Beyond User Ranking: Expanding the Definition of Reputation in Grid Computing

Said Elnaffar
College of IT, UAE University, UAE
elnaffar@uaeu.ac.ae

***Abstract*-Shopping around for a good service provider in a Grid Computing environment is no less challenging than the traditional shopping around in non-virtual marketplace. A client may consult a service broker for providers that can meet specific QoS requirements (e.g., CPU speed), and the broker may return a list of candidate providers that satisfy the client's demands. If this computing platform is backed up by some reputation system, the list of providers is then sorted based on some reputation criterion, which is commonly the user rating. We argue in this paper that judging the reputation of a provider based on user rating is not sufficient. The reputation should additionally reflect how trustworthy that provider has been with respect to complying with the finalized SLA (using a metric called *conformance*) and how consistent it has been with respect to honouring its compliance levels (using a metric called *fidelity*). Accordingly, we perceive the reputation as a vector of three dimensions: user rating, conformance, and fidelity. In this paper, we define these metrics, explain how to compute them formally, and how to use them in the reputation-enabled framework that we describe.**

I. Introduction

Grid computing [5] is a service model in which providers provision computing resources, services and infrastructure management to clients as needed, and charge them per use. As a consequence of the rapid growth of such on-demand computing applications and the profusion of service providers, clients eventually have to select the provider that they can rely on in order to achieve their business objectives [2]. To that end, quality of service (QoS) metrics [1] are typically used to differentiate and rank service providers.

QoS-enabled Grids are typically associated with service level agreements (SLA) that, as binding contracts between the service provider and the service requestor, guarantee application performance quantitatively.

Throughput, Response time, and availability are among the many SLA metrics that can be negotiated and investigated when it comes to assessing the QoS a provider can offer to a client [1]. They comprise the main criteria of compelling a client to favour one provider over others. Unfortunately such set of metrics lack an attribute that measures how often a particular service provider was able to comply with the SLA in the past. Many researchers define reputation (or trustworthiness) of a service provider as the average ranking assigned by the clients who interact with that provider [8]. If a service provider is known to offer certain qualities over a period of time irrespective of its limitations, then it is assumed to have a good reputation. A reputation system is defined as a secure informative system responsible for maintaining a dynamic and adaptive reputation metric for its community. Grid players, such as service providers, brokers, and clients, continuously interact with the reputation system to establish a community ranking mechanism that co-operatively help them make future decisions based on the overall community experiences. However, we contend that designing a reputation system that solely relies on the temporal perspective of humans (i.e., the clients) can expose the system to dishonest ratings caused by the following types of users:

- *Emotional Reactors*: those clients who give non-subjective, inaccurate ratings influenced by some personal (could be temporal) issues with the service provider.
- *Bad Mouthers*: clients who unfaithfully exaggerate by giving negative ratings to service providers.
- *Ballot Stuffer*s: those clients who unfaithfully exaggerate by giving positive ratings to service providers.

Moreover, client views may not necessarily be able to capture the degree of inconsistency (variance) in the service provider's compliance levels over an extended period of time. These views are predominantly influenced by the most recent experience. Such vulnerability gives devious providers the opportunity to manipulate user perceptions psychologically. For examples, and in order to improve their reputation, providers may choose to boost their image in a season that witness soaring competition from their rivals by honouring the compliance levels promised to be delivered to users. As a result, users tend to express their satisfaction through that short period of time, which might be temporary, by giving high rating scores right after that pleasant experience. After that critical season elapses, service provider might opt to return to its relaxed policy with respect to adhering to the SLA. Such behaviour leaves the compliance profile of a given provider with superficial

K. Elleithy (ed.), Advances and Innovations in Systems, Computing Sciences and Software Engineering, 381–386.

positive and negative ratings over various points of time, lacking in-depth analysis.

