Managing Computer Networks Using Peer-to-Peer Technologies

Lisandro Zambenedetti Granville, Diego Moreira da Rosa, André Panisson, Cristina Melchiors,

Maria Janilce Bosquiroli Almeida, and Liane Margarida Rockenbach Tarouco,

Federal University of Rio Grande do Sul

ABSTRACT

Peer-to-peer systems and network management are usually related to each other because the traffic loads of P2P systems have to be controlled to avoid regular network services becoming unavailable due to network congestion. In this context, from a network operation point of view, P2P systems often mean problems. In this article we take a different perspective and look at P2P technologies as an alternative to improve current network management solutions. We introduce an approach where P2P networks are used as flexible interdomain distributed management systems able to provide facilities often absent in traditional management systems. We present three examples of P2P-based network management: how managed network views can be shared in a P2P network, how P2P-based application layer routing can improve connectivity between management entities, and how groups of peers can be used to process management tasks in a balanced way. Our main goal is to show that P2P technologies are a feasible tool for network management, and thus motivate further investigations on the subject.

INTRODUCTION

The management of computer networks is essential to maintaining the health of many modern companies and organizations. Networks that suffer from weak management usually experience performance problems and low availability, which often lead to economic losses. The current, established network management solutions are suitable for monitoring and controlling devices and services located within a single common administrative domain. For example, the Simple Network Management Protocol (SNMP) framework [1], standardized by the Internet Engineering Task Force (IETF), is a widely deployed solution to monitor network traffic on border routers or to map servers with proper SNMP software installed.

With the increasing number of companies interconnecting with one another and the high adoption of the Internet as a transit network, management actions that cross administrative domain boundaries become an important need. For example, for successful high-quality videoconferencing, network bandwidth must be currently allocated in all administrative domains in the path between end users. In each domain, network operators with proper management rights and tools configure the domain's internal devices in order to allocate the required bandwidth. However, since the management rights of each operator are confined to the operator's domain, videoconferencing end-to-end bandwidth allocation requires cooperative management between domain operators. Current management systems, however, fail to provide proper support for such interdomain cooperative management because they were designed to operate in the limited environment of local domains.

Unlike traditional network management systems, peer-to-peer (P2P) systems are designed to allow interaction among users located in different domains. Users of P2P systems consciously or unconsciously cooperate in order to accomplish a P2P task, such as a file download or a distributed calculation. Typically, a P2P user installs an application (i.e., a peer) in his/her desktop to both share local resources (e.g., files or processing power) and access remote resources from other peers, possibly located in remote domains. Peers and their connections form P2P networks where user resources are discovered and shared. Thus, P2P technologies provide opportunities for real-time communication, ad hoc collaboration, and information sharing in large-scale distributed environments, such as the Internet [2].

When P2P systems and network management are investigated together, it is usually because P2P traffic may cause severe network congestion on physical networks when it surpasses regular traffic at dangerous levels. In such conditions P2P traffic needs to be controlled to avoid undesired situations [3]. In this article, however, we look at P2P and network management from another perspective: the services introduced by

P2P networks present innovative features that can be used not only for file sharing or distributed computing, for example, but also for helping in solving problems of other critical disciplines, such as network management. Thus, we take P2P as a management tool able to help operators of different domains to accomplish their tasks in a cooperative fashion.

Our main objective with this article is to introduce and motivate further investigations on applying P2P technologies for network management. Toward this goal, we present how P2P can be used as a management tool through three examples. First, we show that P2P may provide the aforementioned interdomain cooperative management support using management peers. Second, we show that P2P application layer routing can improve connectivity among management entities (e.g., managers and agents) distributed across different domains. Finally, we present how the constantly increasing processing power required to execute management tasks can be provided by flexible, dynamic groups of management peers. At the end, we list additional areas and opportunities for P2P-based network management research, before finishing this article with concluding remarks.

