

On Death, Taxes, and the Convergence of Peer-to-Peer and Grid Computing

Ian Foster^{1,2} Adriana Iamnitchi¹

¹ Department of Computer Science, University of Chicago, Chicago, IL 60615

^M
²M

integration of substantial resources to deliver nontrivial qualities of service within an environment of at least limited trust. For example, NASA's Information Power Grid links supercomputers at four NASA laboratories [26]. In contrast, current P2P systems deal with many more participants (e.g., Limewire [31] reports hundreds of thousands in Gnutella) but offer limited and specialized services, have been less concerned with qualities of service, and have made few if any assumptions about trust.

These characterizations and examples might not suggest commonality of interest. Nevertheless, we shall argue that in fact the two environments are concerned with the same general problem, namely, resource sharing within VOs that may not overlap with any existing organization. Clearly, the two types of system have both conceptual and concrete distinctions, which we shall identify and illuminate, focusing as noted above on characteristics of deployed systems. We shall show that the distinctions seem to be the result of different target communities and thus different evolutionary paths. Grids have incrementally scaled the deployment of relatively sophisticated services and application, connecting small numbers of sites into collaborations engaged in complex scientific applications. As system scale increases, Grid developers are now facing and addressing problems relating to autonomic configuration and management. P2P communities developed rapidly around unsophisticated, but popular, services such as file sharing and are now seeking to expand to more sophisticated applications as well as continuing to innovate in the area of large-scale autonomic system management. We expect the definition of persistent and multipurpose infrastructure to emerge as an important theme.

3.1 Target Communities and Incentives

The development and deployments of **Grid** technologies were motivated initially by the requirements of professional communities needing to access remote resources, federate datasets, and/or pool computers for large-scale simulations and data analyses. Although Grid technologies were initially developed to address the needs of scientific collaborations, commercial interest is growing. Participants in contemporary Grids thus form part of established communities that are prepared to devote effort to the creation and operation of required infrastructure and within which exist some degree of trust, accountability, and opportunities for sanctions in response to inappropriate behavior. At the same time,

the dynamic nature of VO existence and membership and the often-limited engagement of VO participants circumscribe the ability to impose solutions at individual sites (where local VO participants may have only limited authority) and VO-specific administration.

In contrast, **P2P** has been popularized by grassroots, mass-culture (music) file-sharing and highly parallel computing applications [4, 5] that scale in some instances to hundreds of thousands of nodes. The "communities" that underlie these applications comprise diverse and anonymous individuals with little incentive to act cooperatively. Thus, for example, we find that in file sharing applications, there are few providers and many consumers [2]; the operators of SETI@home [4] devote significant effort to detecting deliberately submitted incorrect results; and people tend to intentionally misreport their resources [46]. Thus, rule enforcing mechanisms as well as incentives for good behavior must be provided by the system (more in Section 3.5).

The need for other participation models based, for example, on payment, contracts, markets, and licensing is recognized in both systems [28, 30, 54] but is not yet standard practice.

3.2 Resources

In general, **Grid** systems integrate resources that are more powerful, more diverse, and better connected than the typical **P2P** resource. A Grid resource might be a cluster, storage system, database, or scientific instrument of considerable value that is administered in an organized fashion according to some well-defined policy. This explicit administration enhances the resource's ability to deliver desired qualities of service and can facilitate, for example, software upgrades, but it can also increase the cost of integrating the resource into a Grid. Diversity in architecture and policy makes the publication of resource properties important [16]. Explicit administration, higher cost of membership, and the stronger community links within scientific VOs mean that resource availability tends to be higher and more uniform. In contrast, P2P systems often deal with intermittent participation and highly variable behavior: for example, in the case of Mojo Nation it is reported [54] that average connection time was only 28% and highly skewed (one sixth of the nodes always connected).

Grids integrate not only "high-end" resources: desktop systems with variable availability [34] form a major component of many contemporary Grids [12].

However, the ensemble of all such resources within a Grid are not treated as an undifferentiated swarm of global scope. Rather, they are aggregated within administrative domains via technologies such as Condor [32, 33] to create local resource pools that are integrated into larger Grids via the same Grid protocols as other computational resources.

Resources in Grids, traditionally from c

SETI@home) to second-generation flooding-based (e.g., Gnutella file retrieval) and then third-generation systems based on distributed hash tables. First- and second-generation P2P collaborations have been characterized at the level of both individual nodes (behavior, resources [2, 46]) and network properties (topological properties [42], scale [31], traffic [47]), revealing not only general resilience but also unexpected emergent properties. Third-generation systems have been characterized primarily via simulation studies [41, 50] rather than large-scale deployments. Scalable autonomic management clearly has been achieved to a significant extent in P2P, albeit within specific narrow domains.