In order to reflect providers' degree of compliance and their consistency at providing such compliance, we may need to aggregate some technical performance measures over time. This analysis can be done for two types of provider profiles: *local* profile and *global* profile. Each client locally maintains a private profile for each provider it deals with. This local profile reflects the direct personal experience that client had with a specific provider. The global profile, on the other hand, reflects the aggregate experience of the client community with a specific provider. Therefore, it is typically stored at a central global entity such as a Grid broker. Retaining the two profiles has manifold advantages. For example:

- In case of unavailability of the global profile, clients can still select providers based on their privately retained local profiles
- A service provider selection algorithm can take into account both profiles
- A client has the choice of adopting a selection criterion that assigns different weights to the two profiles. Suspicious clients may choose to give more weight to their personal experience, under the belief that other's ratings cannot be trusted.

The traditional QoS metrics such as throughput and response time are no longer adequate to effectively assess the goodness of provider as seen by the client [4]. Rather, the reputation of a provider should be factored in. However, as we explained earlier, banking on user ranking solely to represent the reputation of a provider has its own shortcomings. We believe that the overall reputation of a provider should be a combination of the following metrics:

- *User Ranking*: reflects the user view of the quality of service received from a provider. It is usually the ultimate means for capturing non-measurable aspects of quality. Therefore, it is deemed a qualitative metric as it is based on the human perception.
- *QoS Conformance*: measures the discrepancies between the projected QoS, as outlined in the SLA, and the eventually delivered QoS. Therefore, it is deemed a system-based metric as it is based on quantitative computations.
- *Fidelity*: is a new notion we present in this work to measure the creditability of a provider by assessing how consistently it abides by the QoS conformance.

The rest of the paper is structured as follows. In Section II, we give background knowledge pertaining to the commonly used QoS metrics in Grid computing and describe related work in reputation systems. In Section III, we formally define how to compute QoS conformance, and introduce the new *fidelity* metric along with its formulation. In Section IV, we describe a reputation-enabled framework in which these concepts fit together. The last section summarizes and concludes our paper.

II. Related Work

QoS is one of the chief aspects in the Grid [1],[2], [3], [5], [6], [7]. Buco [2] advocates that a successful utility computing provider must be able not only to satisfy its customers' demand for high service-quality standards, but also to fulfill its service-quality commitments based upon business objectives (e.g., cost-effectively minimizing the exposed business impact of service level violations). To this end, he presents the design rationale of a business-objectives based utility computing SLA management system. Menasce et al. [1][2] explore some of the relevant issues (e.g., definition of quantitative QoS metrics, and the relationship between resource allocation and SLAs) pertinent to designing grid applications that deliver appropriate QoS. Typically, when a discussion of QoS attributes is raised, the following QoS metrics are considered:

- *Latency*: it is the time elapsed between the moment of receiving the provider a request and the moment of commencing its processing. This time is also called the *waiting time*.
- *Throughput*: is the number of interactions that a service provider can process per time unit.
- *Availability*: it is the percentage of time that a service provider is up and operational through an observed period of time.
- *Response Time*: it is the time elapsed between sending a request by a client and receiving the results from the service provider.
- *Reliability*: this is a metric that represents the success rate (or probability) of processing a client request correctly within a maximum expected time frame.
- *Cost*: it is the monetary fee as set by the service provider.
- *Reputation*: it is an index of trustworthiness that is traditionally derived from the users' assessment for a particular service provider.

In general, reputation indicates how much users have confidence in a provider based on their perceptions. Therefore, the reputation system is one of the important tools that clients should consult as they shop around for a provider. Depending on the reputation algorithm, the reputation system works on aggregating user ratings and distributes (or makes) the results accessible to prospective clients [8]. E-Bay [11] and Amazon [12] are examples of e-business that essentially count on reputation systems that are based on aggregating the numerical ratings.

We believe that reputation systems are vital to Grid environments in order to increase reliability, utilization, and popularity. Reputation serves as an important metric to avert the usage of under provisioned and malicious resources with the help of community feedback; it provides the ability to simplify the selection process while focusing first on qualitative concerns (e.g., what storage or CPU power are available at a particular resource).

