University Of Liège, Belgium

What is the Grid ?
Tentative Definitions
Beyond Resource Coordination

Cyril Briquet and Pierre-Arnoul de Marneffe
Department of EE & CS, University of Liège
Montefiore Institute, B37
B-4000 Liège, Belgium
Email: C.Briquet@ulg.ac.be
and PA.deMarneffe@ulg.ac.be
May 4, 2006

Grid, coordinated resource sharing, V.O. formation

Ian Foster's *What is The Grid ?* paper gives the generally accepted definition of Grid and Grid computing. While sound, it does not precisely define the involved concepts, which may lead to misunderstandings. We propose a tentative set of definitions for Grid, Grid computing and related concepts such as Virtual Organization, Grid resource sharing architectures and policies. Several design parameters are identified, the impact of the resource environment is analysed and effective combinations of design parameters are reviewed.

What is the Grid ?
Tentative Definitions Beyond Resource Coordination

Cyril Briquet and Pierre-Arnoul de Marneffe
Department of EE & CS, University of Liège
Montefiore Institute, B37
B-4000 Liège, Belgium
Email: C.Briquet@ulg.ac.be and PA.deMarneffe@ulg.ac.be

***Abstract*— Ian Foster's *What is The Grid ?* paper gives the generally accepted definition of Grid and Grid computing. While sound, it does not precisely define the involved concepts, which may lead to misunderstandings. We propose a tentative set of definitions for Grid, Grid computing and related concepts such as Virtual Organization, Grid resource sharing architectures and policies. Several design parameters are identified, the impact of the resource environment is analysed and effective combinations of design parameters are reviewed.**

I. Introduction

According to Ian Foster's well-known *What is the Grid ?* paper [1], *"a Grid is a system that*

1) *coordinates resources that are not subject to centralized control,*
2) *using standard, open, general-purpose protocols and interfaces,*
3) *to deliver nontrivial qualities of service."*

Grid computing can also be defined as *"coordinated resource sharing and problem solving in dynamic, multi-institutional collaborations"* [2]. Most authors wanting to define with accuracy what a Grid is have in general referred to these two definitions.

A standardization effort is currently getting under way. The inter-Grid project of Grid community requires some open standards. Open Standards are publicly available and implementable standards. By allowing anyone to obtain and implement the standard, they can increase compatibility between various hardware and software components. This is how Grids can be interoperable. Interoperable Grids require a common negotiation platform as any Grid that collaborates with another has to supply capabilities, security policies, and a set of other requirements to satisfy the requesting Grid. The project of an Open Grid System Architecture is *"the definition of a broadly applicable and adopted framework for distributed system integration, virtualization and management"* [3]. OGSA consists of a set of specifications concerning interfaces, behaviours, resource models and bindings. It provides an abstract definition of the set of requirements it is intended to address. This definition of requirements is based on many representative use cases that makes OGSA the best current solution for building an inter Grid standard.

However, despite the existence of two usually agreed-upon definitions and ongoing standardization efforts, we nevertheless think that a consistent set of precise definitions of Grid computing is still lacking, particularly with regards to its resource sharing aspects. Motivated by this observation, we have extracted and combined common patterns, with an emphasis on the benefits of precisely defining what is coordinated resource sharing.

In this theoretical work, we first analyse Ian Foster's checklist. For each part of the definition, we propose to add extended explanations where clarification would be helpful. The main contribution of this theoretical work is the proposal of a tentative set of definitions of Grid computing and several related concepts, with the aim to structure and consistently integrate current trends of coordinated resource sharing. Building on the proposed set of definitions, the originality of this work is that it goes beyond the usual definition of a Grid as a resource coordination architecture and proposes to distinguish centralized coordination from distributed coordination of resource sharing. Several Virtual Organization (V.O.) resource sharing policies are also presented. Finally, the question of V.O. formation is reviewed in the light of the proposed set of definitions.

In a way, these contributions may be seen as an answer to Ian Foster & al.'s recent call [4] for more integration between Grid computing and Multi Agents Systems.

The rest of the document is structured as follows: Section II analyses the well-known Three-Point Checklist, Section III synthetises a new set of definitions of Grid computing, Section IV formalizes resource sharing architectures, Section V proposes several resource sharing policies, Section VI reviews the question of V.O. formation, and finally Section VII summarizes and concludes.

II. Analysis of the Three-Point Checklist

A. Distributed Control

The first item of the Three-Point Checklist reveals that resources that are shared on a Grid are independent, which is universally agreed-upon.

