have evolved and grown without centralized control and how they self-organize [10]. The Internet has rapidly grown by interconnecting several scales of networks, which use different network technologies. It allows the interlinking of Internet Service Providers (ISPs) who peer and exchange traffic based on several factors in order to provide Internet services to their clients. The structure of the

Internet has demonstrated that it does not impose serious limitations on its growth. Likewise, other examples of sustainable systems exist. The Web enables the exchange of information on a global basis, allows access to any Web page stored on any Web server in any part of the globe and gives rise to a range of business models [25, 26]. An analysis of global infrastructures is provided in this work.

Therefore, the main contributions of this paper are to:

- Propose a model for the InterGrid as an ecosystem that allows the internetworking of islands of Grids to cope with a range of next generation challenging tasks in business and scientific areas.
- Investigate successful global infrastructures such as the World Wide Web (WWW) and the Internet, and identify key aspects that influence their growth.
- Develop an architecture for the InterGrid which allows peering between islands of Grids, with possible policies and incentive mechanisms that can be used to ensure its sustainability.
- Propose a research agenda by identifying several challenges that need to be addressed, such as policy-based peering of Grids, pricing of Grid resources, coordination among Grids, enabling feasible market models, infrastructure for Grid economics, integration of accounting systems, automated resource reservation, inter-Grid resource allocation, among others. Possible solutions to some of these challenges are also necessary to ensure the adoption of Grid computing within industry.

The rest of the paper is organized as follows. Section 2 contains a description and analysis of the structure and principles that form the basis of the Internet and the WWW. Section 3 presents a gap analysis of existing Grid systems. We then present the structure of the InterGrid and the proposed architecture in Section 4. Section 5 then discusses the proposed research agenda on this topic. Section 6 concludes the paper and presents our final considerations on the subject.

2. Global Infrastructures and their Properties

As pointed out by Smarr [12], many infrastructures for current well-known services evolved from isolated initiatives that were connected and put together. In this section, we examine existing infrastructures and draw some lessons from them for building the InterGrid.

2.1 The Internet

The Internet has grown from a small project from DARPA started in 1969, initially linking a few sites in USA. Currently, millions of hosts and networks compose the intricate topology of Internet. An abstract view of this intricate topology is presented in Figure 2a while the interconnection between ISPs is shown in more details in see Figure 2b. The figure shows that hosts are connected to local Internet Service Providers (ISPs) through access networks. In dial-up or broadband services, the local PSTN (Public Switched Telephone Network) loop is commonly used to provide users with access to the Internet. These local ISPs then connect to regional ISPs, which in turn, connect to national and international ISPs, also known as Tier-1 ISPs. Such national and international ISPs, also called National Service Providers (NSPs), represent the highest level of the Internet hierarchy and are connected to each other either directly or through Network Access Points (NAPs), also known as Internet Exchange (IX). Thus, the ISPs can provide services like access, backbone, content, application, and hosting. This structure has allowed the Internet topology to grow quickly and without the endorsement of a central authority [27].

Currently, the Internet presents an intricate structure comprised of a vast number of physical connections established by commercial contracts such as peering agreements. Such agreements are legal contracts, which specify the details of how ISPs exchange traffic. Norton [28] highlights the difference between peering and transit. Peering is the relationship whereby ISPs provide connectivity to each other's transit customers. Transit on the other hand, is the relationship through which one ISP provides access to all destinations in its routing table. The reasons for peering involve social, economical, and technological factors. ISPs can consider their policies, economical advantages and conflicts before establishing agreements. However, there is no common routine for choosing with whom to peer. Such agreements can be of various types, such as private, via IXs or in a relationship between customer and provider. They can specify the amount and proportion of traffic exchanged and the settlements since the traffic between peering ISPs can be asymmetric. Tier-1 providers, also known has having access to the global Internet, generally establish contracts not charging other Tier-1 providers, whereas charge for peering with smaller ISPs.

Another important concept is that of an Autonomous System (AS). An AS in general comprises a network under a single administration and has its own policies to divert traffic or to avoid some peering ASs on the Internet. These policies are enabled by routing protocols such as Border Gateway Protocol (BGP), which allow ASs to advertise the routes that they prefer. An AS can have policies that take into account shortest or most cost-effective paths. Such policy-based routing or peering is also applicable to the InterGrid, where Grids can favor a peering Grid more than others.

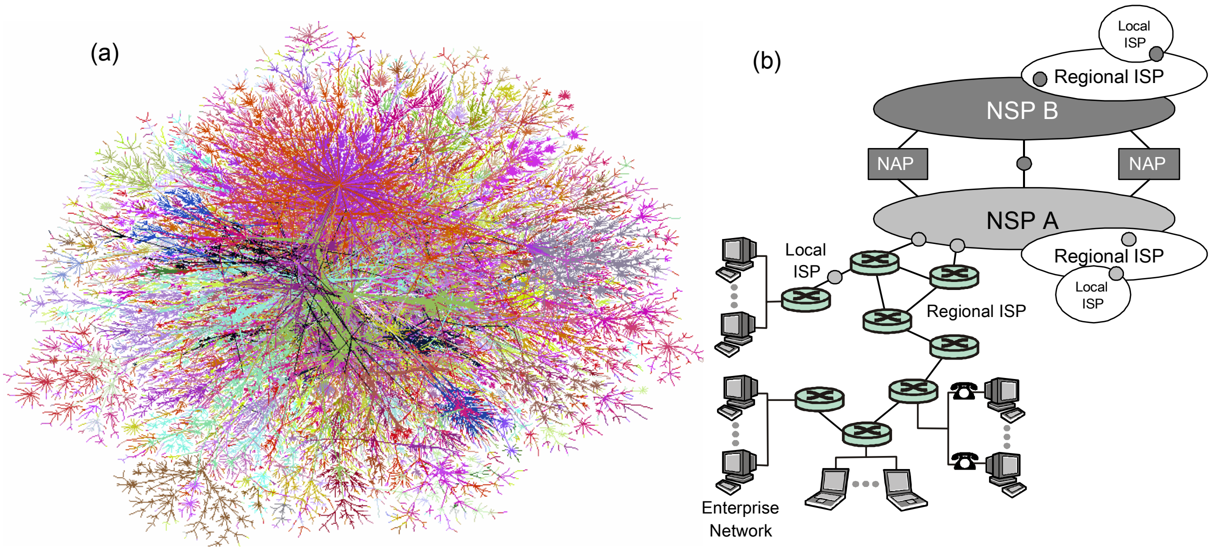

Figure 2. (a) Abstract view of the Internet topology. (b) Interconnection of ISPs.

Some lessons can be learnt by analyzing the structure of the Internet and how it has grown, such as:

- The use of a common set of simple protocols that allow the interoperability among networks with different technologies;
- Even though the Internet has a complex topology, it has a structure that can grow quickly because there is no need of agreements and negotiations involving multiple organizations; a host in the Internet does not need to be directly connected to a large number of networks in order to have access to hosts in other networks;
- A self-healing structure, in which the failure of part of the network does not compromise the whole Internet;
- Although ISPs compete with one another, peering allows peered ISPs to reduce the amount of traffic across an expensive boundary and improve the efficiency for their users [28]. In addition, its business model benefits end-users and compensates service providers;
- Routing protocols that allow traffic to be diverted when it is not allowed or viable to cross a specific network; these protocols allow ASs to deploy a range of routing policies based on internal interests;
- Networks are linked through routers in the Internet, therefore forming a large network of networks.