P2P NETWORKS AND TECHNOLOGIES

P2P networks are popular Internet applications that enable users to share computing resources [2]. Today, P2P, Web browsing, and email can be put together in the same class of importance. Napster is the tool probably most responsible for the popularization of P2P networks due to its music sharing support, which led to a controversial lawsuit with the music industry in 1999. So far, several Napster clones and general file sharing systems have been created, such as KaZaA, BitTorrent, and eMule. Other P2P networks are available on the Internet for other applications, such as distributed processing (e.g., SETI@home), collaborative work (e.g., Groove), instant messaging (e.g., ICQ), and voice over IP (e.g., Skype).

P2P networks are composed of software nodes (peers) installed on the edges of the Internet, usually in inexpensive user hosts along diverse administrative domains. Each peer communicates with other peers directly or indirectly to share local resources (e.g., memory, files, processing power), forming robust distributed systems over the Internet. The P2P communication model clearly differs from the client-server model. In client-server, servers store resources accessed by clients (Fig. 1a); by contrast, in P2P each peer is itself a repository that shares its resources with other peers (Fig. 2b), thus forming environments where resources are distributed among peers acting as both servers and clients.

Client-server addressing is based on the Internet Domain Name System (DNS), while P2P networks employ addressing schemes independent from DNS to index peers, users, and resources. Client-server messages are routed at the network layer by Internet routers (Fig. 1a),


Figure 1. *Routing and resource distribution in: a) client-server; b) P2P.*
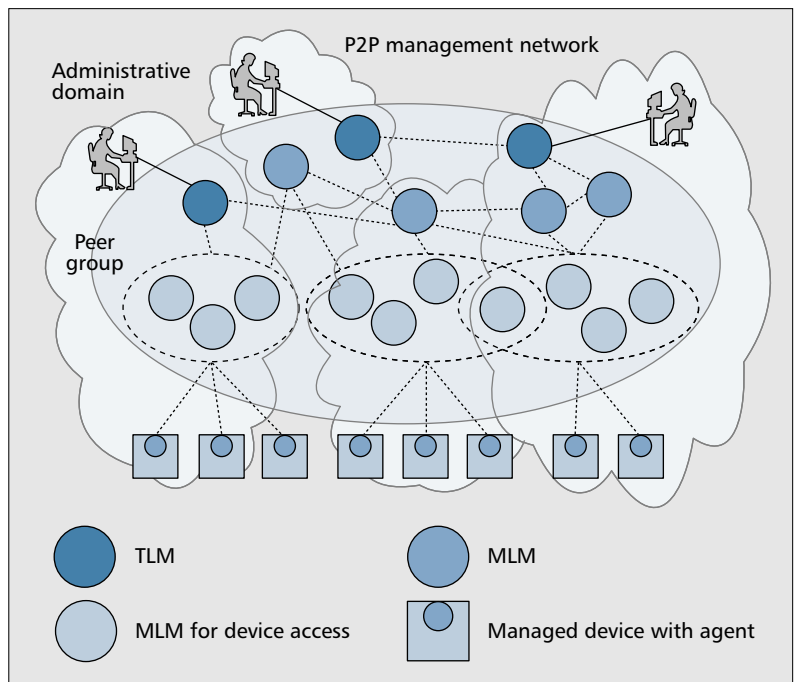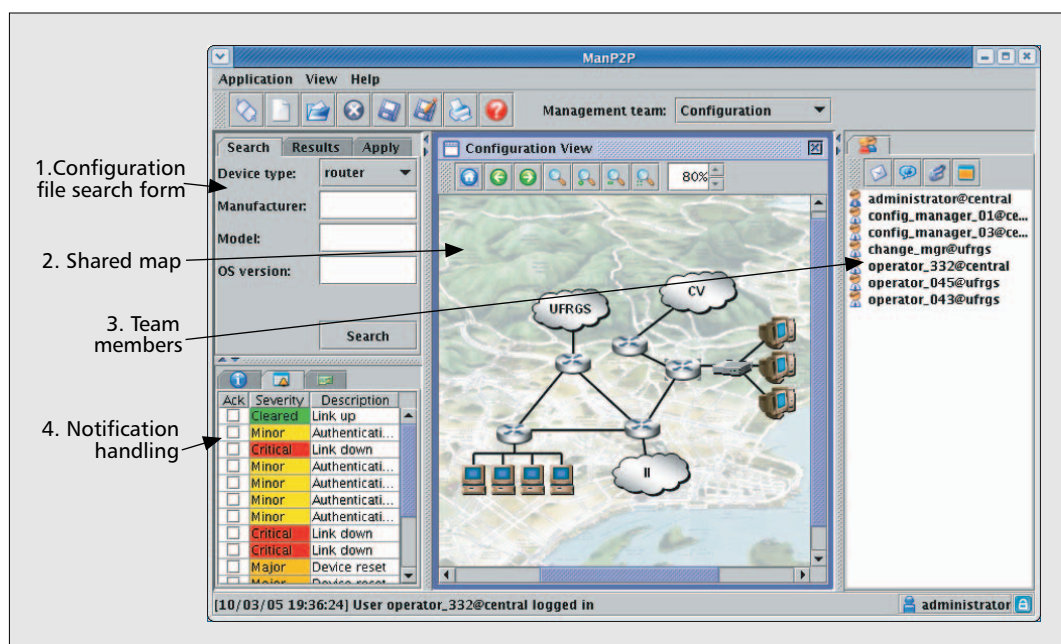

