Flexible QoS-Based Service Selection and Provisioning in Large-Scale Grids

Sebastian Stein, Terry R. Payne, Nicholas R. Jennings
School of Electronics and Computer Science
University of Southampton, Southampton, SO17 1BJ, UK
{ss2,trp,nrj}@ecs.soton.ac.uk

Grid computing is a key enabling technology for the field of e-Science. As scientists rely increasingly on computational support for their work, Grids allow research organisations to pool their resources and spread the cost of expensive hardware. Often, such Grids span across organisational boundaries, thereby giving users access to a vast range of services — from running complex simulations to analysing large data sets [1].

However, as Grids become larger, more open and distributed, new challenges need to be addressed [2]. In particular, the behaviour of services may be highly uncertain, as manifested by frequent failures and execution delays. This uncertainty may be due to network delays, power cuts, competition for resources or machine failures. Furthermore, services will typically be owned by agents that act autonomously and cannot be assumed to behave cooperatively or deterministically [3]. Rather, they follow their own decision-making procedures and may even defect when this increases their private utility.

Clearly, this uncertainty must be addressed, and this is particularly important when consumers execute large workflows of interdependent tasks, when service providers demand remuneration and when workflows have deadlines. To this end, we concentrate on the problem of service provisioning (i.e., deciding which service instances to select for the constituent tasks of an abstract workflow). This allows the consumer to consider non-functional quality-of-service (QoS) parameters (such as the cost or reliability) of functionally equivalent services and to react to failures by dynamic re-provisioning.

Current work does not address the provisioning problem satisfactorily. Typically, QoS parameters are only considered locally for each task, e.g., to select the cheapest or most reliable instance [4]. Other work optimises a weighted sum of global QoS parameters [5]. However, such approaches have several shortcomings. First, they do not reason about failures and their impact on workflow completion. Second, they provision only one service for each task, thus producing brittle workflows that are vulnerable to single failures. Third, they usually rely on either on-demand provisioning or advance provisioning, when in fact a mixture might be appropriate.

To address these shortcomings, we propose a principled decision-theoretic approach to service provisioning. Within this framework, we attach a time-dependent utility to the successful completion of a workflow and then provision services that maximise the consumer's expected utility, as given by:

$$\mathrm{E}(U(p)) = \mathrm{E}(R(p)) - \mathrm{E}(C(p)), \quad (1)$$

where $\mathrm{E}(U(p))$ is the expected utility of the agent when following the provisioning allocation p, $\mathrm{E}(R(p))$ is the respective expected reward and $\mathrm{E}(C(p))$ is the expected cost.

Adopting this formalism allows the consumer to reason explicitly about the value of choosing particular service instances and balances the overall workflow cost with its expected reward. Furthermore, our formulation of p contains two key features that are missing from existing approaches:

- **Redundancy:** For particularly critical workflow tasks, our approach considers the parallel provisioning of several services. This increases the probability of success and typically reduces the task duration.
- **Contingency Plans:** Our approach builds contingency plans to deal with failures, e.g., by deciding to re-provision an equivalent service after a dynamically chosen time-out or by provisioning another more reliable (possibly more expensive) instance.

To illustrate our approach, Figure 1 shows a fully provisioned workflow, where the service consumer has introduced some redundancy (n denotes the number of service instances) and where it reacts to failures by re-provisioning tasks after a chosen time-out (w denotes the time-out).

In the full version of this paper, we provide a detailed overview of our approach (based on [6]) and present three variants that consider different Grid environments. In particular, we first examine Grid systems where the consumer has no specific knowledge about the performance of individual service providers (e.g., in homogeneous clusters or in highly dynamic peer-to-peer systems). In this case, we show that it can still effectively influence its utility by varying the redundancy within workflows. Next, we consider systems where the consumer may have some knowledge to differentiate between providers. Here, we show how the consumer can decide which services to rely on and whether to provision several providers redundantly. Finally, we examine systems where services are offered on a highly dynamic market with changing availability and prices, and where services may be provisioned in advance (by obtaining specific service-level agreements).

To evaluate our flexible approach, we have benchmarked it against the state of the art in service provisioning (both the local and global approaches discussed earlier). Our results are summarised in Figure 2, highlighting that our approach can lead to a significantly higher utility and a larger number of successful workflows than existing strategies when there is some uncertainty in the service providers' behaviour.