2.2 The World Wide Web and Content Delivery Networks

The World Wide Web (WWW) is one of the major network applications that contributed towards the rapid growth of the Internet [29]. Currently, the Web is a merge of network, protocols, and hypertext, which has led to the development of a plethora of scientific and commercial applications and several business models [26].

Although the WWW provides a global system that allows the sharing of several kinds of media on top of the Internet, there have been concerns about the performance and quality of the content delivered to Web users. Content Delivery Networks (CDNs) address some of these issues by providing networks of surrogate servers, which maintain replicas of Web content in strategic locations. Each CDN is set up and operated by providers such as Akamai [24] and Mirror Image [30]. In addition, several content providers have been setting up their own CDNs. However, there is no interaction and peering among these privately operated CDNs. Hence, CDNs can only grow to a certain extent as there are economic and technical reasons to stop a CDN from covering specific regions or maintaining complex and expensive infrastructures [31]. Content internetworking (CDI) [32, 33] through the peering between CDNs, is a solution to this problem as CDN providers can now cover broader areas and minimize their costs with infrastructures. However, CDI poses challenges in defining protocols and policies for internetworking of accounting systems, content distribution and request routing. The challenges imposed by CDI have some similarities to the challenges in the internetworking of islands of Grids. Therefore, we can draw lessons from such endeavors, such as the integration of the accounting systems, the decentralized allocation of resources, and the settlements among CDNs.

Some lessons to be drawn from the WWW and CDNs are as follows:

- The presence of network effect [34] and a cycle that allows the growth of both the Web and its underlying network infrastructure. The size of the Web is able to grow even faster with the increase in amount of information available on the Web, the number of users and servers accessing them and the number of applications.

- The use of ubiquitous protocols.
- No central control.
- It is simple to use.
- By peering, CDNs can share the cost of expensive infrastructure, reach broader areas, and provide a better QoS to their clients by offloading requests to peering CDNs.

2.3 Peer-to-Peer (P2P) Networks

P2P networks allow the sharing of compute resources such as content, storage and CPU cycles without the need of a centralized server or authority [35]. Such kinds of networks generally maintain their own network mechanisms such as addressing, connection and routing, which overlay the Internet. P2P technologies have given good technical insights into Grid computing. P2P has been compared to Grid [36] and its outcomes have been used in some Grid computing technologies [37].

Some aspects of the P2P network technology that are significant due to their relevance to Grid computing are:

- The decentralized architecture and control;
- Overlay addressing and protocols for data replication and resource discovery;
- The self-organizing characteristics;
- Overlay network-oriented approaches that build redundant and fault-tolerant networks on top of the Internet;
- They are in general easy to join. No complex negotiations are required for a new member to join such networks.

3. Issues in Current Grid Systems

When considering a large-scale system such as the InterGrid, a number of challenges arise, such as resource management among different islands of Grids, varying usage and connectivity patterns, different security policies, resource reservation, QoS and SLAs, and formation and management of VOs. In addition, users and providers need incentives to participate in the InterGrid. This section discusses some challenges that need to be addressed in order to realise the InterGrid.

Peering Arrangements: In the Internet, although there are standard protocols, ISPs have policies that define how the peering with other ISPs is performed. ISPs have agreements and implement the peering policies by defining what routes are preferable by considering the economic impact and incentives of peering with other ISPs. Work in Grid has focused on interoperability, but not on the peering between Grids and its economic implications.

Standards: As presented in [38], there are two ways to adopt standards, such as OGSA, for Grids. The first way is to make every single Grid service OGSA compliant. However, there have been different interpretations of OGSA. The second way is to have all service providers providing a standards compliant interface externally, while using their own protocols and interfaces internally. The Grid community has then to agree on a simple set of interfaces that is widely acceptable and easy to implement. Mapping of information through gateways and federations will be necessary until the Grid community does not agree on common interfaces.

Different policies and mechanisms for resource allocation: Besides the interoperability between Grids at the middleware level, interlinking Grids requires advanced and automated mechanisms for inter-Grid resource allocation, reservation, accounting, and scheduling. In addition, issues regarding different and divergent policies for resource usage need to be addressed. Grids follow different policies and mechanisms for resource allocation; the policies and mechanisms implemented by a Grid can be different and incompatible with those adopted by another. This makes difficult to conciliate policies or to create mechanisms to map policies from a Grid to another. An agreement on the standard units for resource usage is also required for allocation or exchange of inter-Grid resources.

Protectionism and collusion: In an open global Grid, local protectionism problems appear. For example, scientists in the Country A can be interested in developing a new technology or running a given set of applications. For doing so, they will look for Grid partners and use resources from the InterGrid. However, Country B has interests in slowing down the development of such technology by Country A. The reasons for that can be: (a) it is also developing similar technology; (b) it considers Country A as a threat; (c) any other political reasons. Therefore, Country B imposes local barriers for using its resources and tries to persuade other countries to do the same. Such political, financial, and cultural issues are not solved in a global Grid scenario. Issues such as collusion and formation of groups to reduce competition should be investigated as well as the development of mechanisms to address such issues. As argued in [39], the potential non-cooperativeness should be modeled and studied at various levels in Grid computing. As protocols such as BGP allow the definition of complex policies by ASs to avoid given networks similar policies may be needed in the InterGrid to allow Grids to avoid other Grids in the network.

Incentives for collaboration and attracting service providers: In the Internet and WWW there are national level projects for deploying relevant technologies to improve usage among the population. The InterGrid needs to provide ways to attract resource providers. Thus, a number of approaches have proposed using economic models to address resource usage and incentives in a Grid [40]. However, in the InterGrid, this approach requires additional globally accessible services such as a Grid bank, a common currency or token exchange mechanism, and the Grids involved need to trust these entities. Either government institutions will be required to fund and maintain such globally available entities, or decentralized mechanisms will be needed.

Pricing of resources and estimation of requirements: Economic models are important to Grid computing because: (i) Grid resources are not shared for free and charging for their usage can provide incentives for resource providers to offer and share their resources in the Grid; (ii) the participants of the Grid can be divided into resource consumers and providers, and resource allocation is achieved through the economic behavior of these actors; (iii) markets can offer a decentralized scheduling approach whereby each participant acts to maximize its own utility (iv) market-based resource allocation provides incentives for users to truthfully reveal how much they value resources [41]. However, another important concern is how resources should be priced and how usage is measured. What would be the basic units of usage for a compute or a storage resource? How do resource providers adjust the price of their resources in a competitive Grid? What are the different possible price mechanisms in a Grid market, considering the local pricing, competitive market and collaboration among Grids? How do the price mechanisms affect the system? How do Grid users and organizations estimate their needs for resources? All these questions need to be answered before economic models can be applied to Grid computing successfully.

Connectivity and interaction patterns: The integration of Grids can enable a large number of interaction patterns, which would be difficult to design in terms of middleware, scheduling, and resource allocation. It is advocated that overlay networks will be important in a large-scale Grid to tackle this heterogeneity and guarantee several interaction patterns [42]. Overlay networks are virtual networks that cover physical infrastructures such as the Internet and add value to them with some features and semantics. They can enable various interaction models through Application Programming Interfaces (APIs) to abstract the middleware from the complexity of the underlying network.