Figure 2. *P2P-based network management scenario.*

but P2P messages are routed at the application layer by routing peers. For example, in Fig. 1b, if the bottom path from the leftmost peer to the rightmost peer is congested, the alternative top path can be used if the uppermost peer is used to route messages.

Having routing and addressing built on top of Internet protocols, P2P networks are quite dynamic environments where new peers can join the network when they are started in user hosts, and leave when users shut peers down. This dynamic nature makes P2P networks both unpredictable and scalable: unpredictable because resources may be unavailable if the peers that host them are down; scalable because the expansion of the network capacity requires only the addition of new peers.

1. Configuration file search form

2. Shared map

3. Team members

4. Notification handling

■ **Figure 3.** *ManP2P TLM peer GUI.*

A SIMPLE MODEL FOR P2P-BASED NETWORK MANAGEMENT

Traditionally, network management is accomplished through the manager-agent model. Managers contact agents to access managed objects on target devices. Managed objects model the managed resources of target devices such that the implementation details of these resources are hidden. Agents, besides reacting to managers' requests, can asynchronously report to communities of managers the status of managed devices on specific events, such as a link down or an intrusion attempt. Standardization of the manager-agent model has been widely supported, for example, by the IETF through its SNMP framework [1], and the International Telecommunication Union (ITU) through the open systems interconnection (OSI) system management X.700 series Recommendations [4].

The manager-agent model can be extended to define more sophisticated architectures. For example, management by delegation (MbD) [5] extends manager-agent by introducing mid-level managers (MLMs) to provide more flexible and scalable management. In MbD, managers delegate the execution of tasks to MLMs located closer to agents (e.g., transferring management scripts), which reduces network bandwidth consumption and decentralizes the execution of management tasks. P2P-based network management extends traditional management models by merging in such models the services introduced by P2P networks. This P2P and network management merging allows the use of P2P features to improve a number of management tasks, such as interdomain cooperative management cited before.

The following sections present three examples of P2P-based network management. We first define a simple P2P-based management model composed of the following entities:

- **Top-level manager (TLM)**: a peer that, reacting to human operator requests, communicates with other management entities in order to accomplish a management task
- **Mid-level manager (MLM)**: a peer that only reacts to TLMs' or other MLMs' requests
- **Agent**: a software daemon that retrieves/changes (upon request) and reports (spontaneously) the status of managed devices

Although top-level manager (TLM), mid-level manager (MLM), and agent are entities well known in the network management field, this simple P2P-based model explicitly defines that TLMs and MLMs are dual-role entities: they are regular managers executing management tasks, and also communicating peers of a P2P management network. This mixed nature allows TLMs and MLMs to manage a computer network using not only conventional management techniques but also the services introduced by P2P networks. Also, we explicitly assume that TLMs react to human operators, thus playing the role of management front-ends, while MLMs react only to TLMs or other MLMs, with no direct human contact, but complementing the management process. This differentiation between TLM and MLM based on human interaction helps to present the TLM and MLM behavior when supporting the following management examples:

- **Human-based cooperative management**: where human network operators interact with one another via TLMs in order to accomplish a management task in a cooperative fashion
- **Improved connectivity for message exchange**: where management entities (TLMs, MLMs, and agents) use a message routing service more flexible and reliable than that available on TCP/IP networks
- **Management tasks load balancing**: where MLMs share their processing power to support the execution of management tasks

Figure 2 shows a scenario where TLMs, MLMs, and agents are part of a P2P management network designed to support the previously listed management examples.

Network operators from Fig. 2 request management services to the P2P management network accessing TLMs. TLMs implement cooperative management, thus enabling operators to exchange management information, regardless whether TLMs are directly or indirectly connected. TLMs also interface with MLMs that route management messages to agents on target devices. MLMs can also execute more complex tasks when required, and since management tasks can consume a large amount of processing power, groups of MLMs can be organized in order to share computing power in a balanced way. In this case management tasks are not delegated to a specific MLM but to a group of MLMs. In the next sections these P2P-based management examples are presented in more detail, and some aspects of a new management tool named ManP2P, which we developed to investigate P2P-based management, are discussed.

P2P FOR HUMAN-BASED COOPERATIVE MANAGEMENT

Human-based cooperative management tools are increasingly required because they help independent operators to accomplish a common task. For example, in large private networks, geographically dispersed operators with complementary roles (e.g., security and change managers) need to cooperate to lead the managed network to a secure state. On the Internet, multimedia applications that require interdomain resource allocation also benefit from cooperative management if operators involved in resource allocation exchange management information in a flexible and fast way.

Observing the scenario of Fig. 2, TLMs can communicate with all peers in the P2P network including all other TLMs, which enables interaction between online operators regardless of their administrative domains. Thus, in this environment traditional P2P applications can be adapted and/or merged to promote cooperative management. File sharing, instant messaging, and voice chat are some examples. Figure 3 shows the graphical user interface (GUI) of the ManP2P tool we have developed to support configuration management, shared network views, and distributed notification handling.

To help in **configuration management**, device configuration files are shared among operators using P2P file sharing facilities similar to those found in KaZaA and eMule. Each ManP2P peer indexes configuration files according to device type, manufacturer, model, and operating system version. An operator can query the ManP2P network entering information in the search form (Fig. 3). Collected results from remote peers are then presented to the operator to download a configuration file. The local copy of a downloaded file can be edited and the final version applied to configure target devices. When indexing, searching, providing, or retrieving configuration files, the ManP2P tool assumes its role of peer; when deploying a configuration, it works as a manager.

Usually, management tools present views of the managed network in network maps. Less usually, however, network maps can be shared among operators. To support **shared network views**, ManP2P peers allow operators to define what local devices are visible to remote peers (Fig. 2). In addition, operators are members of virtual management teams that share common network views. Each member in a team has an access profile (enforced by a team leader) to grant proper administrative privileges when accessing devices in shared maps. Also, views present not only network maps, but also a list of online team members (Fig. 3), which allows them to be in touch using an instant messaging facility similar to ICQ.