The concept of Virtual Organization (V.O.), as proposed by Ian Foster & al. [5], captures the distributed nature of resource control in a Grid. V.O. members dynamically share their resources within a V.O.

B. Resource Coordination

Resource control (i.e. Who owns and has final authority over a resource ?) is different from resource management (i.e. Who decides when and why to use a resource ?).

According to the Three-Point Checklist, a Grid *coordinates* resources. If not further clarified, this may seem to restrict Grid computing to centrally coordinated resource sharing, especially since a lot of current production Grids are based on centrally coordinated resource sharing. However, a Grid can also be based on a distributed management of resource sharing.

Centralized coordination implies the existence of a common resource sharing (e.g. task scheduling) plan that all V.O. members should comply with. This type of coordination is implemented in many production Grids [6]. Distributed or individualized coordination implies that each V.O. member prepares its own resource sharing (e.g. task scheduling) plan and asks other V.O. members to comply with it, which they can accept or refuse. Individualized coordination implies the coexistence of multiple resource sharing (e.g. task scheduling) plans. Indeed, the other V.O. members may or may not exhibit a resource supplying behaviour suitable to a potential resource consumer.

It is of course in the best interest of each individually coordinated V.O. member to generate a resource consumption and supplying plan that is compatible with other V.O. members (e.g. planning resource consumption when no resource is available yields little utility). In the long term, the most profitable behaviour of each V.O. member is to synchronize its resource consumption and supplying with other V.O. members.

As these two forms of coordination are very different, we think that it is important to introduce the distinction between centrally coordinated and individually coordinated resource sharing.

C. Standard Protocols

A Grid is defined by the Three-Point Checklist to be a resource sharing system. However, one can argue that Grid computing is only a technology or platform, an infrastructure or a standard. The Three-Point Checklist does not imply that Grid computing is a standard, but it may lead one to believe that only systems using standard, open, general-purpose protocols and interfaces are actually Grids.

It is true that global interoperability brings benefits and is desirable. However, there might exist experimental or customised systems that exhibit all the characteristics of a Grid but do not implement standard protocols and thus are restricted in interoperability, despite offering true Grid resource sharing. Depending upon commercial interests, people will agree or disagree on the requirement of openness to define a Grid. But yet again, a closed source Grid based on closed protocols may nonetheless offer true Grid resource sharing.

A relevant concept is that of lightweight Grid, a term that has been used by several authors [7], [8]. *"[Being] between multicluster and grid, it can be view[ed] as a simplification of general purpose grid (as envisioned in Globus, EGEE, GLite)"* [8]. The concept of lightweight Grid refers to a lack of compliance with Grid standard functionalities, protocols or interfaces. For the sake of comparison, let's consider TCP/IP. A huge part of current networked equipment and software implement some variant of the TCP/IP stack. TCP/IP is a widespread set of efficient and scalable internetworking protocols but one cannot claim that an interconnection of networks should be labeled as such only if it implements standard, open protocols.

Based on these observations, we argue that standard and open protocols and interfaces are not a strict requirement for a system to be a Grid.

D. Nontrivial QoS

The third item of the Three-Point Checklist states that a system is a Grid if it delivers nontrivial QoS. Although this is of course a very general statement, it is also compact and encompasses several concepts that are usually agreed-upon. For example, a Grid is expected to provide higher levels of availability, reliability, autonomicity, . . . than those that could be achieved by simply adding or aggregating the performance levels of the Grid components.

However, it does not define how this nontrivial QoS is attained, meaning how access to resources of a V.O. member is granted to other V.O. members. Resources can be exchanged or volunteered (i.e. donated) between V.O. members.
Neither does the checklist define to whom is delivered the QoS that can be achieved by sharing resources within a Grid. In the case of centrally coordinated resource sharing, the V.O. members are expected to deliver their resources as planned by a centralized manager. This allows the centralized manager to gather supplied resources that can be used to

attain a certain QoS level, which may in turn be offered back to the V.O. members. There are of course several ways to redistribute the supplied resources. In the case of individually coordinated resource sharing, the V.O. members supply their resources with the expectation of being able to consume resources of other V.O. members when required to attain their self-defined QoS level. In this case, each V.O. member directly manages the level of QoS it wishes to obtain.

Given these observations, we believe that a finer-grained specification of QoS delivery would be valuable.

III. Synthesis of Grid Definitions

A. Core Grid Definitions

The basic component of a Grid is a site. Sites are the sets of resources owned by administratively independent organizations.