Coordination mechanisms: As mentioned in [43], current approaches to resource allocation are non-coordinated. Such approaches can lead to inefficient schedules and worse resource utilization. Coordination mechanisms that allow brokers and resource management systems to exchange information need to be put in place. However, the main challenge is that the InterGrid has Grids with different connectivity patterns. Thus, questions to be answered are: what metaphors should coordination mechanisms follow?; and how can current mechanisms be improved to satisfy the InterGrid's requirements*?*

3.1 Issues in Virtual Organizations (VOs)

Grid computing is also defined as the coordinated resource sharing and problem solving in dynamic and multi-institutional VOs [7]. A VO can be composed of a group of individuals and/or institutions that come together to share resources with a common purpose. According to [44], the life cycle of VOs can be divided into (a) the identification of business opportunities that require VOs to be formed; (b) their formation; (c) their operation and management; and (d) their termination. However, some problems can arise when considering these steps:

Formation of VOs: Currently, organizations define the terms for formation of VOs through multilateral contracts and agreements. Such processes are done in an off-line basis. It is not possible to create VOs in an on-demand and dynamic manner due to security and policy related issues. There is also a lack of mechanisms for the negotiation and establishment of agreements to dynamically form VOs. Moreover, a framework for how the off-line and out-of-the band agreements are defined for composing the source network or physical infrastructure is required. In addition, some legal barriers for the formation of VOs exist; some nations impose restrictions and require detailed information on the nature of collaboration with scientists from other countries. There is, thus, a requirement for change in laws and legal processes for the establishment of such VOs.

Resource management in VOs and across VOs: Providing a fair resource allocation in VOs is troublesome since resource providers can subscribe to multiple VOs and provide different amounts of resources to different VOs. Meta schedulers [39], some taking into account VOs, have been proposed. Dumitrescu *et al* [45], highlight that challenging usage policies can arise in VOs that comprise participants and resources from different physical organizations. Participants want to delegate access to their resources to a VO, while maintaining such resources under the control of local usage policies. In this context, they seek to address questions such as: how usage policies are enforced at the resource and VO levels? What mechanisms are used by a VO to ensure policy enforcement? How the distribution of policies to the enforcement points is carried out; and how policies are made available to VO job and data planners. In [45], they have proposed a policy management model in which participants can specify the maximum percentage of resources delegated to a VO. A VO in turn can specify the maximum percentage of resource usage it wishes to delegate to a given VO's group. However, such policies are defined in an off-line basis and are complex to reconcile.

We believe that resource allocation in static and dynamic VOs could use the metaphor of a corporation. Shareholders that hold the most shares have the right to take decisions regarding how resources are allocated in the VO. The decision taker is chosen in the formation of the VO or as the VO evolves. However, it is important to have one who plays the role of accounting and ethic committee to avoid abuse in the VO. In addition to these problems, VOs can be recursive. That is, a VO can be composed of multiple sub-VOs. Resource allocation among these VOs has to account for the problems previously described, in addition to the allocation problems across these VOs. Regardless of the approach, if using VOs, then inter-VO resource management should be considered because an organization can involve multiple VOs to cope with several problems. However, this imposes challenges regarding the scalability of the VO approach.

Security in VOs: At present, Grid Security Infrastructure (GSI) provides the basis for security in the Grid. At the VO level, VOMS [46] offers support to manage users, groups, roles, and capabilities in VOs. They allow a centralized control of VOs and extend Grid security concepts to a VO level by proving additional services which include: (a) VOMS server that maintains information about users, groups they belong to, roles and permissions; (b) a client that allows the user to create a VOMS proxy certificate; and (c) a VOMS administration service that allows the manager of the VO to setup roles and capabilities. There has been a few works on automated generation and negotiation of access control policies in VOs [47]. However, such security models have to deal with ad hoc Grids and short-lived VOs. Issues regarding the negotiation of access control policies and mapping of existing privileges from a source domain to a target domain have been investigated [48]. However, efforts are still necessary in this area in order to make the Grid a robust infrastructure for commercial applications. In addition, public Grids require additional mechanisms to ensure the trust and security at higher levels like business processes. Thus, reputation based mechanisms are required and key ideas from trust and reputation in P2P networks are applicable.

In the next section, we present the InterGrid architecture that is designed to overcome the above limitations in Grids and promote the establishment of peering arrangements between them.

4. The InterGrid Architecture

4.1 Network of Networks Structure

Through the investigation of existing infrastructures, we note that the concept of network of networks presented by the Internet is missing in Grid computing. In addition, the Internet aims at simplicity and providing a common set of protocols; the Grid is becoming a very complex architecture. Self-healing and benefits from peering, such as reducing traffic, increasing revenues or using services, are reasons adopted by ISPs for peering with one another. From the Web, we can see that the lack of centralized control has allowed for its fast growth. The Web has enabled a range of business models and organizations have reasons for using it. Self-organization, self-healing and decentralization are characteristics of P2P overlay networks that should be incorporated into Grid networks. CDNs can peer to cover a broader area and share costs for an expensive infrastructure thus avoiding over provisioning of resources and minimizing cost.

Based on communities and groups in our society and how they have formed, we see that such structures evolve from locally organized structures to those that are more complex. For example, a group of individuals has a common interest on a given activity. Leaders of this group can look for another similar group and may find it useful to interact with another. After the agreement to co-operate has been settled, interactions can take place, new links can be made, and existing ones can be broken. Some tools have helped people to form communities of interests.

From the different structures analyzed, we can note the following characteristics and needs:

- Small structures are linked to more complex ones through some access point. Examples are in the Internet, routers link networks, in groups, leaders start agreements or collaborations with other groups;
- In joining and forming communities, there are places where people publish not just their capabilities, but also their interests and needs;
- Mechanisms allow one to locate and connect with people or organizations that can fulfill their needs.

4.2 Architectural Requirements

The problem we are coping with here is how islands of Grids can coordinate through peering arrangements to provide a structure that enables the growth of the InterGrid. InterGrid, thus, needs to meet the following requirements:

Incentive-oriented peering arrangements: Although the InterGrid is comprised of islands of Grids with competing organizations with different and competing interests, it will need to provide mechanisms that provide incentives for Grids to peer with one another. Likewise the peering between ISPs, the InterGrid have their how policies for peering and have incentives for trading and using resources of one another.

Standards based: It is important that the research community agree on a small set of standard protocols; these would ensure interoperability at the middleware level and allow automated decisions and policies to be implemented regardless the middleware or low level tools utilized by different Grids. Islands of Grids are then linked together by using an agreed upon architecture and a set of common interfaces such as Open Grid Service Architecture (OGSA) and Web Services [49].

Respect to administrative management and separation: Grids are under different administration, have different resource usage policies and are set up for different objectives. It is important to respect both the internal policies and provide a structure that allows Grids to interlinking whilst respecting the concept of organization. In other words, the concept of networks of Grids is followed.

Deployment of applications that require resources from multiple Grids: Users can submit applications and scientific workflows that require resources beyond the capacity of their Grids, thus requiring inter-Grid resource allocation to be performed. In this scenario, standard protocols are required to allow both the inter-Grid reservation and allocation of resources and the submission of jobs to these resources.

Inter-Grid policies and decentralised resource management: A Grid requires a means to specify its policies, defining which resources are available to other Grids under which circumstances, and make then available to other Grids. However, a Grid will retain ultimate control over its resources and to whom it wants to provide access. We need to implement peering agreements between Grids without the need on global control over the resources. Decentralized approaches for resource allocation such as self-organizing economic models [50, 51] are required by the InterGrid.

Resource allocation, reservation and brokering across Grids: Each Grid has gateways that translate and reply to resource allocation requests originated in local and peering Grids. A gateway, aware of the peering arrangements with other Grids, forwards the request to other Grids able to provide the required resources. We need policies and mechanisms for selecting Grid able to provide the resources and for admission control when accepting requests originated in other Grids. In addition, the InterGrid requires coordination mechanisms between gateways and means for reconciling policies or mapping policies from a Grid to another.

