2009 First International Conference on Evolving Internet

Using Metadata Snapshots for Extending Ant-based Resource Discovery Service in Inter-cooperative Grid Communities

Ye Huang*, Nik Bessis†, Amos Brocco*, Pierre Kuonen‡, Michele Courant* and Beat Hirsbrunner*
*Department of Informatics, University of Fribourg, Switzerland
Email: {ye.huang, amos.brocco, michele.courant, beat.hirsbrunner}@unifr.ch
†Department of Computing and Information Systems, University of Bedfordshire, UK
Email: nik.bessis@beds.ac.uk
‡Department of Information and Communication Technologies, University of Applied Sciences Western Switzerland
Email: pierre.kuonen@hefr.ch

***Abstract*—Much work is under way within the resource management community on issues associated with grid scheduling upon dynamically discovered information. In this paper we tackle the problem by exploiting a bio-inspired resource discovery mechanism, where information is provided by ant-based lightweight mobile agents traveling across a grid network and collecting data from each visited node. We start by providing the current state of the adopted grid scheduler, which is the result of an existing collaborative project named SmartGRID, and its underlying architecture constructed by ant-based mobile agents. We consider the problem of discovering resources in specific grid communities, which are bounded due to different shared community policies, such as diverse ant colonies, different resource discovery approaches, or other issues. Several issues have been raised during the design and implementation of such infrastructure. A notable issue, namely how grid schedulers from various bounded grid communities could be used in a manner which would extend current SmartGRID functionality is identified. Our shared view is that by utilizing already discovered and stored grid nodes's metadata snapshots in the first instance we can facilitate a more convenient and efficient resource discovery operation next time. With this in mind, our paper goes on describing our shared vision with regard to this extended functionality as well as discussing the new conceptual basis and its model architecture.**

***Index Terms*—SmartGRID; Intra and inter cooperative grid architecture; MaGate; Ant-based swarm intelligence; Metadata snapshots.**

I. Introduction

SmartGRID is a cooperative project aiming at increasing the efficiency, robustness, and reliability of heterogeneous grid computing infrastructures [1] concerning volatile and dynamic resources. The proposed grid middleware has been designed as a generic and modular framework supporting intelligent and interoperable grid resource management using swarm intelligence algorithms and multi-type grid scheduling. SmartGRID uses a layered architecture and aims at filling the gap between grid applications, which act as the resource consumers, and the grid resource low-level management systems, which behave as the resource providers. To achieve this goal, SmartGRID uses an autonomic and evolutional grid community composed of its grid schedulers, the MaGate scheduler [2].

Two approaches are currently available for discovering candidate nodes for a specific task. The first approach assumes that each node has partial knowledge of up to six direct neighbor nodes, which are maintained by the resource discovery service of the host node. When the host node requires the discovery of a remote node with the required features, in order to delegate a job, the direct neighbors list is exploited with the node negotiating job delegation with each member.

In the second approach, when the host node joins an existing bounded grid community, it prepares a profile named *agreement offer* to disseminate its capabilities across the community. We assume that the aforementioned public profile is kept up-to-date through nodes lifecycle. This approach is not limited to one community only: each individual node is free to join multiple existing grid communities, thus publishing different capability profiles.

Currently the discovered information for each specific task is discarded after its usage. We aim at extending this model such that each node of a SmartGRID community might also be capable of keeping a metadata snapshot of known remote nodes, in order to facilitate a more efficient and intelligent behavior towards relevant scheduling decisions. Moreover, as the SmartGRID architecture strives to provide intelligent scheduling for the scope of serving the grid community as a whole, not just for a single grid node, our extended work is also concerned with the design of a scheduling strategy supporting the combination of various interoperable bounded grid communities.

In general, the mission of a grid scheduler is to discover appropriate resources for executing jobs across a grid community. Our vision is that of a wider grid community scheduling process is able to exploit resources in large and partially unknown grid communities, and dealing with continuously changing job queues. Thus, since each community node is supposed to receive jobs from both its local and remote grid communities, management of the job queue must deal with a more dynamic, fluid and unexpected environment. It is our goal to ensure robustness, reliability, efficiency, and

978-0-7695-3748-1/09 $26.00 © 2009 IEEE
DOI 10.1109/INTERNET.2009.21

intelligent scheduling response. In this respect, the scheduler must compromise between accepting community jobs and local ones, depending on its workload, *agreement offers*, and the ratio between resources contributed to the local and the global grid communities.