Definition 1. A **site** is a possibly empty or small-sized set of heterogeneous { computing, storage, networking, sensing, software } resources based on different

- hardware (PC, supercomputer, smartphone, . . .),
- platforms (operating system, CPU architecture),
- software stacks (runtime libraries and programming languages),

under the same administrative control.
Note that there is no restriction on the local management of a site's resources.

The following concept, borrowed from the Peer-to-Peer domain, is introduced to refer to the human administrators or software agents controlling a given site:

Definition 2a. A **site Peer** is the entity, composed by the set of people or software agents, that controls a site and acts on behalf of the site owner by sharing the site resources.

The goal of a site Peer is to complete owner-defined objectives under owner-defined constraints.

The resources of a Peer may be shared with other Peers. Consuming the resources of other sites is of course the main motivation for sharing resources. External resource consumption enables one to:

- solve large problems that cannot be solved with one's own resources,
- accelerate computations by temporally aggregating a great number of external resources,
- provide execution stability through redundancy by using external resources.

To consume external resources, a Peer usually supplies its own resources to other Peers. However, a Peer may own no resource (its resource set is empty) or may not wish to supply its resources (resources are restricted to local use). In these situations, resource sharing would be restricted to resource consumption and would have to be compensated by either resource supplying or external (i.e. *out of Grid*) rewarding, or both. A Peer sharing resources with other Peers may at any time be either or both a resource consumer and a resource supplier.

The grouping and connection of several Peers that share resources can be called an infrastructure. However, it should be qualified as virtual, for relationships between Peers exist only through Internet links. This infrastructure is used to produce nontrivial levels of QoS for Peers in ways that remain to be specified.

A last observation is that we did not find any objections to the introduction of recursivity into the concept of a Grid. When a Grid is centrally coordinated, the centralized coordinator may expose the same interface as that of a site Peer. It may then appear as a Peer to another Grid in which it will be integrated. The concept of Grid Peer is therefore defined as follows:

Definition 2b. A **Grid Peer** is the centralized coordinator that controls a Grid and acts on behalf of the V.O. administrator or the Peers that compose the Grid, by sharing the Grid resources.

The term *Peer* designates indistinctly and without restriction both site Peers and Grid Peers. Clients connect to Peers and submit requests to them.

Building upon these and prior observations about resource management, while also removing the requirement for compliance with standards, we propose the following definition:

Definition 3. A **Grid** is a system based on a virtual infrastructure of independent sites and Grids that adaptively share their heterogeneous resources, in a centrally (imposed) or individually (negotiated) coordinated way, through resource exchange or resource volunteering, in order to meet nontrivial multicriteria objectives.

The concept of Virtual Organization can then be formally defined as follows:

Definition 4a. A **Virtual Organization member** (V.O. member) is the union, taken as a whole, of a Peer and the site or Grid it controls.

Definition 4b. A **Virtual Organization** (V.O.) is a community of V.O. members that share their resources within a Grid.

This definition is compliant with the expectation that Grids should enable scalable V.O.

B. Grid Computing

We define the term gridification to refer both to the inclusion of a site into a Grid and to the adaptation of a software to derive benefits from a Grid.

Definition 5a. Gridification is the equipment of a site with decision making capabilities to enable the centrally or individually coordinated sharing of its resources.

Definition 5b. Gridification is the software development paradigm that enables software applications to transparently leverage Grid resources.

Definition 6. A Grid **Resource Management System** (RMS) is a middleware that allows the gridification of a site.

Building on the definition of gridification, we now give a precise meaning to Grid computing:

Definition 7. Grid computing is a form of distributed processing based on the gridification of the involved software and hardware.

C. Grid Interoperability

As the compliance with established standards has been removed from our proposed definition of a Grid, we take it into account in the following way.

Definition 8. A **gridificated** or **grid-enabled resource** is a networked resource/system equipped with a middleware that allows its gridification with other such resources/systems and its management by Resource Management Systems. Grid sites are composed of grid-enabled resources.

Definition 9. The Grid is the global interconnection of all the world Grids with standard, open and general-purpose protocols and interfaces (defined, for example, by an international body such as GGF, EGA).

D. Coordinated Resource Sharing

We now propose definitions for the two alternatives of Coordinated Resource Sharing that might be used:

Definition 10a. Centrally Coordinated resource sharing is centralized coordination of resource consumption and supplying decisions of Grid sites. Each site must comply to a common plan prepared by a central RMS (see figure 1).

Definition 10b. Individually Coordinated resource sharing is distributed coordination of resource consumption and supplying decisions of Grid sites. Each site follows its own plan established by its own RMS (see figure 2).