For **notification handling**, once a device is selected in the current map, a list of received notifications is displayed (Fig. 4). Active alarms are shown in a different fold of the same notification area. For different operators accessing the same device, different alarms can be presented according to the operator profile. For example, the security manager would not be able to receive notifications or see alarms related to performance management. Again, which notifications are received by which operators depends on each operator profile defined by team leaders.
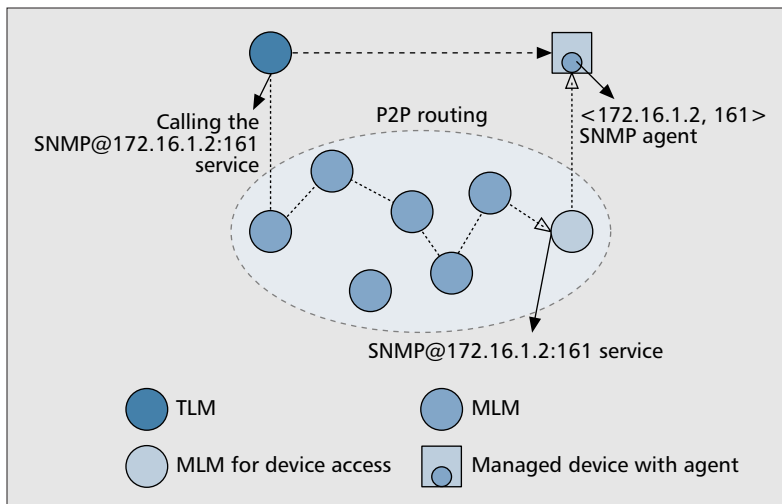
Team leaders are important to the distribution of management roles among team members. In fact, different levels of leadership can exist. For example, some members can be leaders only able to include new operators in a team, while other leaders would define the roles of new members. Thus, each member ends up with two profiles: one defining the administrative privileges when accessing managed devices, and another defining the member role in the team leadership hierarchy. Team support improves traditional management systems by including the software support to define very dynamic communities of independent and cooperative operators sharing common goals.

IMPROVED CONNECTIVITY FOR MANAGEMENT INFORMATION EXCHANGE

In traditional network management, managers and agents use network layer routing to communicate with one another. When considering inter-domain management, however, network layer routing can lead to low or even absent connectivity between management entities located in different domains.

Low connectivity comes from the fact that although several physical paths may exist between a manager and an agent in IP networks, only one route is selected for message delivery. If this route fails, manager and agent communication collapses until the network recovers from the fault. Also, if the route is congested, manager and agent are unable to select an alternative uncongested route. Dedicated management networks can be used to solve these low connectivity problems, but they are often expensive and confined to a single domain. This means that several adjacent domains all employing dedicat-

■ Figure 4. *Improved connectivity using P2P routing.*

ed management networks to support interdomain management is unlikely to happen.

Total lack of connectivity happens when management entities are isolated inside administrative domains due to security issues. For example, SNMP messages rarely cross domain boundaries because Internet firewalls block SNMP traffic to avoid unauthorized external investigation of the contents of unencrypted SNMP messages. Lack of connectivity is also the result of introducing boxes that break the network layer logic, such as Network Address Translation (NAT). If the addresses of internal devices are mapped to a smaller set of external valid network addresses, there is no way for external managers to discover which valid address has to be used to reach the target internal device. Again, network layer routing renders interdomain management difficult.

In order to bypass network layer routing, application layer routing could be used instead. In this case managers and agents can exchange information using intermediate management proxies [6], which can be seen as MLMs that receive management messages from TLMs and forward them to other MLMs toward the target agent.

Application layer routing is actively investigated in P2P research to support, for example, fast resource discovery [7]. P2P routing can select, between two peers, not only one route but several routes. In addition, routes may be defined based on criteria such as network bandwidth, peer processing power, and storage. Considering these advances, P2P networks can be used as a cheaper and more reliable mechanism to deliver management information between management entities distributed along different domains.

The ManP2P tool presented before can run in two different modes: standard and daemon modes. In standard mode it works as a TLM providing the management front-end. In daemon mode it becomes an MLM that executes delegated tasks to decentralize the network management from TLMs, and forwards management messages acting as an application layer router. The ManP2P tool is implemented using JXTA [8], a P2P framework that has protocols for application layer routing able to bypass Internet boxes such

as firewalls and NAT. In addition, JXTA provides modules to support interdomain secure communication taking into account privacy, authentication, and authorization issues. Although there is no study comparing JXTA and SNMP security, it is reasonable to figure that JXTA tends to better treat the interdomain security issue because its security functions have been designed to operate in the Internet: in contrast, the SNMP security model, only available in the not widely deployed SNMPv3, does not explicitly consider interdomain communication questions.