The goals of the presented work are multifold. First, a discussion on the current state of the implementation of our interoperable grid scheduling approach and its underlying ant-based resource discovery architecture is provided. Next, we describe how the ant-based resource discovery can be extended across diverse bounded grid communities in a manner which would enhance its current functionality yet maintaining its robustness, reliability, and efficiency.

In particular, we propose to exploit already discovered grid nodes and store metadata snapshots that would facilitate more convenient, efficient and intelligent subsequent resource discovery operations. With this in mind, the paper goes on to describe the concepts of our improved solution as well as to discuss an extended model architecture to be used for the next generation grid scheduler.

The remainder of the paper is organized as follows: current work concerning the SmartGRID project is introduced in the next section. Section III details a shared vision for extending the model, along with some insights to future development. Finally, the presented work is summarized in Section IV.

II. Current Work

Currently, the SmartGRID architecture comprises of three parts: the Smart Resource Management Layer (SRML), the Smart Signaling Layer (SSL), and the Data Warehouse Interface (DWI).

The SRML is an interoperable grid scheduler community composed of engaged decentralized grid schedulers named on MaGates (Magnetic Gateway), designed to be modular and emphasizing scheduler interoperation. With the infrastructure information retrieved from the DWI, MaGates discover and connect to each other, so as to collaborate in order to bridge heterogeneous grid systems with a consensual view. The grid community evolves dynamically, and is able to automatically recover from failure situations. Information about available resources and network status is gathered by the SSL, and stored into DWI's distributed data storages. The SSL maintains an overlay network of Nests that provide the runtime environment for the execution of bio-inspired ant algorithms. This approach provides an adaptive and robust signaling mechanism, supporting both grid resource discovery as well as monitoring. The layered architecture of the SmartGRID is shown in Figure 1.

The SRML is responsible for grid level dynamic scheduling and interoperation that provides grid applications with scheduling decisions on dynamic discovered computing resources. Being the core of SRML, the MaGate is also in charge of propagating resource discovery related tasks to the SSL, analyze the returned results, and decide future operations.

The SSL [1] represents the interface from and to the network of the SmartGRID architecture, by providing access

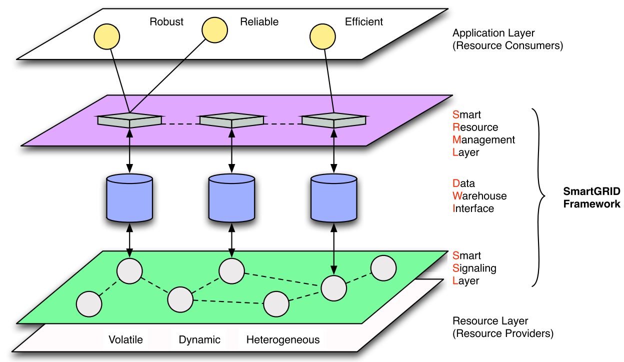

Fig. 1. SmartGRID architecture overview

to Virtual Organization (VO) resource. The SSL is controlled by the SRML, and provides information about the availability of other resources on other nodes, as well as their status. From the SSL point of view, each node has some partial knowledge of the underlying logical network. Remote nodes that fall into this partial view are called *direct neighbors*, because they are considered as having good connection with the host node. The SSL hides the complexity and instability of the underlying network by offering reliable services based on distributed ant algorithms. Ant algorithms do not require centralized control, and are known to be robust and adaptable, thus well suited for dynamic networks. Ants are defined as lightweight mobile agents traveling across the network, collecting information on each visited node: a distributed middleware named Solenopsis [3] provides an environment for the execution of ant colony algorithms, in particular the specific designed BlåtAnt collaborative ant algorithm [4] [5]. The activity of the SSL can be either reactive or proactive. Reactive behaviors are controlled by incoming requests from the SRML: information is asynchronously transmitted through a data warehouse interface, and fetched by the local nest. The same interface is used to provide feedback and results on the execution of algorithms. Continuous pro-active activities, such as network monitoring, are used to enhance the QoS of provided services for the SRML, and the robustness of the whole system.

III. A Shared Vision for Extending the Model

This section aims at addressing two specific conceptual aspects, including a contrast between current and extended SmartGRID topologies with regard to the logic of resource discovery. On one side, we focus on additional resource discovery logic as a method to enhance results and improve job delegation, while on the other side we describe its implementation within the MaGate.