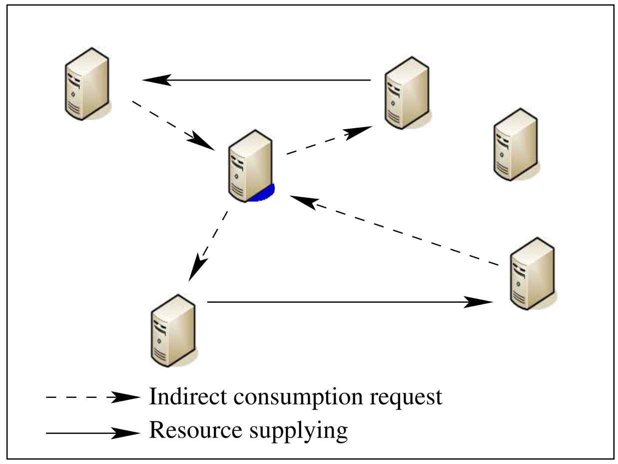

Fig. 1. Centrally Coordinated Resource Sharing

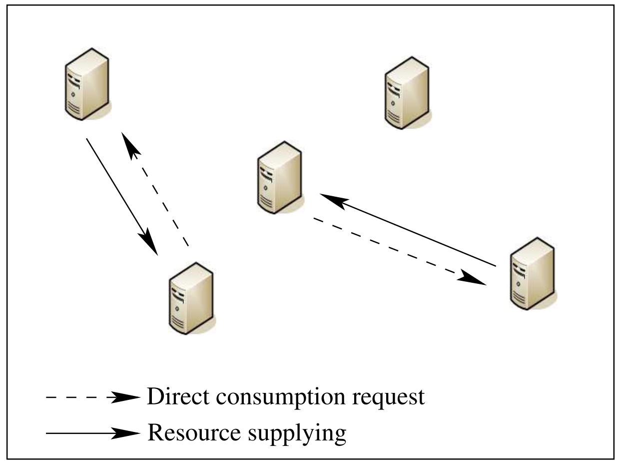

Fig. 2. Individually Coordinated Resource Sharing

To contrast Centrally and Individually Coordinated Resource Sharing, it is important to understand that sites accepting a centralized coordination of their resources accept that the decisions to consume and supply their resources are surrendered to the wisdom of a centralized manager. With either a totally stable resource environment or instantaneous and perfect information about it, centralized resource coordination could generate better utility for every site. On the other hand, individually coordinated resource sharing enables sites to adopt a more flexible and reactive resource sharing behaviour, and to take better care of their own interests.

It must be noted that centrally coordinated resource sharing does not absolutely guarantee higher Quality of Service (QoS) levels than individually coordinated resource sharing. In a stable resource environment, centrally coordinated resource sharing will attain better performance and higher QoS levels. In an unstable resource environment, individually coordinated resource sharing will offer more stable QoS because the careful planning of centrally coordinated resource sharing is unceasingly destabilized. **Therefore, choosing between**

centrally and individually coordinated resource sharing is a strategic decision that depends upon (the evaluation of) the stability of the resource environment.

It also must be noted that centrally and individually coordinated sharing may be combined to a certain extent: in an unstable environment, a centralized manager could advise sites managers on optimal resource sharing decisions.

IV. Resource Sharing Architectures

To complete the proposed definitions of the main Grid concepts, classic forms of Grid resource sharing architectures are now precisely defined.

The first architecture historically comes from the domain of cluster computing/supercomputing: *"Grid computing typically involves using many resources [...] to solve a single, large problem that could not be performed on any one resource"* [2]. Several authors have named it Virtual Supercomputing [9] with the augmented meaning that it is not dedicated to a single problem but to many applications. The Globus Toolkit [10] and gLite [6] are middlewares enabling this form of Grid computing.
⇒ This is the prime example of coordinated resource sharing: several sites supplying resources, all managed by one centralized resource manager (enabling centrally coordinated resource sharing) and multiple sites consuming resources.

Another more recent architecture is emerging as an interesting alternative. It is a form of Virtual Supercomputing crossed with Peer-To-Peer or Multi Agents System technologies. More simply, in this architecture, the latter is fully distributed, as opposed to the aforementioned Virtual Supercomputing with a centralized management of resource sharing. Each Peer negotiates resource consumption and supplying directly with other Grid Peers, usually through an RMS agent. OurGrid [11] is an excellent example of middleware enabling this architecture of resource sharing.
⇒ There are several (possibly a large number) sites supplying resources, each managed by its own resource manager (enabling individually coordinated resource sharing) and each consuming resources.