Reputation systems have been the focus of several researchers [9], [10], [13], [14], [15]. The work in [9] suggests a reputation management framework for Grids to facilitate a distributed and efficient mechanism for resource selection. The framework uses an algorithm that combines the two known concepts of EigenTrust [13] and global trust integration [14]. In [15] the authors present a trust model for Grid systems and show how the model can be used to incorporate the security implications into scheduling algorithms.

A quantitative comparison of reputation systems in the Grid environment is conducted in [10]. This study shows that using a reputation system to guide service selection can significantly improve client satisfaction with minimal overhead. We noticed that the majority of the studied systems rely primarily on basic statistical computation (e.g., averaging) on user rating. What we advocate in this paper is to maintain system-based time series statistics that help infer how providers are likely to behave in the future, leading to improving the quality of selection. Among these statistics are the *Conformance* and *fidelity*, as explained next.

III. REPUTATION = (USER RANKING, CONFORMANCE, FIDELITY)

As shown in the previous section, the majority of reputation systems count ultimately on ratings given by the end users. Such interpretation of trustworthiness remains merely a user perception that fails to express the real performance quality of a provider or its service. Confining reputation to user ranking creates a vulnerable metric that is sensitive to manipulative providers (e.g., *bad mouthers* and *ballot stuffer*) and to devious practices by service providers such as delivering extremely high quality of service at specific times of the year in order to manipulate the perception of the end users hoping that they forget the provider's past shortcomings. Reputation should rather be an indication of the *truthfulness* of a service provider to deliver the promised performance along with user's opinions. Therefore, we express reputation as a three dimension vector comprising *User Ranking*, *Conformance*, and *Fidelity*. User ranking gives users the opportunity to express their opinion. Conformance gauges the compliance of a provider with the agreed SLA. Fidelity reflects how much clients can trust a specific provider with respect to delivering *consistent* conformance levels.

A. Local vs. Global Reputation Profiles

For each provider, the three metrics introduced above can be computed for two reputation profiles: *local* and *global*. The local profile is internally maintained by each client to reflect its sole view and experience with certain service provider. The global profile, on the other hand, is a publicly shared profile, typically maintained by the Grid broker, that reflects the collective view and experience of all clients with that provider. Relying on the global profile solely denies the client from factoring in its personal past experience (good or bad) with a particular provider. Furthermore, the local profile can be the last resort in case the global profile is inaccessible for some reason.

Having the two reputation profiles entails having two versions for each metric. *lRank*, *lConformance*, and *lFidelity* to denote the reputation components maintained in the local profile. Similarly, *gRank*, *gConformance*, and *gFidelity* denote the reputation components maintained in the global profile.

To obtain reliable statistics, reputation components (user ranking, conformance, fidelity) should be calculated based on observing a number, n, of client-provider interactions that occurred over a time window (*assessment window*). The client and the Grid broker can control the size of this window, which can be as big as days, weeks, months, or even year. The subsequent sections explain how to compute the three reputation components formally.

B. User Ranking

When a transaction is completed, each user is given the opportunity to rate the provider with regard to the quality of the transaction. Depending on the ranking system and the aggregation algorithm [10], a rating can be, for example, either +1 for satisfied, -1 for dissatisfied, or 0 for neutral; a user's feedback score is the sum of these individual ratings. Other aggregation algorithms can produce a normalized ranking metric as follows:

$$lRank = \frac{r - s}{s + r} \tag{1}$$

Where r and s represent the number of positive and negative ratings assigned to the provider respectively over a given assessment window.

The *gRank* can be computed as follows:

$$gRank = \frac{p - v}{v + p} \tag{2}$$

Where p and v respectively represent the number of clients that voted positively and negatively through the assessment window of the last n clients interacted with the service provider.

C. Conformance

Conformance is the concept that indicates the degree by which a service provider complies by the SLA. Bhoj et al. [16] view a compliance as a contract template and a system dictionary used to verify a contract. The specific data of a client are used to fill out the contract template. Such data get verified against the system measurements that were collected while forming the contract. The contract is evaluated and the compliance results are produced as customized reports for each client. What we propose in this paper is to quantitatively express the discrepancy between what the provider promised to deliver in its SLA and the actual quality attained.