4.3 The Proposed Architecture

Our architecture for InterGrid is show in Figure 3. Similar to the Internet, each Grid establishes peering arrangement with its neighboring or associated Grids depending on their peering policies. This peering arrangement will be managed by components called InterGrid Gateways (IGGs). Some of the elements in this architecture are discussed below.

Grid Resource Provider (RP): A RP is responsible for providing resources in the IntraGrid and to users from other Grids. To respect the idea of administrative separation, it would be interesting to minimize the number of Grids to which a resource provider can subscribe.

IntraGrid Resource Manager (IRM): An IRM is responsible for the allocation and management of resources in an IntraGrid and uses local protocols in order to communicate with resources provided by RPs to an IntraGrid. Different IRMs can use different policies to allocate Grid resources. Examples of technology to enable such IRMs include Shirako [52], Virtual Grids [53], VioCluster [54] and Virtual Workspaces [55].

Grid Resource Broker (GRB): Clients wanting to allocate Grid resources can utilize a GRB to do so. The GRB uses resources from the Grid and InterGrid when demand surpasses the resources its Grid can offer. It should then interact with IRM in order to obtain access to resources from other islands of Grids. Examples of existing brokers that can be adapted to such a scenario include Gridbus Resource Broker [56] and Service Manager of Shirako [52].

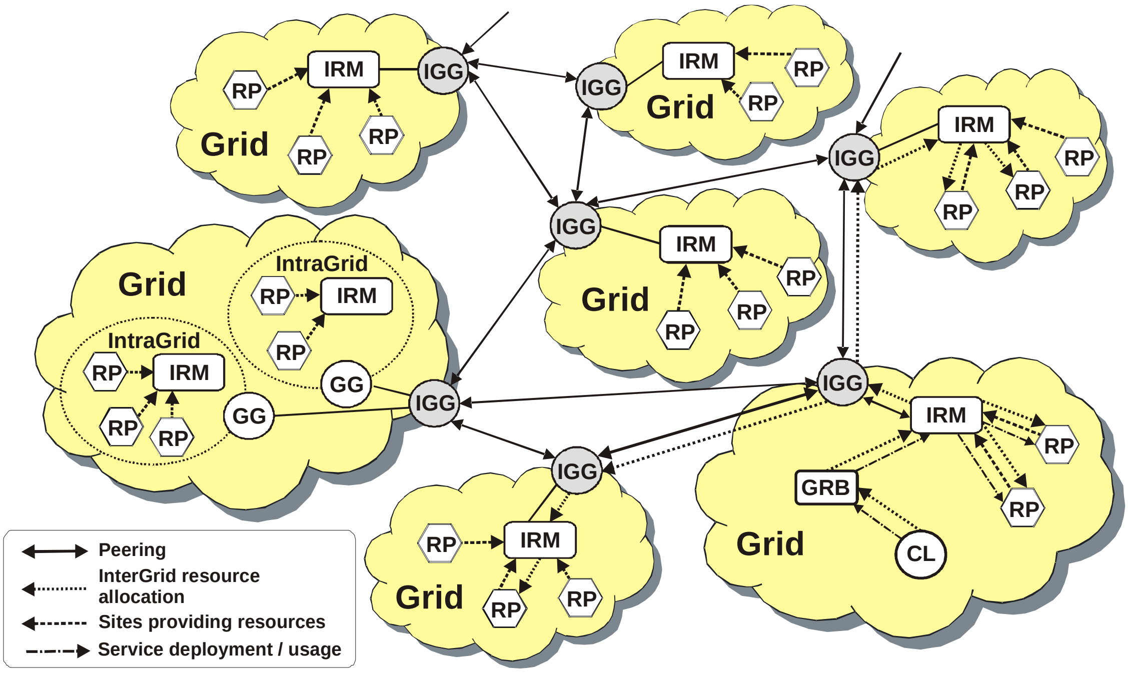

Figure 3. Abstract view of the structure of the InterGrid.

Client (**CL**) A CL or Service Creator interacts with the IRM via a Grid Resource Broker to allocate resources to deploy a service. When a CL requires more resources than what their internal RPs can provide, the IRM can contact the GG or IGG, which is then responsible for redirecting resource requests to other Grids where there is supply. An IRM can contact the IGG when the client requires more resources that its Grid is able to provide or when it is more cost-effective to allocate resources from other Grids.

InterGrid Gateway (IGG): An IGG is aware of agreements with other IGGs and thus acts as a Grid selector taking into account policies defined in the Grid to which it belongs and making the necessary conversions. These peering agreements can be established in an off line manner through SLAs. An IGG can also be in charge of allocating inter-Grid resources and establishing new peering agreements when there is a need for users from a Grid to gather resources from another Grid.

Grid Gateway (GG): A GG is responsible for similar functions as an IGG, but within a single Grid. A GG provides functions similar to those of Virtual Organization Membership Services (VOMS) [46], such as managing the membership of IntraGrids. An IntraGrid can be a Grid composing one of the islands of Grid participating in the InterGrid.

In addition, the following components are required even though they are not included in the architecture:

Standard Interfaces: The InterGrid uses common interfaces for accessing resources and the deployed services must comply with existing standards. Such common interfaces have to follow standards such as the ones proposed by the Open Grid Forum (OGF) [14].

InterGrid Directories and Marketplaces: An InterGrid Directory can be a database with information regarding Grids, Grid projects, their goals and capabilities, proposals for collaboration and requirements by Grid projects. The current facilitators for virtual organizations such as OSG and EGEE maintain InterGrid Directories with information that can be shared, such as existing VOs and their Grid projects.

As illustrated in Figure 4, we propose the internetworking of accounting systems, means to route requests for allocation of inter-Grid resources to deploy services and advertise resources among Grids. As an example we consider the following scenario. A client needs resources to deploy a service and then contacts an IRM via a GRB, which then checks whether the local Grid can provide the required resources. If the local Grid does not have the resources, the client is able to utilize inter-Grid resources available through peering agreements with other Grids. Once resources have been allocated, GRB can deploy the service or applications on the resources. IRMs maintain accounting systems to account for how much resources have been allocated and utilized. Resources are also advertised through messages sent to peering Grids informing them about availability of resources. Another way to advertise and negotiate the usage of resources is through marketplaces. For example, this scenario could be a situation in which the client as a content provider needs to allocate compute and data resources in different Grids to build a CDN. Such a request would be satisfied by forwarding the user's requirements to Grids which can offer the resources and are peered to the provider's local Grid.