To distinguish between these two architectures, they will respectively be referred to as **Centrally Coordinated Virtual Supercomputing (CCVS)** and **Individually Coordinated Virtual Supercomputing (ICVS)**.

Another trend that has attracted considerable attention has been named **Internet Computing** [9] or, more figuratively, **Volunteer Computing** [12]. Famous examples of this form of Grid computing include SETI@home [13], Folding@home [14]. Cycle Stealing projects like Condor [15] could also be included.
⇒ There are several (a huge number) sites supplying resources, one centralized resource manager (enabling centrally coordinated resource sharing) and usually only one site consuming resources (the site that hosts the Grid manager). However, it should not be ruled out that such initiatives as the Compute feature of the Google Toolbar [16] pool benevolent home users' resources and offer them to more than one scientific project. Therefore, it must be considered that Internet Computing may also have several resource consumers.

Another recent trend currently generating much interest has been named **Desktop Grid** [17], which is very similar to Internet Computing. There are differences in the existing definitions of Desktop Grid, but also some common patterns:

1) the sites all belong to the same real-world organization rather than to individuals scattered on the Internet;
2) idem, but every site may consume resources;
3) there is only one site of which the resource set is composed of several desktop PC.

As it will now be shown, the Volunteer Computing and Desktop Grid architectures can be modeled as variants of the main architectures (CCVS, ICVS) combined with specific resource sharing policies.

V. V.O. Resource Sharing Policies

A fundamental requirement of Grid computing [3] is resource sharing across independent organizations. The Grid virtual infrastructure is composed of Peers representing independent organizations that share various levels of trust, agreement and affinity. In this section, decision making in resource sharing is explored and several policies (Philanthropy, Mutualism and Individualism) are proposed, each corresponding to a different expectation of reciprocity. A Peer that supplies resources to other Peers can expect to get back (i.e. to be able to consume) the same amount of resources, a small but reasonable amount or none. Let's follow the resource trail ...

The first proposed resource sharing policy is Philanthropy.

Definition 11a. A **philanthropic resource sharing policy** means that a centralized coordinator maximizes the utility of one Peer (usually itself) without giving any utility to the other Peers. No accounting of resource exchange (which is unilateral) is kept.

The main benefit of Philanthropy is the aggregation of huge amounts of resources that allows one Peer to run large scale applications.

Philanthropic resource sharing is the policy typically used in Volunteer Computing, where most Peers are supplying resources and one Peer consumes them. It must be noted that this policy is not incompatible with the ICVS (individually coordinated) architecture, as free resource supplying may also happen within a pair of Peers only.

The second proposed resource sharing policy is Mutualism.

Generally speaking, a mutualistic organization is created to provide its members with the best possible service and maximum return on investment, without keeping any benefit for itself. This kind of business may even be owned by its members. Moreover, it can be conceptually thought of as a health care insurance.

Definition 11b. A **mutualistic resource sharing policy** means that a centralized coordinator globally maximizes Peers utility and resource utilization, without keeping long-term resource exchange accounting.

Members of mutualistic organizations are not expecting to be rewarded for all their resource supplying, but expect instead to get some proportional rewarding (i.e. schedule priority). It really is load balancing.

The main benefit of mutualism is that if one Peer suffers some trouble (e.g. resource failure, transient request overload) and cannot supply enough resources for some time, it will still be able to consume resources, but less than it could in usual operating conditions. When a few Peers are facing an instability of their resource and/or request environment, mutualism enables load balancing between all Peers. However, all the other Peers will be penalized but as the burden will be equally shared, they will be able to consume resources in only a slightly smaller amount than they could in usual operating conditions. It can be seen as a form of fault tolerance where performance penalties are shared among components in the system.

Another important aspect is that the Peer that receives help from the other Peers is not penalized because there is no long-term accounting of resource exchange. Such a policy is therefore highly suggested when

- there is strong trust between Peers (e.g. Peers belong to the same enterprise or association) and
- the total amount of consumed resources within the V.O. is smaller than the total amount of supplied resources jointly consumed by the Peers.

Indeed, if there is no trust, there is a high risk of free-riding [18]. And if there are not enough idle resources, the form of redundancy proposed by a mutualistic policy will not be possible.