To compute the conformance, the reputation system should progressively record the actually delivered attribute values for each interaction between the client and the provider. The conformance of each quality attribute (e.g., response time) in the SLA is computed as the average of the normalized difference between the projected and actual values of the attribute. The normalized difference could be positive or negative depending on whether the agreed upon value was greater than or lesser than the delivered value, respectively. A positive average indicates a positive compliance which means the agreed values have been delivered without violations. A negative average indicates a negative compliance and that the provider failed to deliver the agreed values. An average value of zero is ideal compliance indicating the delivered values being exactly equal to the agreed values.

Formally, the conformance of a provider is computed from the conformance values of all SLA attributes. At interaction (or service invocation) number I between the client and the provider, there are two vectors of attributes values:

$$P = (p_{1i}, p_{2i}, \ldots, p_{qi})$$
$$A = (a_{1i}, a_{2i}, \ldots, a_{qi})$$

Vector P represents the projected values of all q attributes negotiated in the SLA. Vector A represents the actually delivered values of these attributes. I denotes the interaction number observed within the assessment window of the most recent n interactions ($1 \le i \le n$).

The normalized difference, c_{xi}, between any pair a_{xi} and p_{xi} ($1 \le x \le q$) can be positive or negative and it denotes the compliance of the provider after executing interaction I with respect to attribute x:

$$c_{xi} = \frac{a_{xi} - p_{xi}}{p_{xi}} \quad (3)$$

This leads to producing the conformance vector ($c_{1i}, c_{2i}, \ldots, c_{qi}$). However, since all dimensions are normalized and we are interested in obtaining the scalar value, C_i, that reflects the collective conformance of the provider upon the completion of interaction i, we can define

$$C_i = \frac{1}{q} . \sum_{x=1}^{q} Cxi \quad (4)$$

To compute *lConformance* over an assessment window of n interactions, we can write:

$$lConformance = \frac{1}{n} . \sum_{i=1}^{n} Ci \quad (5)$$

To compute *gConformance*, the broker views the assessment window as the last m clients that interacted with a certain provider. Therefore, we can write:

$$gConformance = \frac{1}{m} . \sum_{i=1}^{m} Ci \quad (6)$$

where C_i is the conformance value reported by client i resulting from its last interaction with the provider.

D. Fidelity

Conformance is useful to assess the overall compliance of a provider; however, it does not show the degree of adherence to that compliance. The role of the *fidelity* metric proposed in this paper is to assess the degree of consistency in the compliance levels of a service provider, adding a new dimension to the assessment of the reputation of providers.

The notion of fidelity is based on computing the variance in the conformance levels. The lower the variance is, the more successful the provider is in delivering consistent performance levels. Fidelity is computed over a range of conformance measures obtained from past interactions with the service provider. Doing so gives an insight into providers' historical performance by progressively assessing their conformance levels over a range of past interactions (assessment window). In other words, fidelity is deemed a measure of truthfulness and verity.

Since fidelity represents the degree of variation in the conformance values achieved by a service provider, we can use the statistical standard deviation (σ) to define fidelity. The standard deviation can assess the degree of spread in the conformance values of a service provider, which indicates how efficiently and how often the guaranteed levels of quality are met. However, and since we are in need for a normalized metric that enables ranking providers among themselves objectively, we further normalize the standard deviation by the mean (μ) producing what is statistically known as the *coefficient of variance* (COV), which is basically σ / μ.