3.5 Services and Infrastructure

The technologies used to develop Grid and P2P applications differ both in the specific services provided and in the emphasis placed on persistent, multipurpose infrastructure.

Much work has been expended within the **Grid** community on both technical and organizational issues associated with creating and operating persistent, multipurpose infrastructure services for authentication [23], authorization [37, 52], discovery [16], resource access, data movement, and so forth. (Perhaps because of the relatively self-contained nature of early Grid communities, less effort has been devoted to managing participation in the absence of trust, via accounting, reputation management, and so forth, although these issues are increasingly discussed.) We use the term *persistent* to indicate that services are operated by participants over extended periods as critical and often highly available infrastructure elements, like DNS servers; and *multipurpose* to indicate that the same services are used for multiple purposes (e.g., the same monitoring and discovery service [16] is used by a wide range of higher-level functions, such as computation scheduling, data replication, and fault detection). These services operate as overlays on resources and services maintained by participating institutions. Gatewaying from these overlay structures to local mechanisms and policies is a significant concern.

Many Grid communities use the open source Globus Toolkit [20] as a technology base. Significant effort has been channeled toward the standardization of protocols and interfaces to enable interoperability between different Grid deployments. The Open Grid Services Architecture (OGSA) [22] is such an effort: it integrates Grid and Web services technologies to define a service-oriented architecture within which all

services adhere to a set of standard interfaces and behaviors (some required, others optional) for such purposes as service creation, registry, discovery, lifecycle, service data query, notification, and reliable invocation.

P2P systems have tended to focus on the integration of simple resources (individual computers) via protocols designed to provide specific vertically integrated functionality. Thus, for example, Gnutella defines its own protocols for search and network maintenance. Such protocols do, of course, define an infrastructure, but in general (at least for second- and third-generation systems) the persistence properties of such infrastructures are not specifically engineered but are rather emergent properties. Over time, experience with these emergent properties has revealed the need for new services, such as anonymity and censorship resistance [48], incentives for fair sharing and reputation management [14], and result checking [45]—important issues that have not tended to arise to date in Grid computing, because of different underlying trust assumptions.

JXTA [27], XtremWeb [18], and BOINC [9] have been proposed as standard service infrastructure for P2P systems but, to date, have seen little adoption and no interoperability. Thus, for example, a user participating in Gnutella, KaZaA, SETI@home, and FightAIDSatHome must run four independent applications, each coded from scratch and running its own protocols over its own overlay networks. This lack of interest in standard infrastructure may perhaps derive from the simple nature of current P2P applications and the fact that a typical user does not encounter a substantial administrative burden even when running multiple applications. As functionality requirements increase, so presumably will the benefits of standard tooling and infrastructure.

While P2P and Grid service requirements overlap in many regards, there are also important distinctions. First, some services are specific to particular regimes: for example, mechanisms that make up for the inherent lack of incentives for cooperation in P2P. Second, functionality requirements can conflict; for example, Grids might require accountability and P2P systems anonymity. Third, common services may start from different hypotheses, as in the case of trust.

4 Future Directions

Grid and P2P are both concerned with the pooling and coordinated use of resources within distributed communities and are constructed as overlay structures that operate largely independently

of institutional relationships. Yet despite these commonalities, there are also major differences in communities, incentives, applications, technologies, resources, and achieved scale.

Nevertheless, we argue that the vision that motivates both Grid and P2P computing—that of a worldwide computer within which access to resources and services can be negotiated as and when needed—will come to pass only if we are successful in developing a technology that combines elements of what we today call both P2P and Grid computing. This technology will address *failure* (death) at a fundamental level, using scalable self-configuring protocols such as those emerging from P2P research. It will also provide persistent and multipurpose *infrastructure* (at some cost, justified because amortized over many uses and users), which like DNS and routing tables will be supported in an organized and distributed fashion and will exploit heterogeneity (whether naturally occurring or artificially imposed) in its environment to achieve goals of robustness, performance, and trust. Diverse discovery, negotiation, and maintenance protocols constructed on some common base will be used to deliver a wide spectrum of services and qualities of service.

The Grid and P2P communities are approaching this nirvana from different directions. Over time, the scale of Grid systems is increasing as barriers to participation are lowered and as commercial deployments enable communities based on purely monetary transactions. For example, the International Virtual Data Grid Laboratory [8] is deploying to scores of sites and many thousands of resources both nationally and internationally, and interest is growing in utility computing models [28] that establish sharing relationships based on commercial transactions rather than common interest. Both trends lead to a greater need for scalability, trust negotiation, self-configuration, automatic problem determination, and fault tolerance [11, 16]—areas where P2P has much to offer. OGSA definition work, proceeding within the Global Grid Forum, and an OGSA-based Globus Toolkit 3 are stimulating much work on service definition and implementation.