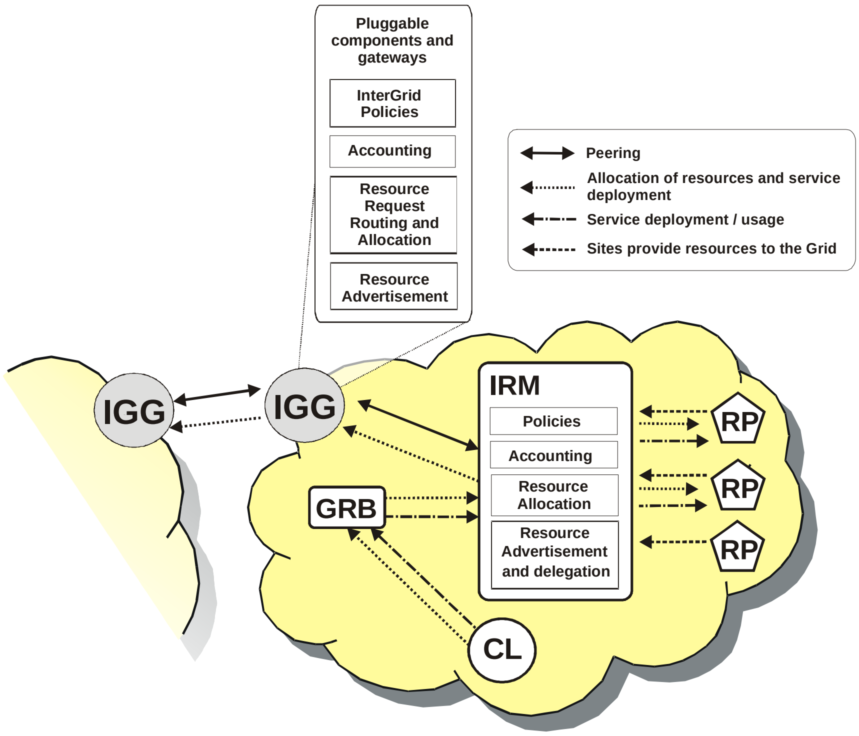

Figure 4. Internetworking of accounting systems and resource allocation.

As illustrated in Figure 4, we propose the internetworking of accounting systems, means to route requests for allocation of inter-Grid resources to deploy services and advertise resources among Grids. As an example we consider the following scenario. A client needs resources to deploy a service and then contacts an IRM via a GRB, which then checks whether the local Grid can provide the required resources. If the local Grid does not have the resources, the client is able to utilize inter-Grid resources available through peering agreements with other Grids. Once resources have been allocated, GRB can deploy the service or applications on the resources. IRMs maintain accounting systems to account for how much resources have been allocated and utilized. Resources are also advertised through messages sent to peering Grids informing them about availability of resources. Another way to advertise and negotiate the usage of resources is through marketplaces. For example, this scenario could be a situation in which the client as a content provider needs to allocate compute and data resources in different Grids to build a CDN. Such a request would be satisfied by forwarding the user's requirements to Grids which can offer the resources and are peered to the provider's local Grid.

It is also worth mention that a VO in this context is a dynamic group of resources formed by the request of clients for resources for a common purpose. A VO is thus a virtual Grid or a share of the resources provided by the InterGrid.

5. Research Agenda

As the internetworking of islands of Grids requires that fundamental issues be investigated and solved by the research community, we propose a research agenda to guide these efforts. The following agenda is defined based on issues identified and lessons learnt from our analysis of existing global ecosystems and their evolution. This agenda aims at grouping and defining some of the research topics that, in our opinion, deserve special attention in order to enable the vision of an InterGrid.

Similar to the Internet, Grid computing has been moving from a research curiosity to a support to commercial applications. In the Internet, ISPs are in the business to make profit - they see one another as competitors or sources of revenue – whereas users are interested in using services at low price. It has been shown that new design principles to the Internet have to accommodate such economic aspects [57]. Similarly, resource providers and consumers in different islands of Grids will have different interests, which can be adverse to one another. These parties will act to favor their own interests and will share resources expecting financial compensations. Economy based models are relevant to the InterGrid because resource allocation can be achieved through the economic behavior of the involved parties, markets can provide incentives for providers to offer their resources to the InterGrid, and can provide the settlements necessary between Grids. Therefore, in the research agenda we focus on economic based policies and mechanisms for interlinking of Grids.

5.1 Formation, Adaptation, Coordination Mechanisms and Self-Organization for the InterGrid

In this work, we advocate the need of Grid Gateways as entities that are aware of agreements or peering arrangements among islands of Grids. These gateways also look for partners when there is a need to collaborate or form a VO. In the InterGrid, a large number of islands of Grids or independent resource providers and consumers will coexist and interoperate. Thus, there is a need to specify how the whole system behaves and how components interact and peer in such an ecosystem. Current approaches neither define how the peering arrangements will take place nor specify how the InterGrid will self-adapt. Thus, new methods and mechanisms are necessary to enable the peering of islands of Grids through the automated and responsive formation, operation, maintenance, and dissolution of VOs or other kinds of alliances in the InterGrid. In this case, we consider that VOs are formed through the need of allocating a share of resources in the InterGrid. Hence, once the peering arrangements are defined, it is important to predict how VOs will form, self-maintain, and dissolve in this environment. In addition, mechanisms to enable the formation and evolution of self-adaptable structures for VOs are necessary. For example, data applications require peers to be located closer or links to be built dynamically among the peers involved [58].

In this context, there is a need to solve complex tasks such as the selection of peers with affinities for a given collaboration. Also important is the formation and adaptation of these shares according to application requirements and the consequent searching and selection of partners for forming VOs. Thus, self-adapting and self-organizing shares are necessary.

Considering these problems, some questions to be answered are:

- What are the mechanisms necessary to enable the formation and self-organization of overlay topologies arrangements for the InterGrid?
- What are the metaphors and techniques to be used to promote the fast convergence and evolution of the InterGrid?
- What are the mechanisms necessary to enable the formation and operation of dynamic VOs in the InterGrid?
- What would be the growth patterns of the InterGrid based on the metaphors and mechanisms chosen in the previous questions?

Coordination is necessary in the InterGrid [59]. Resource allocation approaches in Grids are currently non-coordinated and different domains have their own resource brokers, objectives, and QoS requirements. Such divergent approaches can lead us to a scenario with bad schedules and inefficient resource allocation. The problem to be addressed is how to coordinate and organize such disparate islands of Grids considering the aforementioned Grid internetworking scenario. Hierarchical coordination mechanisms can be required since islands of Grids can be formed by several other organizational Grids. Scalable coordination mechanisms for the InterGrid can allow the fast growth of such an ecosystem. On the other hand, such an environment can become complex and thus exhibit requirements for self-organization [51, 60, 61]. In this case, the global behavior emerges from local actors designed to interact locally without the sense of global control or a centralized system [61]. Therefore, there is the need for engineering and developing economic-aware self-organizing brokers and schedulers that act locally without centralized control [62].

We plan to apply computational economy as a metaphor for the internetworking of islands of Grids and formation of VOs in the InterGrid. Thus, in this scenario, some questions that need to be answered are:

- What type of market model is suitable for the InterGrid?
- What kind of coordination mechanisms needs to be implemented to allow the scalability of the InterGrid?
- What are the requirements and issues of such coordination mechanisms? If the coordination of resource and service brokers spanning different administrative domains is necessary, are there any protocols that can be used for exchanging information between brokers?
- Is it possible to achieve equilibrium through entities only engineered to achieve local self-organizing behavior?
- What are the metaphors and models that can be used for self-organization in the Grid? How can they be applied in terms of design and development?

5.2 Peering Agreements, Policies and Settlement Among InterGrid Gateways

In the Internet, ASs interconnect and carry out interdomain routing based on routing policies, thus placing varying costs on routes [63, 64]. Even though ASs can have transit or peering agreements among themselves, they can have complex policies to assess routes that lead to an AS to decide in favor of one specific route over another. There are mainly two important principles in this scenario: the incentives and viability of peering; and the minimization of costs by choosing one peering partner over another. Similar principles from policy-based routing are applicable to Grid internetworking in activities including offloading and redirection of resource allocation requests to peering Grids.

We advocate the need for Grid gateways as entities aware of agreements between islands of Grids. QoS and SLA guarantees are essential between service providers and consumers. However, the investigation of peering agreements, mechanisms and settlements among IGGs is required. The specification of policies, policy-based selection and allocation of resources from peering Grids, and game-theories of peering Grids are thus part of this research agenda. The interaction between IGGs needs to take into account the issues listed in previous topics of this research agenda, including the interoperation of accounting systems and inter-Grid resource allocation.