Hence, the local fidelity can be formulated as:

$$lFidelity = \frac{\sigma}{\mu} = \frac{\sqrt{\frac{1}{n}\sum_{i=1}^{n}(C_i - \mu)^2}}{\mu} \quad (7)$$

Where n is the number of the conformance measures spawn from the last n interactions (i.e., size of the assessment window as set by the client). C_i is the conformance level measured upon the completion of the i^{th} interaction. μ is the mean conformance, that is

$$\mu = \frac{\sum_{i=1}^{n} C_i}{n} \quad (8)$$

Similarly, the global fidelity can be computed as follows

$$gFidelity = \frac{\sqrt{\frac{1}{m}\sum_{t=1}^{m}(C_t - \psi)^2}}{\psi} \quad (9)$$

Where m is the number of the conformance measures collected from the last m clients that interacted with the provider. Therefore, m represents the assessment window size as set by the Grid broker. C_t is the conformance obtained from the t^{th} client upon the completion of its interaction. ψ is the mean conformance, that is

$$\psi = \frac{\sum_{t=1}^{m} C_t}{m} \quad (10)$$

E. Reputation (Rep) of Providers

Since we define reputation as a combination of user ranking, conformance, and fidelity, the local reputation profile of a provider can be represented by the vector *lRep*=(*lRank*, *lConformance*, *lFidelity*) that is maintained at the client side. Likewise, the global reputation profile of that provider can be represented by the vector *gRep*=(*gRank*, *gConformance*, *gFidelity*) that is maintained at the broker side. The overall reputation (*Rep*) of the provider as seen by the client can be expressed as the weighted sum of the local and global reputation vectors:

$$Rep = W_l * lRep + W_g * gRep \quad (11)$$

Such that

$$W_l + W_g = 1$$

Where W_l and W_g represent the significance (or trust) weights of the local reputation and the global reputation, respectively, as arbitrarily assigned by each client. Therefore, we can write

$$Rep = (Rank, Conformance, Fidelity) \quad (12)$$

where

$Rank = W_l * lRank + W_g * gRank$
$Conformance = W_l * lConformance + W_g * gConformance$
$Fidelity = W_l * lFidelity + W_g * gFidelity$

IV. A Framework for Reputation-Based Grids

In light of the new definition of reputation, Fig 1 depicts the reputation-based framework that we suggest for the Grid. As we can see, we have the three typical main players in this framework: clients, Grid brokers, and service providers. A provider publishes its services by registering them, along with their associated QoS specifications, with a broker. The

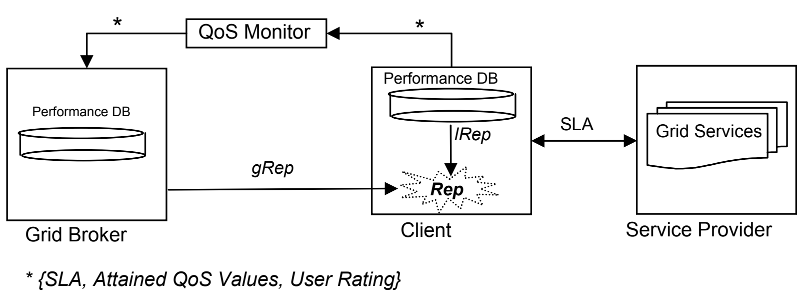

Fig 1. Reputation-Based Framework for the Grid

broker maintains a database of historical performance measures (e.g., user ratings and conformances) for each service provider. The broker uses this database to compute the global reputation (*gRep*) of a provider, as explained in the previous section. Likewise, each client maintains a database that stores a time series of performance measures experienced with providers that client dealt with. This internal database is referenced when the client needs to compute the local reputation (*lRep*) of a given provider.

Upon a request from a client, the broker returns a list of candidate providers. The client, based on the *Rep* vector, selects its prospective provider and negotiates the relevant QoS attributes with it. A copy of the finalized SLA is sent to the broker in order to compute reputation components (e.g., *gConformance*).

The execution of client-provider interactions gets monitored by an independent module *QoS Monitor* [17] in order to intercept the eventually delivered QoS attribute values and relay them to the performance database of the broker. It is important to have such independent monitoring entity in order to prevent clients from tampering with real performance readings attainted by service providers. In addition to the finally attained performance data, each client sends its own personal rating to the broker in order to compute the global ranking (*gRank*), which is required to complete the computation of the global reputation (*gRep*) of a provider.