Meanwhile, developers of P2P systems are becoming increasingly ambitious in their applications and services, as a result of both natural evolution and more powerful and connected resources. We expect that the developers of such systems are going to become increasingly interested in standard infrastructure and tools for service description, discovery, and access, as well in standardized service

definitions and implementations able to support different mixes of logical and physical organizations.

This analysis suggests to us that the Grid and P2P communities have more in common than is perhaps generally recognized and that a broader recognition of key commonalities will tend to accelerate progress in both disciplines—which is why we wrote this article.

Acknowledgments

We are grateful to H. Casanova, K. Ranganathan, and M. Ripeanu for comments. Our work was supported in part by the NSF GriPhyN project.

References

1. Berkeley Open Infrastructure for Network Computing. <http://boinc.berkeley.edu>, 2002.
2. DZero Experiment. www-d0.fnal.gov.
3. JXTA. www.jxta.org.
4. Limewire. www.limewire.com.
5. Abramson, D., Sosic, R., Giddy, J. and Hall, B., Nimrod: A Tool for Performing Parameterised Simulations Using Distributed Workstations. *Proc. 4th IEEE Symp. on High Performance Distributed Computing*, 1995.
6. Adar, E. and Huberman, B.A. Free Riding on Gnutella. *First Monday*, 5 (10). 2000.
7. Allen, G., Dramlitsch, T., Foster, I., Goodale, T., Karonis, N., Ripeanu, M., Seidel, E. and Toonen, B., Supporting Efficient Execution in Heterogeneous Distributed Computing Environments with Cactus and Globus. *SC'2001*, 2001, ACM Press.
8. Anderson, D.P., Cobb, J., Korpella, E., Lebofsky, M. and Werthimer, D. SETI@home: An Experiment in Public-Resource Computing. *Communications of the ACM*, 45 (11). 56-61. 2002.
9. Anderson, D.P. and Kubiawicz, J. The Worldwide Computer. *Scientific American* (3). 2002.
10. Annis, J., Zhao, Y., Voelckler, J., Wilde, M., Kent, S. and Foster, I., Applying Chimera Virtual Data Concepts to Cluster Finding in the Sloan Sky Survey. *SC'2002*, 2002.
11. Anstreicher, K., Brixius, N., Goux, J.-P. and Linderoth, J.T. Solving Large Quadratic Assignment Problems on Computational Grids. *Mathematical Programming*, 91 (3). 563-588. 2002.
12. Avery, P., Foster, I., Gardner, R., Newman, H. and Szalay, A. An International Virtual-Data Grid Laboratory for Data Intensive Science, 2001. www.griphyn.org.
13. Casanova, H., Obertelli, G., Berman, F. and Wolski, R., The AppLeS Parameter Sweep Template: User-Level Middleware for the Grid. *Proc. SC'2000*, 2000.
14. Chervenak, A., Deelman, E., Foster, I., Guy, L., Hoschek, W., Iamnitchi, A., Kesselman, C., Kunszt, P., Ripeanu, M., Schwartzkopf, B., Stockinger, H., Stockinger, R. and Tierney, B., Giggle: A Framework for Constructing Scalable Replica Location Services. *SC'02*, 2002.
15. Chien, A., Calder, B., Elbert, S. and Bhatia, K. Entropia: Architecture and Performance of an Enterprise Desktop Grid System. *Journal of Parallel and Distributed Computing*, To appear.
16. Clarke, I., Sandberg, O., Wiley, B. and Hong, T.W., Freenet: A Distributed Anonymous Information Storage and Retrieval System. *International Workshop on Designing Privacy Enhancing Technologies*, Berkeley, CA, USA, 2000, Springer-Verlag.
17. Cornelli, F., Damiani, E., Capitani, S.D., Paraboschi, S. and Samarati, P., Choosing reputable servants in a P2P network. *International World Wide Web Conference*, Honolulu, Hawaii, USA, 2002, ACM Press, 376 - 386.
18. Crowcroft, J. and Pratt, I., Peer to Peer: peering into the future. *IFIP-TC6 Networks 2002 Conference*, Pisa, Italy, 2002, Springer Verlag.

19. Czajkowski, K., Fitzgerald, S., Foster, I. and Kesselman, C., Grid Information Services for Distributed Resource Sharing. *10th IEEE International Symposium on High Performance Distributed Computing*, 2001, IEEE Press, 181-184.
20. Fedak, G., Germain, C., Néri, V. and Cappello, F., XtremWeb : A Generic Global Computing System. *Workshop on Global Computing on Personal Devices (CCGRID2001)*, Berlin, Germany, 2001, IEEE Press.
21. Foster, I. The Grid: A New Infrastructure for 21st Century Science. *Physics Today*, 55 (2nd y