5.3 Pricing Resources in the Grid Economy

Buyya *et al.* [65] present a discussion on Grid Economy. The adoption of economic principles to Grids comes from observing the success of economic institutions in the real world as a sustainable model for regulating resources, goods and services. However, the adoption of such economic approaches requires the study of pricing of Grid resources and/or agreement on pricing mechanisms. Therefore, if an economic approach is used by the InterGrid, detailed studies have to be done in areas such as resource pricing, modeling consumer's utility, resource provider's marginal cost, and benefit in providing resources.

Some of the questions that need to be answered are:

- What resources should be free of charge and what resources should be priced in the Grid market?
- What are the policies that define in which circumstances resources should or should not be shared in the Grid?
- How to price the resources in the Grid?
- What kinds of issues related to the price setting for the resources arise in the Grid?
- How do resource providers adjust the price of their resources in the Grid in order to achieve the price that maximizes their profits in a competitive market, yet maintaining the equilibrium of supply and demand?
- How do the price setting mechanisms differ from one another when considering the local pricing, a competitive market and collaboration among Grids?
- How do different price mechanisms impact the system?
- How to model the resource price variation process to predict the future price of resources in the Grid?

Grid economy can become quite complex when considering the case of InterGrid. Thus, the study of pricing of resources and its effect in the Grid economy should be worked out in different steps as shown in Figure 5. For now, we are not interested in studying the impact of resources that are not priced, such as files in some P2P networks. Pricing of resources involves the following steps/issues:

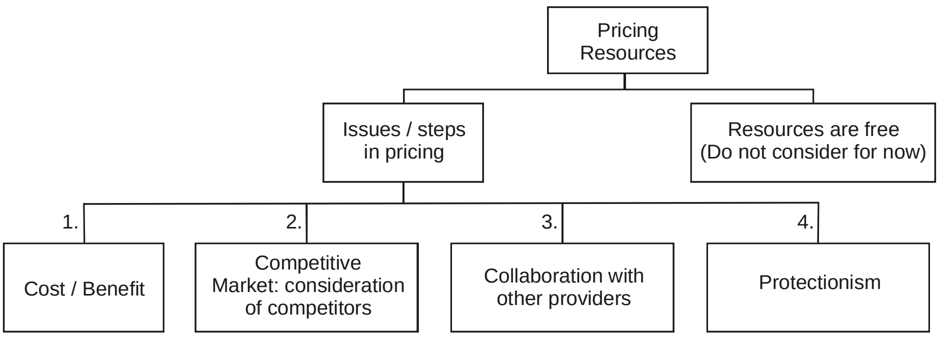

Figure 5. Pricing of resources steps/issues in the Grid.

1. Cost/Benefit - First, the pricing of resources within a local Grid is studied. The challenge here is how resource providers should calculate cost of producing the resources and benefit of providing the resources, in order to determine the prices of resources in the local Grid.

2. Competitive market: consideration of competitors - A competitive market is taken into account. Thus, resource providers must consider different ways of adjusting their prices, while maximizing their profits. The target price in the market should be the one that is the equilibrium of supply and demand and maximizes the overall welfare. In this regard, equilibrium theories are considered. Models to be developed should include resource providers of several types (e.g. providers with lower prices and better resources; lower prices and same resources; lower prices and worse resources; and vice-versa) and achieve fair allocation of resources.

3. Collaboration with other providers - Studies the effect of resource providers collaborating in reducing or increasing their prices, to eliminate competitors at an inter-Grid level. In this regard, mechanisms to avoid unexpected behavior in the economy and the emergence of monopoly or oligopoly would be studied.

4. Protectionism: The impact of local protectionism or government bodies that aim to control prices or impose barriers for using resources from other Grids would be considered. The macroeconomic aspects in price and exchange rate also need to be taken into account. These characteristics can take us to a scenario in which the Grid is a complex and unstable system because of the competition and the chaos generated by such rules, and protectionism. Thus, the main goal is to achieve equilibrium in the InterGrid for at least some time. Macroeconomics principles can be applied to achieve this goal.

5.4 Infrastructure for Grid Economics and Estimation of Requirements for Resources

Economic approaches are useful for coping with problems like providing Grid resources to different users with diverging QoS requirements and of how to reward resource suppliers. However, it is not clear whether the Grid economy should use real or virtual currency [66]. Economic models may also require globally trusted entities for several activities such as accounting, usage quota enforcement and charging. Trust federations [67] would also be required to ensure minimum levels of trust in these entities. Trying to fill the gap of global trust, the International Trust Federation aims at promoting harmonization and synchronization of regional Policy Management Authorities (PMA) policies.

Although the design of economic institutions for accounting, Grid banking and charging for resource usage is needed, it may not be possible in practice since it requires interlinking of accounting systems. In addition, if each islands of Grids adopts its own virtual currency, the detailed study of a money exchange system and its impact is important. Furthermore, electronic payment infrastructures for the InterGrid are also difficult, since countries can have different policies regarding the flow of money.

Resource exchange among VOs and Grids is also important. However, it is difficult to agree on standard units for resource usage. For example, it is difficult to evaluate how much storage is equivalent to 30 CPU hours.

Most of the current economic approaches for Grid computing assume that users know how to estimate their needs for resources, which is a fallacious assumption because users often have trouble estimating their needs [68]. Current approaches assume that the time taken for jobs to execute is known, which is generally not the case in practice. In addition, it is presumed that resource providers know how to estimate the cost for their infrastructure and that capacity planning is not an issue, which is untrue in practice [69]. Thus, these issues should be addressed correctly to effectively apply economic principles in Grids.

5.5 Integration of Accounting Systems

The integration and connection of accounting systems is necessary in CDI, where multiple CDNs interconnect in order to replicate content from their clients and need means to charge one another either based on the content replicated or the user requests satisfied by each CDN. Similarly, in Grid internetworking the peering between two Grids will require the definition of settlements and in this scenario, the integration of accounting systems is necessary. Grids therefore need to agree upon the measurements for resource usage and provide interoperation of their accounting systems. A work towards this goal is the Extensible and Economics-Inspired Open Grid Computing Platform (EGG) [70]. EGG is a macroeconomics inspired open Grid platform that promotes the use of resources across Grids through exchange rates, thus allowing each Grid to have their own local currency.

5.6 A Way Forward

The internetworking of islands of Grids and their evolution to the InterGrid level is full of challenges. It not only requires fundamental research in business models, methodologies and mechanisms that enable creation of an ecosystem for interlinking Grids, but also an open infrastructure that supports standard protocols and interoperability between Grids. As presented in the research agenda, it is important to address issues such as price setting of compute resources, inter-Grid resource allocation and internetworking of accounting systems. The research community has been addressing some issues in these respects, but a massive effort is still needed to realize the internetworking of islands of Grids and develop mechanisms that allow the InterGrid to grow in a sustainable manner. Moreover, the development of applications that can harness the capabilities of such Grids is also a challenge. The design of scalable applications that can utilize of the capabilities of networks of Grids is also of utmost importance and a challenging task.

6. Conclusions

In this work, we have presented a case and model for the InterGrid as an evolvable and sustainable Grid ecosystem. We started with the analysis of current global ecosystems, how they have evolved and what lessons can be drawn from them and applied to Grid computing in order to enable the vision of the InterGrid. An abstract architecture for the InterGrid is then presented with the aims to realize it. A discussion on current issues that arise when linking islands of Grids is also provided before a gap analysis of current technologies is given.