V. Conclusion

Finding a service provider that can satisfy the client's computing requirements, such as CPU power or storage capacity, is the first step towards forming the initial set of candidate providers. The client has yet to rank these candidates according to their past compliance with the agreed SLA. This compliance and providing consistent conformance with the promised QoS attributes constitute provider's reputation. We contended in this paper that counting exclusively on the typical ratings given by users is not a reliable means of assessing the goodness of a provider due to numerous pitfalls. Rather, we view reputation as a function of three combined qualities: user ranking, conformance, and fidelity. We explained how to produce local and global reputation vectors and described the general framework in which the new concepts fit together in order to establish a reputation-based computing platform.

The next item in our research agenda is to quantitatively compare our reputation system with others. We intend to build a generic Grid simulation framework in which we can seamlessly plug into it any arbitrary reputation policy that we intend to evaluate.

Acknowledgment

We are grateful to the research support we got from the Scientific Research Council of the UAE University through research grant #02-06-9-11/06.

References

[1] D. Menascé and E. Casalicchio, Quality of Service Aspects and Metrics in Grid Computing, *Proc. 2004 Computer Measurement Group Conference*, Las Vegas, NV, December 5-10, 2004.

[2] D. Menascé & E. Casalicchio, QoS in Grid Computing, *IEEE Internet Computing 8*(4), July/August 2004.

[3] M. J. Buco, R. N. Chang, L. Z. Luan, C. Ward, J. L. Wolf & P. S. Yu, Utility Computing SLA Management Based Upon Business Objectives, *IBM Systems Journal 43*(1), 2004, pp. 159-178.

[4] A. Bouch, A. Kuchinsky & N. Bhatti, Quality is in the Eye of the Beholder: Meeting Users' Requirements for Internet Quality of Service. *Proc. SIGCHI on Human Factors in Computing Systems*, April 2000.

[5] I. Foster & C. Kesselman, *The Grid: Blueprint for a New Computing Infrastructure* (2nd ed., Morgan Kaufmann, 2004).

[6] A. Leff, J. Rayfield & D. Dias, Service-Level Agreements and Commercial Grids, *IEEE Internet Computing 7*(4), 2003, pp. 44-50.

[7] J. Nabrzyski, J. Schopf & J. Weglarz, eds., *Grid Resource Management - State of the Art and Future Trends*, (Kluwer Academic Publishers, 2004).

[8] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara, Reputation Systems, *Communications of the ACM 43*(12), December 2000.

[9] K. Alunkal, I. Veljkovic, G. von Laszewski & K. Amin, Reputation-Based Grid Resource Selection, *Workshop on Adaptive Grid Middleware (AGridM 2003)*, New Orleans, LA, USA, Sept. 28, 2003.

[10] J. Sonnek & J. Weissman, A Quantitative Comparison of Reputation Systems in the Grid, *Proceedings of the Sixth ACM/IEEE International Workshop on Grid Computing*, 2005.

[11] Ebay Web Page [online]. Available: http://www.ebay.com.

[12] Amazon Web Page [online]. Available: www.amazon.com.

[13] S. Kamvar, M. Schlosser & H. Garcia-Molina, The Eigentrust Algorithm for Reputation Management in P2P Networks, *12th International World Wide Web Conference 2003*, Budapest, Hungary, ACM Press, May 2003.

[14] F. Azzedin & M. Maheswaran, Evolving and Managing Trust in Grid Computing Systems, *IEEE Canadian Conference on Electrical Computer Engineering*, *IEEE Computer Society Press*, May 2002.

[15] F. Azzedin and M. Maheswaran, Integrating Trust into Grid Resource Management Systems, *Proceedings of the International Conference on Parallel Processing*, pages 47-54, 2002.

[16] P. Bhoj, S. Singhal & S. Chutani, SLA Management in Federated Environments, *HP Labs Technical Report*, December, 1998.

[17] A. Keller & H. Ludwig, Defining and Monitoring Service Level Agreements for Dynamic e-Business, *Proceedings of the 16th USENIX System Administration Conference (LISA '02)*, November 2002.