A mutualistic policy is typically used in Desktop Grids. The architecture is that of an enterprise-level Grid where cycle stealing is performed on idle desktop PC. There is V.O. with several Peers representing the various departments/units of the firm. These Peers share their resources with a mutualistic policy. It must be noted that this policy is incompatible with ICVS (individually coordinated) architecture because of the requirement of a centralized coordinator able to balance load between Peers.

The third proposed resource sharing policy is Individualism.

Definition 11c. An **individualistic resource sharing policy** means that a centralized coordinator or distributed coordinators maximize Peers utility and maintain long-term resource exchange accounting.

With an individualistic policy, an accounting of resource exchange is maintained independently by each Peer. Peers can then consume as many resources as they supply and do not have to supply more than they consume. The goal of the individualistic policy is to separate concerns of the Peers and maximize their utility independently.

The main benefit of an individualistic policy is avoidance of free-riding, which depends upon the accuracy of resource exchange accounting [19]. This policy will incite Peers to supply resources as they know undue overconsumption of their resources by other Peers will be limited.

It must be noted that this policy is compatible with both CCVS (centrally coordinated) and ICVS (individually coordinated) architectures. It is of course the policy of choice used within the latter.

Within an ICVS architecture, use of this policy requires a careful selection of a bootstrapping strategy [20] as Peers may deny resource supplying requests. Indeed, if all Peers wait to have consumed some resources before supplying their own, not much exchange will take place. Every Peer will remain idle, ridden by fear of being free-ridden. An option to overcome this initial lack of trust is that Peers randomly accept a small, yet nonzero, percentage of supplying requests from Peers that do not have a good resource exchange history (i.e. they did not themselves supply resources to the considered Peer). Another option is that Peers accept all supplying requests as long as there are no pending requests to supply Peers with higher priority [11] (i.e. with better resource supply history). With this second option, resource utilization is promoted as idle resources will be supplied so as to build trust with other Peers. The counterpart of both reviewed bootstrapping options is a risk of free-riding, but it will be limited either to a small percentage of resource utilization or to periods when resource utilization is low anyways.

An interesting observation about these three resource sharing policies is that philanthropic, mutualistic and individualistic policies may be viewed as discrete points in a continuum of policies ranging from long-term to opportunistic relationships, with accounting of resource supplying ranging from nonexistent to loose to accurate.

Finally, there are a couple of aspects of resource sharing that are orthogonal, yet relevant, to all three reviewed resource sharing policies: the import and export of resources *out of Grid*, meaning the consumption and supplying of resources enabled by rewards that are external to a Grid. *Out of Grid* rewards include real money, feel-good, or an external agreement between Peers administrators.

Definition 12a. Import of resources may take place when some Peers do not own any resource or have exhausted their consumption potential, cannot supply any more resources and still want to consume resources. A Peer may still offer an *out of Grid* reward to augment its consumption potential.

For example, so-called Utility Computing offerings [21] offer Grid resource supplying for sale. In another context, scientific projects of general interest [13], [14] allow home users to supply their resources against a feeling of taking part in a project useful to mankind. In yet another context, human administrators of a Grid may decide to lend access to their resources to the administrators of another Grid, and therefore transiently share some resources with another Grid. Whatever its nature, *out of Grid* rewarding can be used with all three resource sharing policies previously reviewed.

Definition 12b. Export of resources may take place when, within a centrally coordinated architecture, the centralized manager also follows objectives of its own and consumes some resources of V.O., possibly to supply them to an external entity.

This arrangement can be modeled as the application of both a philanthropic and another (mutualistic or individualistic) policy. After the philanthropic policy has been applied, the other policy is applied to the resources that were not exported.

VI. V.O. Formation

V.O. formation remains an open question [4] and has not been explored thoroughly until very recently. In this section, the lifecycle of a V.O. is explored and links are established with other Grid aspects that were discussed.

A first observation is that a Peer needs to be motivated to enter a V.O. As already stated, solving large problems, accelerating computations and augmenting execution stability are the main motivations to gridificate one's resources and share them with other Peers. Indeed, through resource sharing, Grid computing enables transient collaborations as well as lasting partnerships, which promote, respectively, dynamic resource consumption opportunities [22] and stabilization of the resource environment.