Existing projects aiming at creating national and continental Grids are discussed. However, applications are currently requiring amounts of resources achievable through Grids of Grids. Existing projects have tried to federate Grids and provide means to enable VOs to solve several problems. However, the cyberinfrastructure to cope with these and next generation challenges will not be realized given that today's Grids follow organizational models and mechanisms that prevent them from internetworking.

Current technologies do not allow the InterGrid vision due to conceptual and technological drawbacks such as the lack of coordination mechanisms. As argued in this work, there is the need for an architecture that allows Grids' structures to evolve from local to the InterGrid and enables the easy development of Grid applications for e-Science and e-Business. In addition, many issues related to cultural, social, and political divergences have to considered or even solved. Like the Internet, the InterGrid will comprise of numerous self-interested stakeholders and the design of its architecture has to consider these aspects [57]. Our contribution in this work is of identifying key problems in realizing a true InterGrid and delineating a research agenda on the topic. As the research agenda is a massive endeavor, we invite the global research community to address some of the issues and work collaboratively in realizing the proposed grand vision for the InterGrid.

Acknowledgements

The authors are thankful to Krishna Nadiminti, Chee Shin Yeo, Hussein Gibbins and Srikumar Venugopal from the University of Melbourne for helping to improve the language and expression in this paper. We also thank Geoffrey Fox (Indiana University), Omer F. Rana (Cardiff University), Jianhong Xia (RMIT University), Fernando Koch (Utretch University), Sushil K. Prasad (Georgia State University), Vanessa Teague (University of Melbourne), and Torsten Eymann (Universität Bayreuth) for sharing their thoughts on the topic.

References

[1] M. Baker, R. Buyya, and D. Laforenza, "Grids and Grid Technologies for Wide-Area Distributed Computing," *Software: Practice and Experience*, vol. 32, pp. 1437-1466, December 2002.

[2] L. D. Paulson, "News Briefs: Putting a Business Suit on Grid Computing," *IEEE Computer*, vol. 38, p. 24, July 2005.

[3] A. Abbas, *Grid Computing: A Practical Guide to Technology and Applications*. Hingham, Massachusetts: Charles River Media, 2004.

[4] P. A. Freeman, D. L. Crawford, S. Kim, and J. L. Muñoz, "Cyberinfrastructure for science and engineering: promises and challenges," *Proceedings of the IEEE*, vol. 93, pp. 682-691, March 2005.

[5] R. Nandkumar, "International cyberinfrastructure: activities around the globe," *CTWatch Quarterly*, vol. 2, February 2006.

[6] S. Emmott and S. Rison, "Towards 2020 Science Report," Microsoft Research, Cambridge March 2006.

[7] I. Foster, C. Kesselman, and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," *International Journal of Supercomputer Applications*, vol. 15, 2001.

[8] B. Boghosian, P. Coveney, S. Dong, L. Finn, S. Jha, G. E. Karniadakis, and N. T. Karonis, "Nektar, SPICE and Vortonics: Using Federated Grids for Large Scale Scientific Applications," in *Workshop on Challenges of Large Applications in Distributed Environments (CLADE) - in conjunction with the 15th International Symposium on High Performance Distributed Computing (HPDC-15)*, Paris, France, 2006.

[9] S. A. Kauffman, *At home in the universe: the search for laws of self-organization and complexity*. New York: Oxford University Press, 1996.

[10] M. Girvan and M. E. J. Newman, "Community structure in social and biological networks," *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, vol. 99, pp. 7821-7826, June 2002.

[11] R. Kumar, P. Raghavan, S. Rajagopalan, and D. Sivakumar, "The Web as a graph," in *19th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS '00)*, Dallas, Texas, United States, 2000, pp. 1-10.

[12] L. Smarr, "The Grid 2: Blueprint for a New Computing Infrastructure," I. Foster and C. Kesselman, Eds. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2003, pp. 3-12.

[13] "Enabling Grids for E-sciencE (EGEE) project," http://public.eu-egee.org, 2005.

[14] "Open Science Grid," http://www.opensciencegrid.org, 2005.

[15] "Teragrid," http://www.teragrid.org/, 2006.

[16] C. Catlett, P. Beckman, D. Skow, and I. Foster, "Creating and Operating National-Scale Cyberinfrastructure Services," *CTWatch Quarterly*, vol. 2, pp. 2-10, May 2006.

[17] "K*Grid," http://www.gridcenter.or.kr/.

[18] K. Miura, "Overview of Japanese science Grid project NAREGI," *Progress in Informatics*, vol. 3, pp. 67-75, 2006.
[19] N. Andrade, L. Costa, G. Germoglio, and W. Cirne, "Peer-to-peer grid computing with the OurGrid Community," in *23rd Brazilian Symposium on Computer Networks (SBRC 2005) - 4th Special Tools Session*, 2005.
[20] "Asia Pacific Grid," http://www.apgrid.org/, 2005.
[21] "INWA Project - Australia, China, United Kingdom Come Together via Grid," *Grid Today - Daily News and Information for the Global Grid Community*, vol. 4, August 2005.
[22] I. Foster and C. Kesselman, "Globus: A Metacomputing Infrastructure Toolkit," *International Journal of Supercomputer Applications*, vol. 11, pp. 115-128, Summer 1997.
[23] J. Almond and D. Snelling, "UNICORE: uniform access to supercomputing as an element of electronic commerce," *Future Generation Computer Systems*, vol. 15, pp. 539-548, October 1999.
[24] J. Dilley, B. Maggs, J. Parikh, H. Prokop, R. Sitaraman, and B. Weihl, "Globally distributed content delivery," *IEEE Internet Computing*, vol. 6, pp. 50-58, September-October 2002.
[25] J. Magretta, "Why Business Models Matter," *Harvard Business Review*, pp. 3-8, May 2002.
[26] P. Timmers, "Business Models for Electronic Markets," *Journal on Electronic Markets*, vol. 8, pp. 3-8, April 1998.
[27] J. F. Kurose and K. W. Ross, *Computer Networking: A Top-Down Approach Featuring the Internet*, 2nd ed. Boston: Addison Wesley Professional, 2002.
[28] W. B. Norton, "Internet service providers and peering," in *19th North American Network Operators Group Meeting (NANOG 19)* Albuquerque, New Mexico, 2000.
[29] T. Berners-Lee and M. Fischetti, *Weaving the Web: the past, present and future of the World Wide Web by its inventor*. London: Orion Business, 1999.
[30] "Mirror Image Internet. About the Company," http://www.mirror-image.com/company, 2006.
[31] L. Amini, A. Shaikh, and H. Schulzrinne, "Effective peering for multi-provider content delivery services," in *23rd Annual Joint Conference of the IEEE Computer and Communications Societies*, 2004, pp. 850-861.
[32] P. Rzewski, M. S. Day, and D. Gilletti, "Content Internetworking (CDI) Scenarios," RFC 3570, 2003.
[33] M. S. Day, B. Cain, G. Tomlinson, and P. Rzewski, "A Model for Content Internetworking (CDI)," RFC 3466, 2003.
[34] N. Economides, "The Economics of Networks," *International Journal of Industrial Organization*, vol. 16, pp. 673-699, October 1996.
[35] S. Androutsellis-Theotokis and D. Spinellis, "A survey of peer-to-peer content distribution technologies," *ACM Computing Surveys*, vol. 36, pp. 335-371, 2004.
[36] I. Foster and A. Iamnitchi, "On Death, Taxes, and the Convergence of Peer-to-Peer and Grid Computing," in *2nd International Workshop on Peer-to-Peer Systems (IPTPS'03)*, Berkeley, CA, 2003.
[37] O. Ardaiz, F. Freitag, L. Navarro, T. Eymann, and M. Reinicke, "CatNet - Catallactic Mechanisms for Service Control and Resource Allocation in Large Scale Application-Layer Networks," in *Workshop on Global and Peer-to-Peer Computing on Large Scale Distributed Systems, 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid*, Berlin, Germany, 2002.
[38] G. Fox and D. Walker, "e-Science Gap Analysis," UK e-Science June 2003 2003.
[39] Y.-K. Kwok, S. Song, and K. Hwang, "Selfish Grid Computing: Game-Theoretic Modeling and NAS Performance Results," in *International Symposium on Cluster Computing and the Grid (CCGrid'05)*, Cardiff, UK, 2005.
[40] R. Buyya, "Economic-based Distributed Resource Management and Scheduling for Grid Computing," Melbourne, Australia: Monash University, 2002.
[41] K. Lai, "Markets are Dead, Long Live Markets," HP Labs, Palo Alto, CA, USA February 2005.
[42] P. Grace, G. Coulson, G. Blair, L. Mathy, W. K. Yeung, W. Cai, D. Duce, and C. Cooper, "GRIDKIT: Pluggable Overlay Networks for Grid Computing," in *Distributed Objects and Applications (DOA'04)*, Cyprus, 2004.
[43] R. Ranjan, R. Buyya, and A. Harwood, "A Case for Cooperative and Incentive-Based Coupling of Distributed Clusters," in *7th IEEE International Conference on Cluster Computing*, Boston, Massachusetts, USA, 2005.
[44] M. Anastasiou, A. Bifulco, P. Findeisen, M. Hannus, I. Karvonen, H. Löh, A. Plüss, M. Ollus, and M. Weidemann, "VO Guidelines Draft v.19," Virtual Organisations Cluster, Deliverable D54.1, May 2004.
[45] C. Dumitrescu, M. Wilde, and I. Foster, "A Model for Usage Policy-Based Resource Allocation in Grids," in *Sixth IEEE International Workshop on Policies for Distributed Systems and Networks*, 2005, pp. 191-200.