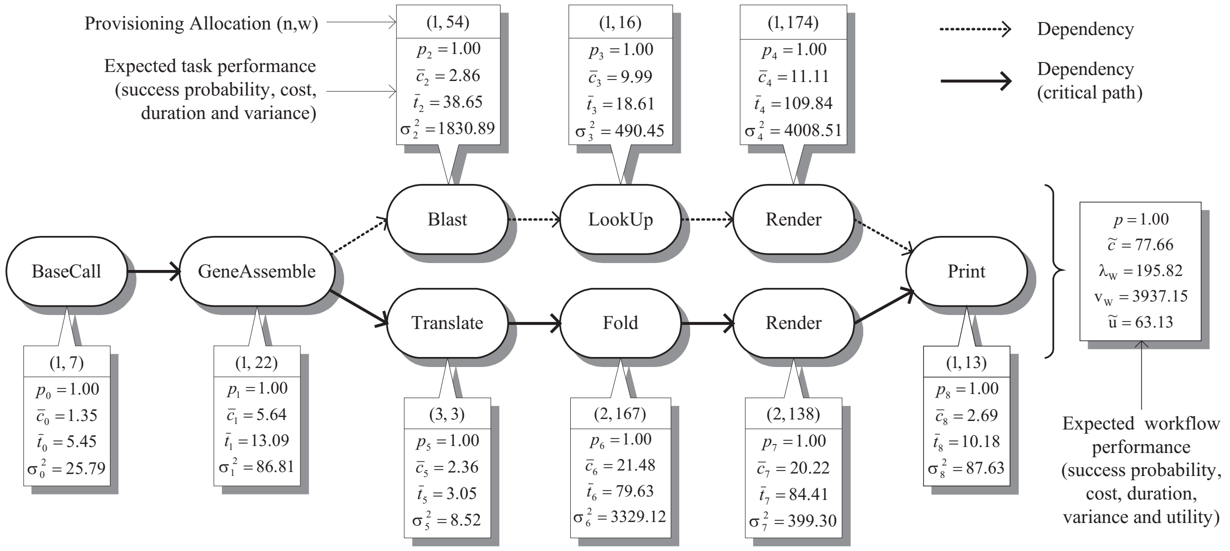

Fig. 1. An example bioinformatics workflow that has been provisioned using our approach. Here, a number of parallel providers have been provisioned in parallel for each task (n), and the consumer has determined a suitable time-out for re-provisioning each task (w). Based on these, local performance parameters are given for each task and these are aggregated to yield global workflow parameters (using a critical-path technique). In this example, the consumer attaches a value of £150 to the workflow and attempts to finish within a deadline of 240 minutes.

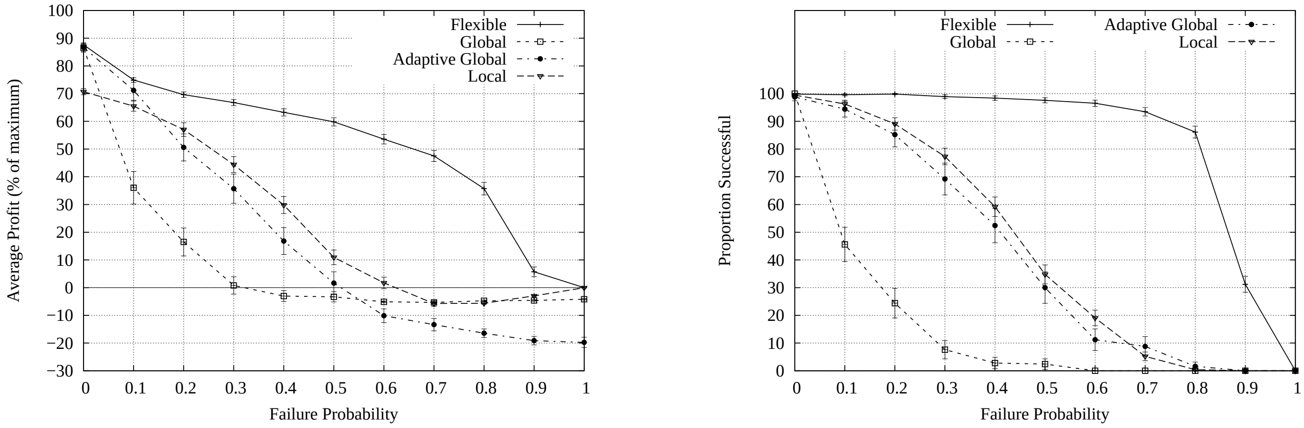

(a) Average Profit

(b) Proportion of Successful Workflows

Fig. 2. Performance of the considered strategies when provisioning random workflows in environments where providers are increasingly unreliable. Here, the *local* strategy provisions services based purely on local task decisions, the *global* strategy provisions the entire workflow, the *adaptive global* strategy behaves as the global, but re-provisions in case of failures, and the *flexible* strategy uses a decision-theoretic provisioning approach.

REFERENCES

[1] I. Foster, C. Kesselman, J. Nick, and S. Tuecke, "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration," Open Grid Service Infrastructure WG, Tech. Rep., 2002.

[2] I. Foster, N. R. Jennings, and C. Kesselman, "Brain meets brawn: Why Grid and agents need each other," in *Proceedings of the 3rd International Conference on Autonomous Agents and Multi-Agent Systems*, 2004, pp. 8–15.

[3] N. R. Jennings, "An agent-based approach for building complex software systems," *Communications of the ACM*, vol. 44, no. 4, pp. 35–41, April 2001.

[4] S. Vazhkudai, S. Tuecke, and I. Foster, "Replica selection in the globus data grid," in *Proceedings of the 1st IEEE/ACM International Symposium on Cluster Computing and the Grid*, 2001, pp. 106–113.

[5] L. Zeng, B. Benatallah, A. H. Ngu, M. Dumas, J. Kalagnanam, and H. Chang, "QoS-aware middleware for web services composition," *IEEE Trans. Softw. Eng.*, vol. 30, no. 5, pp. 311–327, 2004.

[6] S. Stein, "Flexible service provisioning in multi-agent systems," Ph.D. dissertation, School of Electronics and Computer Science, University of Southampton, 2008.