A second observation is that there are multiple stakeholders who may have an interest in creating a V.O. or making a Peer enter a V.O. Several classes of stakeholders have been identified [18]:

- "*end users making use of Grid applications [...],*"
- "*resource administrators and owners,*"
- "*and V.O. administrators and policy makers.*"

There are basically two ways to form a V.O. [23]: top-down and bottom-up. Top-down creation of a V.O. may be initiated by stakeholders who want to control a Grid. Top-down creation is necessarily initiated *out of Grid*. A centrally coordinated architecture would usually be selected because it would be in the interest of rational human stakeholders to create a V.O. from scratch with mostly stable resources. Most current production Grids are created top-down. Bottom-up creation of a V.O. may be initiated by stakeholders who want to be part of a Grid. Bottom-up creation is necessarily initiated by Peers. An individually coordinated architecture would usually be selected because this bottom-up creation is the scenario of choice for Peers among which there is no or little trust and where human administrators are not related, thus implicating a possibly unstable resource environment. Though a V.O. created top-down could be viewed as being owned by its V.O. administrators, a V.O. created in a bottom-up manner certainly does not belong to anyone, just as today's Internet does not either.

We now propose a formal definition of the structure of a V.O. that includes the openness and recursivity features that were discussed. A directed acyclic graph (DAG) structure allows the grouping of several V.O. members or V.O. to form a new V.O., enabling seamless, multi-level V.O. formation.

Definition 13. The **structure of a Virtual Organization** is a directed acyclic graph $G = (V, E)$, where V is a set of vertices and where $E = V \times V$ is the transition relation between the vertices (see figure 3). Let $V_L \subset V$ be the set of the leaves of the graph, and $V_I \subset V$ be the set of its internal nodes. Each leaf vertex $v_L \in V_L$ corresponds to a V.O. member, whereas each internal node $v_I \in V_I$ describes the structure of a V.O. An edge $(v, v') \in E$ models the fact that the V.O. v' is a member of the V.O. v (either V.O. or V.O. members). As a shorthand, let $C(v) = \{v' \mid (v, v') \in E\}$ be the set of members of a V.O. v. As a V.O. is composed of at least two members, it is required that each internal node has at least two successors. Precisely, for each internal node $v_I \in V_I$, the equation $|C(v_I)| \geq 2$ must hold.

A third observation about V.O. policies is that in V.O. that are created bottom-up, the expected initial lack of trust calls for the use of an individualistic resource sharing policy. However, after stability in resource exchange has been achieved and maintained for some time, leading to a building of trust, Peers in such a V.O. could consider switching to a mutualistic resource sharing policy. This would guarantee that a Peer experiencing transient abnormal conditions, precluding its expected resource supplying, would be helped by other

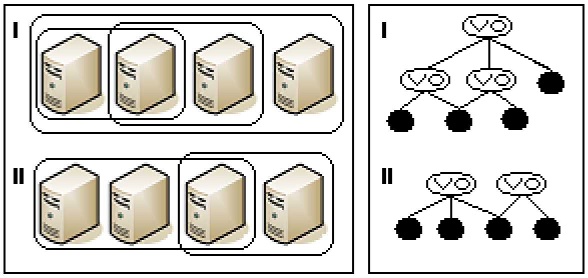

Fig. 3. Structure of Virtual Organizations (1, then 2 top-level V.O.)

Peers. From this perspective, switching to mutualism can be seen as a rational decision enabling the preservation of stable resource exchange patterns. As some authors argue that stability is a form of robustness [24], there are clearly incentives for such a policy switch.

Another way to promote stability is for Peers belonging to a given V.O. to reserve a small part of their resource supplying for external Peers. In other words, it can be interesting for Peers to belong to several V.O. at the same time, in order to foster new possibilities of resource exchange. By continually maintaining several possibilities of resource consumption, even at low levels, Peers would be more resilient to the collapse of their primary V.O. and would recover faster in such worst-case scenarios. In regular scenarios, Peers would also be able to augment the utilization of their own resources.

An open question is the openness of a V.O.: under what conditions is it profitable to allow a new Peer into a given V.O. ? Given that temporally aggregating external resources is a main motivation for a Peer to share resources, there is the consequence that Peers should associate with Peers that have complementary expertise. It means that resource supplying patterns in a V.O. should probably be temporally heterogeneous. Another open question is the tolerance of a V.O.: under what conditions is it profitable for a V.O. and most of its Peers to keep a Peer which regularly exhibits failures of resource supplying ? Yet another open question is the dissolution of a V.O.: under what conditions is it profitable for Peers to remain in a V.O. when most Peers regularly exhibit failures of resource supplying ?

VII. Conclusions

In this theoretical work, Ian Foster's checklist has been analysed and the parts that may benefit from an extended definition have been discussed. A tentative set of definitions were then proposed for Grid and Grid computing, as well as for related concepts such as V.O., Grid resource sharing architectures and policies.