[46] R. Alfieri, R. Cecchini, V. Ciaschini, L. dell'Agnello, Á. Frohner, K. Lörentey, and F. Spataro, "From gridmap-file to VOMS: managing authorization in a Grid environment," *Future Generation Computer Systems,* vol. 21, pp. 549-558, April 2005.
[47] B. Nasser, A. Benzekri, R. Laborde, F. Grasset, and F. Barrère, "Access Control Model for Grid Virtual Organizations," in *7th International Conference on Enterprise Information Systems (ICEIS 2005)*, Miami, USA, 2005, pp. 152-158.
[48] B. E. Carpenter and P. A. Janson, "Abstract Interdomain Security Assertions: A Basis for Extra-Grid Virtual Organizations," *IBM Systems Journal,* vol. 43, pp. 689-701, 2004.
[49] I. Foster, C. Kesselman, J. M. Nick, and S. Tuecke, "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration," Open Grid Service Infrastructure WG, Global Grid Forum, 2002.
[50] F. A. von Hayek, *The Fatal Conceit: The Errors of Socialism*, First ed. vol. 1. London: Routledge, 1988.
[51] D. De Roure, "On self-organization and the semantic grid," *IEEE Intelligent Systems,* vol. 18, pp. 77-79, July-August 2003.
[52] D. Irwin, J. Chase, L. Grit, A. Yumerefendi, D. Becker, and K. G. Yocum, "Sharing Networked Resources with Brokered Leases," in *USENIX Technical Conference*, Boston, Massachusetts, 2006.
[53] R. Huang, H. Casanova, and A. A. Chien, "Using Virtual Grids to Simplify Application Scheduling," in *IEEE International Parallel & Distributed Processing Symposium*, 2006, p. 10pp.
[54] P. Ruth, P. McGachey, D. Xu, and X. Jiang, "VioCluster: Virtualization for Dynamic Computational Domains," in *IEEE International Conference on Cluster Computing (Cluster 2005)* Boston, Massachusetts, USA, 2005.
[55] K. Keahey, I. Foster, T. Freeman, and X. Zhang, "Virtual Workspaces: Achieving Quality of Service and Quality of Life in the Grid," *Scientific Programming Journal,* To Appear 2006.
[56] S. Venugopal, R. Buyya, and L. Winton, "A grid service broker for scheduling distributed data-oriented applications on global grids," in *Proceedings of the 2nd Workshop on Middleware for Grid Computing*, New York, NY, USA, 2004, pp. 75-80.
[57] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," *IEEE/ACM Transactions on Networking,* vol. 13, pp. 462-475, June 2005.
[58] A. Singh and M. Haahr, "Creating an adaptive network of hubs using Schelling's model," *Communications of the ACM,* vol. 49, pp. 69-73, March 2006.
[59] T. W. Malone and K. Crowston, "The interdisciplinary study of coordination," *ACM Computing Surveys,* vol. 26, pp. 87-119, 1994.
[60] F. Heylighen and C. Gersherson, "The Meaning of Self-Organization in Computing," *IEEE Intelligent Systems,* vol. 18, pp. 72-77, July-August 2003.
[61] F. Heylighen, *The Science of Self-organization and Adaptivity* vol. Knowledge Management, Organizational Intelligence and Learning, and Complexity, 2003.
[62] P. Pattnaik, K. Ekanadham, and J. Jann, "Autonomic computing and Grid," in *Grid Computing: Making the Global Infrastructure a Reality,* F. Berman, G. Fox, and T. Hey, Eds.: John Wiley & Sons Ltd, 2003, pp. 351-361.
[63] J. Feigenbaum, R. Sami, and S. Shenker, "Mechanism design for policy routing," in *23rd Annual ACM Symposium on Principles of Distributed Computing (PODC '04)*, New York, NY, USA, 2004, pp. 11-20.
[64] J. Feigenbaum, D. R. Karger, V. S. Mirrokni, and R. Sami, "Subjective-cost policy routing," in *1st International Workshop on Internet and Network Economics (WINE 2005)*, Hong Kong, 2005, pp. 174-183.
[65] R. Buyya, D. Abramson, and S. Venugopal, "The Grid Economy," *Special Issue on Grid Computing, Proceedings of the IEEE,* vol. 93, pp. 698-714, March 2005.
[66] T. Sandholm, P. Gardfjäll, E. Elmroth, L. Johnsson, and O. Mulmo, "An OGSA-based accounting system for allocation enforcement across HPC centers," in *2nd International Conference on Service oriented computing (ICSOC '04),* New York, NY, USA, 2004, pp. 279-288.
[67] "International Grid Trust Federation," http://www.gridpma.org/, 2005.
[68] H. K. Bhargava and S. Sundaresan, "Contingent Bids in Auctions: Availability, Commitment and Pricing of Computing as Utility," in *37th Annual Hawaii International Conference on System Sciences (HICSS'04)*, Washington, DC, USA, 2004, p. 80217.3.
[69] R. Krishnan and K. Hosanagar, "Challenges in designing grid marketplaces," in *3rd International Workshop on Grid Economics and Business Models (GECON 2006)*, Singapore, 2006, pp. 61-69.
[70] J. Brunelle, P. Hurst, J. Huth, L. Kang, C. Ng, D. Parkes, M. Seltzer, J. Shank, and S. Youssef, "EGG: an extensible and economics-inspired open grid computing platform," in *3rd International Workshop on Grid Economics and Business Models (GECON 2006),* Singapore, 2006, pp. 140-150.