We introduce the notion of *critical friends* as a mean of describing the interaction between nodes in a wider, larger-scale and unknown grid community. Particular attention is also given to the functional features of resource discovery of each MaGate scheduler. The idea is to empower the existing resource discovery mechanism, based on ant algorithms, by using a cache of previously visited remote nodes. This additional information enables MaGates to take more intelligent scheduling decisions. To discuss the extended model architecture, an aggregative case scenario is then presented.

A. Resource Discovery Topologies

The current SmartGRID network topology implies that propagated ants are looking towards resource discovery within a specific grid community, which is bounded due to various reasons, such as shared community policy, trust issues, geographical location, etc. Let us now label this bounded grid community as a Virtual Organization 1 (VO_1). In a similar vain, let us assume that there are a number of separated VOs across a wider (larger-scale and thus unknown) grid community (VO_1, ... VO_n). Let us also assume that an individual node is member in more than one VO (e.g., VO_1, VO_2) and that each node within a VO can be a service consumer, service provider or both.

This view enables a network topology, which clearly extends the aforementioned SmartGRID topology. This is mainly due to the fact that the current SmartGRID framework takes job delegation decisions on the basis of ants searching across one and only one VO (e.g., VO_1) and it does not take the full advantage of the fact that a node in a VO_1 can be also a node in a VO_2. In such a grid community, a node, for example, n_1 in VO_1 can communicate with another neighbor node, n_5 in the same VO_1. The rationale is that a node, n_5 can be also a member of another VO_2, which in turn leads to the rationale that n_5 can communicate with another neighbor node such as n_9 (that is also a member of VO_2 and VO_3) and so on. The assumption here is that communicating and/or delegating a job to a node belonging to a different VO will not result to a conflict of interest between parties. Having said that, the assumption is valid given the fact that a VO should not allow membership of a distinct node been part of two conflicting VOs unless it is unknown or there is a certain level of trust. In the case of the latter point, the assumption is still valid given the fact that the associated *agreement offer* and policies explicitly specify the range of act of a job delegation (what is acceptable).

Our vision is also based upon the very important notion of *Self-led Critical Friends* (SCF). This concept is built upon relations between nodes, and the knowledge that each node constructs about some neighbor node (either a member of the same or of a different VO) based on previous (direct and/or indirect) interactions, such as communications and delegations. The notion of previous interactions between neighbor nodes determines the strength of the relation and ultimately the level of Critical Friendship. We thus consider a topology of a wider dynamic grid community, based upon a series of SCF relations between nodes from different communities; the strength of a relation between two nodes can be either constant or can evolve over time and influence decisions of the job delegation task. This idea is similar to the one proposed in [6], and effectively creates a second level overlay that can be exploited for efficient resource discovery.

The concept of SCF mirrors the notion of relationships occurring in the real world. If a person (node in our case) is looking for a specific service and they do not know how to find it, they will ask some of their friends (neighbor nodes) who may know it (decision based on past experience). If they do not, these friends will pass the query on to their own friends with the view that someone across the "friends" network (neighbors network in our case) will know and have information relevant to the original request about the specific service. Based upon this information a decision can be made. We purposely moved away from a centralized authority as this could lead to a single point of failure. To explain further, a centralized authority could be compromised by an external entity, and if all users are dependent upon this entity then the functionality of the whole network could fail very quickly. This extended resource discovery topology clearly increases scalability and thus is the most suitable solution for wide grid communities.

Figure 2 shows the current and extended views of these resource discovery topologies. In brief, we aim at extending resource harvesting methods and transform the current intra-cooperative, e.g., VO_1, capacity of the SmartGRID architecture by incorporating an inter-cooperative, e.g., $VO_{1,2}$, capacity. This alternative is not designed to be mandatory, but needs to integrate well into the existing architecture.