TLMs and MLMs are reached in the ManP2P network using the JXTA addressing scheme, where each peer is addressed by a uniform resource identifier (URI) that includes a unique peer identifier. Agents, on the other hand, are not ManP2P peers and can only be reached through network layer routing and addressing. In SNMP, for example, each agent is addressed using a <network address, transport port> tuple. To be accessed from a ManP2P TLM, however, each agent is mapped to a P2P service provided by an MLM *gateway peer* associated with the target device. In Fig. 4, for example, SNMP agent <172.16.1.2, 161> is mapped to service SNMP@172.16.1.2:161 exposed by a ManP2P MLM gateway peer. A TLM wanting to access the target agent must call this service on the MLM gateway peer. The intermediate MLMs, in turn, use JXTA routing support to deliver the messages between the TLM and the MLM gateway peer.

Although P2P improves connectivity between management entities, it may also impact management performance. Figure 5 compares the network traffic generated in downloading, to a ManP2P TLM located in a domain A, the routing table (with 7–120 rows) of a managed device located in domain B using SNMP and an MLM gateway peer (also located in B). P2P traffic is higher than SNMP because JXTA messages are text-encoded, while SNMP employs shorter binary-encoded messages. Figure 6 presents the retrieval delay. For more than 23 rows, P2P traffic performs better than SNMP because the MLM gateway locally exchanges several SNMP messages with the target device to build up a single P2P reply sent to the TLM. Without P2P, however, the caller TLM has to remotely interact with the target device with no intermediate MLM. Since more hops now separate the TLM from the target device (in our tests, TLM and target device are 10 hops distant from each other), the retrieval delay for each row increases, thus increasing the total delay. Although these results are restricted to routing table retrieval, they indicate that with proper manipulation (e.g., additionally compressing P2P messages), P2P may not introduce drastic performance drawbacks.

DISTRIBUTING MANAGEMENT TASKS USING GROUPS OF PEERS

An interesting feature of P2P networks is that a service can be provided by a group of peers instead of by a single peer. The services of a group are available to the remaining P2P net-

work as long as there is at least one peer active in that group. Groups of peers then improve the availability of services employing the redundancy of peers in a group. Thus, groups can provide better availability for network management services too. To strongly avoid management services unavailability, a P2P management network additionally needs to check that groups will not become empty by losing all of their peers. In this case groups themselves can check, using a self-monitoring process, for a reasonable number of active peers. If this number falls to a critical value, groups notify TLMs indicating that the provisioning of services is in danger; active peers may further try to borrow peers from other groups in order to have more peers providing the services in danger.

The main advantage of using groups of peers, however, is related to performance. When a single peer is responsible for managing several network devices, performance problems can easily arise. Tasks such as network analysis, script execution, and accounting can demand a high number of computing resources. The use of peer groups allows the distribution of the device management load among peers. If a group of management peers needs more computing resources, other peers can be instantiated in additional computers or borrowed from other groups.

Peers in a group can themselves control the distribution of management tasks. This contrasts with traditional management by delegation, where TLMs are the entities responsible for selecting MLMs in order to have a proper balanced distribution of management tasks. In the case of P2P-based management, groups of peers provide such balancing on their own by selecting which MLM will perform a task when it arrives, freeing TLMs from this decision. From the perspective of TLMs, this is an interesting P2P feature because a group of peers is seen as a single high-processed MLM able to execute more processing-demanding tasks delegated by TLMs.

Some techniques for load balancing (e.g., Weighted Round-Robin Scheduling and Weighted Least-Connection Scheduling) can be used within a group of peers [9] to internally distribute delegated tasks. The ManP2P tool uses the Weighted Round-Robin Scheduling algorithm to distribute the management load. Each peer informs all other group members about its computing resources. Peers with more available resources are selected to execute management tasks first. Just after initiating a new task, a peer has fewer resources available, which may cause the selection of another peer to execute the next task.