The impact of initial V.O. formation and resource environment on the consistent selection of a Grid resource sharing architecture and a resource sharing policy has been thoroughly explored, and effective combinations of Grid design parameters have been reviewed. It is our hope that the tentative set of definitions that have been presented will prove valuable for researchers in Grid computing.

Acknowledgement

We want to thank Christophe Noël and Krzysztof Rzadca for their most useful suggestions, Valérie Leroy for the production of the drawings, and Sébastien Jodogne for his insights. We also want to thank Claire Kopacz and Elisabeth Spilman for the linguistic review.

References

[1] I. Foster, "What is the grid ? a three point checklist," *Grid Today*, July 2002.

[2] J. Nabrzyski, J. Schopf, and J. Weglarz, Eds., *Grid Resource Management: State of the Art and Future Trends*. Kluwer Academic Publishers, 2003.

[3] "The open grid services architecture, version 1.0." [Online]. Available: http://www.gridforum.org/documents/GWD-I-E/GFD-I.030.pdf

[4] I. Foster, N. Jennings, and C. Kesselman, "Brain meets brawn: Why grids and agents need each other," in *Proc. Int. Conference on Autonomous Agents and Multi-Agent Systems*, New York, 2004.

[5] I. Foster, C. Kesselman, and S. Tuecke, "The anatomy of the grid: Enabling scalable v.o." *Int. J. Supercomputer App.*, vol. 15(3), 2001.

[6] "glite project." [Online]. Available: http://glite.web.cern.ch/

[7] "Stripped-down grid: A lightweight grid for computing's have-nots." [Online]. Available: http://www.hpl.hp.com/news/2005/jan-mar/grid.html

[8] O. Richard, "Grid 5000 report." [Online]. Available: https://www.grid5000.fr/mediawiki/index.php/Special:UserReports?menu=renderreport&username=Orichard

[9] F. Capello and al., "Grid'5000: A large scale, reconfigurable, controlable and monitorable grid platform," in *Proc. Grid'2005 workshop*, Seattle (WA, USA), 2005.

[10] "The globus alliance." [Online]. Available: http://www.globus.org/

[11] N. Andrade, W. Cirne, F. Brasileiro, and P. Roisenberg, "Ourgrid: An approach to easily assemble grids with equitable resource sharing," in *Proc. Workshop on Job Scheduling Strategies for Parallel Processing*, Seattle (WA, USA), 2003.

[12] D. Anderson and G. Fedak, "The computational and storage potential of volunteer computing," in *Proc. CCGRID 2006*, Singapore, 2006.

[13] "Seti@home." [Online]. Available: http://setiathome.berkeley.edu/

[14] "Folding@home." [Online]. Available: http://folding.stanford.edu/

[15] R. Raman, M. Livny, and M. Solomon, "Matchmaking: Distributed resource management for high throughput computing," in *Proc. HPDC7*, Chicago (IL, USA), 1998.

[16] "Google toolbar." [Online]. Available: http://toolbar.google.com/

[17] R. Olejnik, B. Toursel, M. Tudruj, and E. Laskowski, "Dg-adaj: a java computing platform for desktop grid," in *Proc. Cracow Grid Workshop*, Cracow (Poland), 2005.

[18] I. Foster and C. Kesselman, Eds., *The Grid 2: Blueprint for a New Computing Infrastructure*. San Francisco (CA, USA): Morgan Kaufmann, 2004.

[19] R. Santos, A. Andrade, F. Brasileiro, W. Cirne, and N. Andrade, "Accurate autonomous accounting in peer-to-peer grids," in *Proc. Workshop on Middleware for Grid Computing*, Grenoble (France), 2005.

[20] B. Chun, Y. Fu, and A. Vahdat, "Bootstrapping a distributed computational economy with peer-to-peer bartering," in *Proc. Workshop on Economics of Peer-to-Peer Systems*, Berkeley (CA, USA), June 2003.

[21] "Sun grid compute utility." [Online]. Available: http://www.network.com/

[22] J. Patel and al., "Monitoring, policing and trust for grid-based virtual organisations," in *Proc. UK e-Science All Hands Meeting*, Nottingham (UK), 2005.

[23] K. Amin, G. von Laszewski, and A. Mikler, "Grid computing for the masses: An overview," in *Proc. Workshop on Grid and Cooperative Computing*, Shangai (China), 2003.

[24] R. Leus, "The generation of stable project plans," Ph.D. dissertation, Catholic University of Leuven, Leuven (Belgium), 2003.