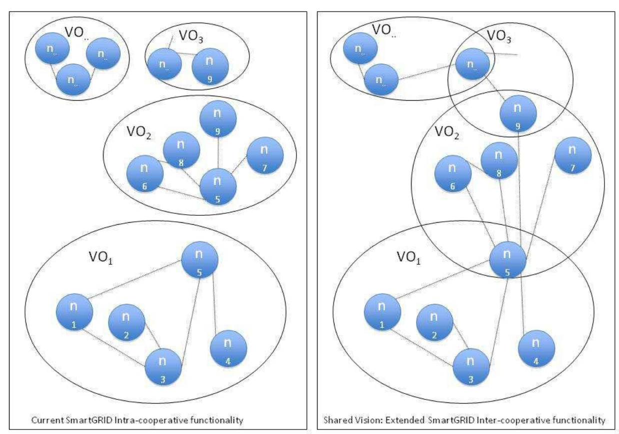

Fig. 2. Current and Extended View of Resource Discovery Topologies

B. Metadata Snapshots

This section addresses the second aspect, namely how to utilize current functionalities of the MaGate to extend the resource discovery technique. That is, using metadata snapshots stored in previously visited nodes' cache.

As mentioned earlier, every node within the grid community must publish its capabilities; moreover, the assumption is that this public profile is kept continuously up-to-date. We also explained that the MaGate scheduler generates directives to propagate ants in the SSL in order to provide nodes with availability and status information. Current MaGate resource discovery service either utilizes partial knowledge stored in a node's cache, or delivers search queries to match individual job delegation requirements with published node profiles from the known bounded grid community. To achieve this, the SSL is employed to hide the complexity and instability of the underlying network by utilizing ant algorithms. These ants function as lightweight mobile agents traveling across the grid network, collecting information on each visited node.

We now want to extend current functionality by enabling the MaGate resource discovery service to utilize more knowledge that can be made available from the visited nodes. In this respect, we propose that for every time a node, n_1, delegates a job to another node (with no regard to the possibility that it might be a member in multiple VOs) such as node n_3, node n_1 is required to keep an instance profile with regard to the parameters which have been used to discover it as a resource originally, as well as, the quality of service provided by node n_3. In a similar way, it is expected that every time a node, n_1, delegates a job to n_3, node n_1 has to update the profile about n_3 in its cache. We suggest that this is a bi-directional commodity and thus, we expect that node, n_3 will also keep an instance profile about n_1 in its cache. We also assume that a node n_1 stores as many profiles in its cache as the number of previously visited nodes. Our shared vision is to enable nodes storing meaningful information that can help assist them and their critical friends at a later stage.

Apparently, these profiles residing in each cooperating node can be sustained or even evolve over time and act as the tool towards decision making for delegating a job next time. That is, a calculation as an aggregative and weighted value representing the (strength or else) critical friendship relationship between nodes that previously cooperated, could significantly improve the decision making towards a job delegation to a particular node (or cluster of nodes) that is available from a pool of discovered resources across the wider grid community. Such a notion clearly provides a richness not seen in any other resource discovery or job delegation model.

C. SmartGRID Extended Model Architecture

We are interested here in describing the novel inter-cooperative process and the related events sequence in order to illustrate the extended SmartGRID architecture. This involves ants as agents acting on behalf of neighboring nodes in order to enable the MaGate scheduler to discover, decide and assign job delegations on suitable resources. To achieve this, we are using the aggregative case scenario (ACS) below. Figure 3 illustrates this low-level architecture.

ACS: Let us assume that a VO_1, consists of 8 nodes (nodes n_1 ... n_8). Let us assume that a VO_2, consists of 6 nodes (nodes n_7, n_9, ... n_{13}). Let us also assume that a node, n_1 in VO_1 wishes to delegate a job to another node $n_?$, using ants a_2, a_4 and a_7, which are propagated according to queries issued by the MaGate resource discovery service. The following sequential steps describe this novel, inter-cooperative process:

1) Node n_1 in VO_1 invokes its ants a_2, a_4 and a_7;
2) Ant a_2 contacts the neighboring node n_2 (that is a critical friend: $cf_{1,2}$), ant a_4 contacts the neighboring node n_4 (that is a critical friend: $cf_{1,4}$) and ant a_7 contacts the neighboring node n_7 (that is a critical friend: $cf_{1,7}$, note that node, n_7 is a member of both VO_1 and VO_2);
3) Ant a_2 reads and collects the public availability profile, as well as its metadata snapshots about previous job delegation activities completed in node n_2 that are available from the cache of node n_2;
4) With discovered metadata snapshots (available from the cache of node n_2) by ant a_2, MaGate n_1 realizes that node n_3 is a $cf_{3,2}$; moreover, n_3 has the capacity to take the job delegation task jd_1;
5) Ant a_4 reads and collects the public availability profile of n_4, as well as its metadata snapshots about previous job delegation activities completed in node n_4 that are available from the cache of node n_4;
6) With discovered public availability profile by ant a_4, MaGate n_1 realizes that node n_4 has no capacity to take the job delegation task jd_1;
7) With discovered metadata snapshots (available from node n_4 cache) by ant a_4, MaGate n_1 realizes that node n_9 is a $cf_{9,4}$; moreover, n_9 has the capacity to take the job delegation task, jd_1;
8) Ant a_7 reads and collects the public availability profile of n_7, as well as its metadata snapshots about previous job delegation activities completed in node n_7 that are available from the cache of node n_7;
9) With discovered metadata snapshots (available from the cache of node n_7) by ant a_7, MaGate n_1 realizes that node n_9 is a $cf_{9,7}$; moreover, n_9 has the capacity to take the job delegation task jd_1;
10) Ants a_2, a_4 and a_7 collect profiles about nodes n_3 and n_9 for MaGate n_1, which reports such information to a virtualized data warehouse;
11) We now assume that a calculation as an aggregative and weighted value representing the (strength or else) critical friendship relationship between previously cooperating nodes, $cf_{3,2}$, $cf_{9,4}$ and $cf_{7,9}$ has significantly improved the decision making towards jd_1 to node n_9. This is due to the aggregative cf values that have suggested that node n_3 has been delegated x number of jobs and the satisfaction (confidence) level was significantly less than the satisfaction (confidence) level provided for an equal number of past delegated jobs in node n_9.

D. The Shared Vision Roadmap and Plan of Work

Our shared vision roadmap and plan of work involve a three phase implementation strategy.

During phase 1, we will focus on implementing a fully decentralized grid information discovery environment. The work will be based on and extend the current integrated

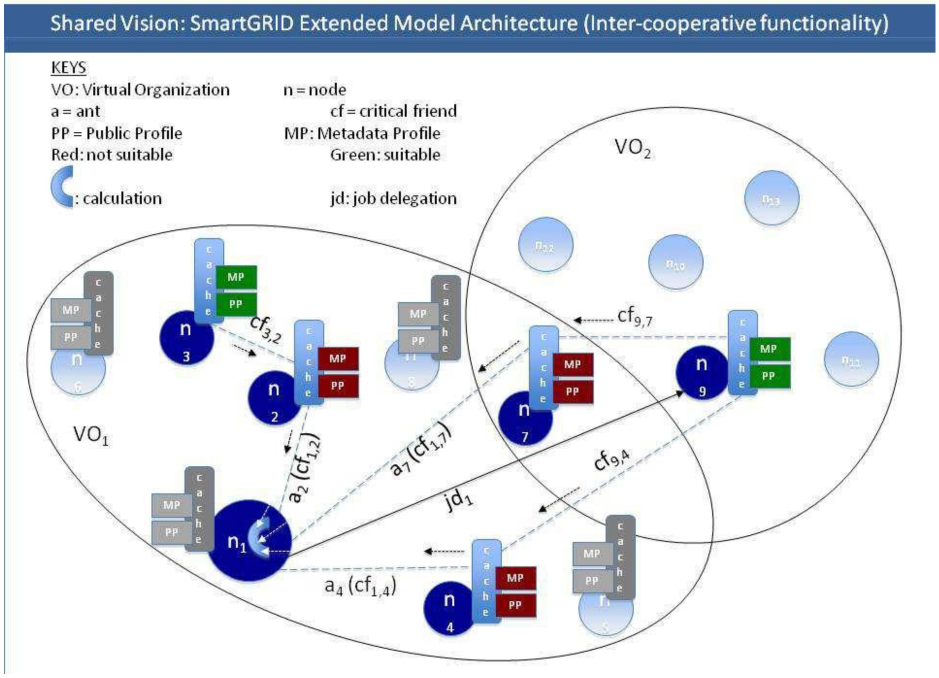

Fig. 3. SmartGRID Extended (Low-level) Model Architecture

simulation prototype of SmartGRID, that is currently being used to prove the functionality of scheduler interoperation and job delegation within the existing SmartGRID community. Our extended work here involves the additional development of the scheduler interaction protocol. We are currently considering the deployment of a WS-Agreement based approach for this purpose.