ADDITIONAL ISSUES IN P2P-BASED NETWORK MANAGEMENT

The previous sections have presented three examples of P2P-based network management. Other issues may represent further interesting research opportunities.

For example, peer software is more easily updated than the operating systems of network devices. Consequently, P2P networks are more flexible environments where traditional, well-

■ **Figure 5.** *Network traffic to retrieve a routing table.*

■ **Figure 6.** *Routing table retrieval delay.*

known management technologies could be revisited. Mobile agents for network management [10], for example, have been intensively investigated but have never seen the light of day in industry due to the infrastructure required on target devices. With P2P, however, such infrastructure can be provided by management peers able to execute mobile code. This approach may effectively enable mobile agents for network management, and investigations are required, for example, to evaluate the advantages/drawbacks of running mobile agents on a layer above the managed network.

Due to its ease of use, P2P can also be the basis for user-based management, where network users participate in the management process executing restricted actions. Users could use a light version of a TLM peer to join a P2P management network and participate as a passive party (e.g., checking the available end-to-end bandwidth) or active party (e.g., initiating backup DNS servers at the Internet service provider). Investigations could determine, for example, the management rights users may receive and the consequences of delegating such rights to these users.

P2P also has the potential to allow more flexible distributed monitoring, where peers incor-

Due to its flexibility, scalability, and ease of use, P2P can be used to build up overlay management networks composed of peers installed on inexpensive computers. Such P2P-based management networks can evolve and introduce facilities that are absent in traditional management systems.

porate the role of network monitors. More specialized peers could also support the definition of new distributed detection models for denial of service attacks.

SUMMARY

Currently, the management of computer networks can be accomplished by several technologies. SNMP is the de facto solution widely deployed, and XML and Web-services-based management have been attracting the attention of both academia and industry. In this article we have introduced the use of P2P technologies for network management.

Due to its flexibility, scalability, and ease of use, P2P can be used to build up overlay management networks composed of peers installed on inexpensive computers. Such P2P-based management networks can evolve and introduce facilities absent in traditional management systems. Thus, P2P management does not replace traditional models but improves them by enabling effective interdomain management through the deployment of management peers along different administrative domains, thus forming a flexible environment where traditional technologies can be revisited.

Three examples where management can be improved through P2P-based solutions have been presented. We have seen how network operators can team up to cooperatively administrate devices visible on shared maps, how connectivity between management entities can be improved using routes selected by P2P routing services, and how groups of peers can accomplish resource demanding tasks.

Our work is intended to be an initial step toward the introduction of P2P technologies in network management. P2P is an area of ongoing study and development, and other issues may represent interesting research opportunities in the context of network management.

REFERENCES

[1] D. Harrington, R. Presuhn, and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks," RFC 3411, Std. 62, Dec. 2002.
[2] S. Androutsellis-Theotokis and D. Spinellis, "A Survey of Peer-to-Peer Content Distribution Technologies," *ACM Comp. Surveys*, vol. 36, no. 4, Dec. 2004, pp. 335–71.
[3] S. Seng and J. Wang, "Analyzing Peer-To-Peer Traffic Across Large Networks," *IEEE/ACM Trans. Net.*, vol. 12, no. 2, Apr. 2004, pp. 219–32.
[4] ITU-T Rec. X.700, "Management Framework for Open Systems Interconnection (OSI)," Sept. 1992.
[5] G. Goldszmidt and Y. Yemini, "Distributed Management by Delegation," *15th Int'l. Conf. Distrib. Comp. Sys.*, May/June 1995, pp. 333–40.
[6] E. P. Duarte *et al.*, "Improving the Dependability of Network Management Systems," *Int'l. J. Net. Mgmt.*, vol. 8, no. 4, July/Aug. 1998, pp. 244–53.
[7] H. Balakrishnan *et al.*, "Looking Up Data in P2P Systems," *Commun. ACM*, vol. 46, no. 2, Feb. 2003, pp. 43–48.
[8] L. Gong, "JXTA: A Network Programming Environment," *IEEE Internet Comp.*, vol. 5, no. 3, May 2001, pp. 88–95.
[9] H. Bryhni, E. Klovning, and O. Kure, "A Comparison of Load Balancing Techniques for Scalable Web Servers," *IEEE Network*, vol. 14, no. 4, July/Aug. 2000, pp. 58–64.
[10] A. Bieszczad, B. Pagurek, and T. White, *IEEE Commun. Surveys and Tutorials*, vol. 1, no. 1, Sept. 1998.

BIOGRAPHIES

LISANDRO ZAMBENEDETTI GRANVILLE (granville@inf.ufrgs.br) is an associate professor at the Institute of Informatics of the Federal University of Rio Grande do Sul (UFRGS), Brazil. He received his M.Sc. and Ph.D. degrees, both in computer science, from UFRGS in 1998 and 2001, respectively. He has served as a TPC member (2003–2005), General Co-Chair (2004), and Steering Committee member (2005–2006) for the Brazilian Symposium on Computer Networks (SBC/LARC SBRC). Currently, he is chair of the special interest group on Computer Networks and Distributed Systems of the Brazilian Computer Society (SBC), and executive director of the Brazilian Laboratory on Computer Networks (LARC). He has served as a TPC member for IFIP/IEEE IM 2005 and IFIP/IEEE DSOM 2005, and is Poster Co-Chair of IEEE/IFIP NOMS 2006 and General Co-Chair of LANOMS 2005. His main areas of interest include policy-based network management, management using of Web services, network monitoring and configuration, and P2P-based services and applications.

DIEGO MOREIRA DA ROSA (diegoro@inf.ufrgs.br) received a B.Sc. in computer engineering from UFRGS in 2004. Currently he is completing a Master's degree in computer science at the same university. His main topics of interest include telecommunications network management, distributed management, and human-based cooperative management.

ANDRÉ PANISSON (panisson@inf.ufrgs.br) holds a B.Sc. degree (2003) in computer science from UFRGS). He is an M.Sc. candidate in computer science at the Institute of Informatics of UFRGS. His topics of interest are peer-to-peer computing, distributed management by delegation, and load sharing in network management.

CRISTINA MELCHIORS (cmelchiors@inf.ufrgs.br) is a Ph.D. student at the Institute of Informatics, UFRGS. She currently works on cooperative management using peer-to-peer technologies. She received a B.Sc. degree in informatics from the Pontifical Catholic University of Rio Grande do Sul (PUCRS) in 1995, and an M.Sc. degree in computer science from UFRGS in 1999. Her main research interests include cooperative management, distributed management, artificial intelligence applied to network management, and peer-to-peer computing.

MARIA JANILCE BOSQUIROLI ALMEIDA (janilce@inf.ufrgs.br) received her M.Sc. and Ph.D. degrees from UFRGS in 1986 and 1994, respectively. She is an associate professor of computer science at the Institute of Informatics of UFRGS, and chair of the Advanced Training Center of UFRGS in Gramado. She has been serving as a TPC member for the SBC/LARC Brazilian Symposium on Computer Networks (SBRC) since 1997, and was General Co-Chair of SBRC 2004. Her main topics of research include protocol specification, multimedia, and QoS.

LIANE MARGARIDA ROCKENBACH TAROUCO (liane@penta.ufrgs.br) is a full professor at the Center for Innovation and Use of New Technologies in Education of UFRGS, Brazil. She holds an M.Sc. degree (1976) from UFRGS and a Ph.D. degree (1990) from the University of São Paulo, both in computer science. She chaired the 7th Brazilian Symposium on Computer Networks (SBC/LARC SBRC 1989) and has been serving as a TPC member since 1983. She was chair (2002–2003) of the Educational Applications Workgroup of the National Research and Education Network, and has been chair of the Human Resources Workgroup of the Brazilian Internet Steering Committee (CGI-BR) since 1995. From 1983 to 1992 she was a representative of Brazil at IFIP TC6 (Communication Systems). Her main areas of interest are distributed network management and expert systems to support fault correlation.