During phase 2, we will work towards the mechanism of WS-Agreement based job delegation between host node and remote nodes. As mentioned earlier, an agreement is initiated by the initiator (e.g., host node), and confirmed by the responder (e.g., remote node). The capability of negotiating job delegation between the agreement initiator and responder will require a simple and flexible workflow, which could decouple each side. Based on the decoupled architecture, details of resource configuration and information representation are hidden, even if different agreement initiators will get different information feedback from the same responder. In addition, the workflow would be flexible enough in order to automatically contact neighboring nodes with abundant community knowledge (router node) to get alternative candidate nodes in case of failure of the original responder.

Finally, during phase 3, we will work towards the creation of metadata snapshots, namely, identification and optimization of the representational form for data stored in each node cache. Moreover, meaningful utilization of metadata snapshots and thus, collecting and aggregating information from different neighboring nodes based on the critical friends notion will be developed to generate a series of dynamic, replicable, non-redundant, asynchronous and self-evolvable virtualized data warehouses for the larger-scale wider grid community. We will also consider developing an additional proactive mechanism using distributed agent push technology [7] in order to implement recommended resource discovery across the aforementioned wide grid community.

IV. Conclusion

We have conceptually extended a work related to the SmartGRID project, which has been developed to be a generic and modular framework to support intelligent and interoperable grid resource management using swarm intelligence algorithms. In this paper we offered a state-of-the-art regarding the adopted grid scheduler and its underlying architecture constructed by ant-based mobile agent technology within the context of discovering resources in bounded grid communities.

Our work herein addressed a notable case, namely how grid schedulers from various bounded grid communities could be used in a manner that would extend current SmartGRID functionality.

Our shared view is that by utilizing already discovered grid nodes that have stored metadata snapshots of past queries would facilitate a more convenient and efficient operation next time the MaGate resource discovery service starts looking for a suitable resource. With this in mind, our paper described our shared vision and discussed the new conceptual basis and its model architecture. This work has extended resource harvesting available methods and transformed the current intra-cooperative capacity of the SmartGRID architecture to the option of functioning in a novel approach, namely, functioning in an inter-cooperative capacity when and if it is required.

Acknowledgements

The work of SmartGRID is financially supported by the Swiss Hasler Foundation[1], in the framework of the ManCom Initiative (ManCom for Managing Complexity of Information and Communication Systems), project Nr. 2122.

References

[1] Y. Huang, A. Brocco, P. Kuonen, M. Courant, and B. Hirsbrunner. SmartGRID: A Fully Decentralized Grid Scheduling Framework Supported by Swarm Intelligence. In *Seventh International Conference on Grid and Cooperative Computing, 2008. GCC '08*, pages 160–168, China, 2008. IEEE Computer Society.

[2] Y. Huang, A. Brocco, M. Courant, B. Hirsbrunner, and P. Kuonen. MaGate: an interoperable, decentralized and modular high-level grid scheduler. Submitted to International Journal of Distributed Systems and Technologies (IJDST), 2010.

[3] Amos Brocco, Béat Hirsbrunner, and Michèle Courant. Solenopsis: A framework for the development of ant algorithms. In *Swarm Intelligence Symposium*, pages 316–323, Honolulu, Hawaii, April 2007. SIS, IEEE.

[4] Amos Brocco, Fulvio Frapolli, and B. Hirsbrunner. Blatant: Bounding networks' diameter with a collaborative distributed algorithm. In *Sixth International Conference on Ant Colony Optimization and Swarm Intelligence*, Bruxelles, September 2008. ANTS, Springer.

[5] Amos Brocco, Fulvio Frapolli, and Beat Hirsbrunner. Bounded diameter overlay construction: A self organized approach. In *IEEE Swarm Intelligence Symposium*, Nashville, Tennessee, USA, April 2009. SIS 2009, IEEE.

[6] Vicent Cholvi, Pascal Felber, and Ernst Biersack. Efficient search in unstructured peer-to-peer networks. In *SPAA '04: Proceedings of the sixteenth annual ACM symposium on Parallelism in algorithms and architectures*, pages 271–272, New York, NY, USA, 2004. ACM.

[7] Nik Bessis. *Grid Technology for Maximizing Collaborative Decision Management and Support: Advancing Effective Virtual Organizations*, chapter Model Architecture for a User Tailored Data Push Service in Data Grids. IGI, 2009.

[1] http://www.haslerstiftung.ch/