

Robust Execution of Service Workflows using Redundancy and Advance Reservations (Online Appendix)

Sebastian Stein, Terry R. Payne and Nicholas R. Jennings

1 INTRODUCTION

THIS online appendix accompanies the paper “Robust Execution of Service Workflows using Redundancy and Advance Reservations” [1]. It illustrates how the algorithm proposed in that paper is applied to a real workflow from the bioinformatics domain. In doing so, it highlights the applicability of the algorithm to a real application, and it complements the high-level description in the original paper by giving a concrete example.

The appendix is organised as follows. In Section 2, we briefly describe the bioinformatics workflow that serves as the example throughout the appendix, and in Section 3, we detail the performance characteristics of the available services. In Section 4, we provide some examples of the high-level strategies that are available to the consumer agent. Then, in Section 5, we detail the process of reserving offers for the workflow when the consumer associates a low utility and long deadline with it, and in Section 6, we discuss how this changes when the workflow is more urgent. We conclude in Section 7.

2 BIOINFORMATICS WORKFLOW

Throughout this appendix, we consider a typical workflow from the bioinformatics domain¹ — an area that relies heavily on computationally intensive services and that has increasingly seen the establishment of large distributed Grid systems for sharing resources, as exemplified by the GriPhyN [3], myGrid [4] and CombeChem [5] projects. Bioinformatics, and, more generally, data processing applications in e-science, are also prime examples of a domain where execution uncertainty and failures are rife. More specifically, computational tasks




Fig. 1. Example bioinformatics workflow, based on workflows described by [7], [8] and [9].

in this domain often take a long time to complete, sometimes running for hours or days. This directly increases the probability that hardware failures or software problems disrupt the workflow, especially when this contains many such tasks. As a concrete example of this uncertainty, [6] describes a real scientific workflow, which routinely encounters multiple task failures during a single run.

Now, although providers may be able to mitigate the uncertainty by using replication, as we argued in Section 2, this requires a considerable effort on the provider’s side, which may simply not be economical. As an example, it may often be more profitable for a provider to also sell processing time on a newly acquired mainframe (and thereby potentially doubling its revenue), rather than to use it for replication, in order to increase reliability by a few percent. Clearly, some providers *may* choose to increase reliability in this way and charge a premium for this, but our approach will then make the appropriate choice for the consumer, and, critically, it also addresses systems where all providers are unreliable.

Against this background, for our example we assume that a scientist has just sequenced a previously unknown gene of a bacterium, and is now interested in visualising the shape of the associated protein. For this, she has to carry out a number of tasks, which are shown in Figure 1.

Her initial data comprises a large set of overlapping DNA fragments in the form of chromatograms, as is common in shotgun DNA sequencing [10]. These show characteristic light traces at different wavelengths, corresponding to the four bases found in a DNA sequence. As these traces typically contain some noise and errors,

- Stein and Jennings are with the School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK. E-mail: {ss2,nrj}@ecs.soton.ac.uk.
- Payne is with the School of Electrical Engineering, Electronics and Computer Science, University of Liverpool, Liverpool, L69 3BX, UK. E-mail: T.R.Payne@liverpool.ac.uk.

1. This workflow was first described in [2].

the scientist first needs to run a base-calling service (*BaseCall*). This translates the chromatograms to the corresponding base sequences, attaching a quality value to each base in the process that denotes how accurate the assignment of the base is. The resulting base sequences are then assembled to a single continuous DNA sequence by identifying and merging overlapping fragments, using the quality values to find and repair errors. This task is performed by a sequence-assembling service, which also identifies and isolates the coding region of the gene (*GeneAssemble*).

When the coding region of the gene has been assembled, it is then translated to the corresponding amino acid sequence using a simple translation service (*Translate*). As the primary structure of the protein, this forms the input to the computationally-intensive folding service (*Fold*), which predicts the 3-dimensional shape of the protein based on a search for the conformation with the lowest free energy. The output of this — a file containing the tertiary structural data — is then rendered in high resolution using an appropriate graphics service (*Render*). In parallel with the folding simulation, the scientist is also interested in comparing the new gene to previously discovered sequences. To this end, she searches through public collections of known proteins to find the closest match using a specialised service (*Blast*), and then accesses commercial database services to retrieve structural information about the protein (*Lookup*). This is rendered again, and both images are printed as part of a report on a local printer (*Print*).

In the following, we describe how offers for these services are generated at run-time.

3 SERVICE MARKET

For the purpose of illustrating our algorithm, we detail the operation of a dynamic market for the services described in the previous section in a similar manner as done in Section 4 of [1]. In more detail, we keep a list of currently available offers associated with each time step (representing a minute in real time), from the current step t to $t+60$. Hence, the consumer may reserve services up to one hour in advance. During the simulation, at the beginning of each time step, we first generate new offers that become available in the market by drawing the number of new offers and their parameters from random distributions. These distributions are detailed in Table 1 ($\mathcal{U}_c(a, b)$ is a continuous uniform distribution over the interval $[a, b]$) and depend on the number of time steps the offer is generated in advance. This time dependency allows us to include performance differences in offers when they are reserved with varying advance notice periods. It is expressed here by including two rows of distributions for each service type — the first indicates the performance of offers when reserved at short notice (as given by the *advance time* column), and the second gives the performance when offers are reserved with a long advance notice period.

To elaborate this, for each service type and for each possible time step from t to $t + 60$, we first generate the number of new offers by drawing a sample from a Poisson distribution with a mean given by the respective *birth* rate². Then, for each such generated offer, we draw its failure probability, reservation cost, execution cost and service duration from the relevant distributions (depending on how far the offer is generated in advance). When the offer time lies between the two extremes corresponding to the two rows for each service type, we interpolate linearly between the distribution parameters. For example, when generating an offer for the *BaseCall* service type for time step $t + 2$, we draw its failure probability from $\mathcal{U}_c(0.18, 0.44)$. If the service type is marked as repaying in the table, we set the failure penalty to the combined execution and reservation costs (this means that providers for these services always refund a consumer in case of failures). Otherwise, it is set to 0. Finally, at the end of each time step, we remove offers in a similar way as above by drawing a random sample from a Poisson distribution with its mean given by the *death* rate. This models the demand for such services and we randomly remove the generated number of offers from that time step (or all offers if the number exceeds the current supply).

The service parameters in Table 1 were chosen to represent a range of heterogeneous services with some, such as *Fold*, being expensive and time-consuming, while others, such as *BaseCall*, are fast and cheap. Also, the services display significant performance differences depending on the time of reserving them. For example, some services are generally more reliable and faster when reserved in advance (*BaseCall* and *GeneAssemble*), some offer a far better service overall but are also significantly more expensive and need a long advance notice period (*Fold*), and a few display no difference in quality over time (*Look Up* and *Print*).

Given this service market, we describe a number of example high-level task strategies in the following section.

4 HIGH-LEVEL TASK STRATEGIES

As described in Section 3.2 of the main paper, we assume that the service consumer has already obtained a library of atomic high-level strategies by observing the market over some time. To illustrate this, Figure 2 shows a number of example strategies and their statistics for the *Fold* service type. In the centre, we use the notation $\omega = \langle 1, 20, 10, b \rangle$ to represent a strategy ω that reserves a single offer ($n(\omega) = 1$) 20 time steps before it is required ($t_a(\omega) = 20$), with the consumer considering 10 consecutive time steps ($t_w(\omega) = 10$) and selecting the offer that best balances all performance characteristics

². As in the main paper [1], we chose the Poisson distribution here because it is a common distribution for modelling random arrival events [11].

Service Type	Advance Time (minutes)	Failure Probability	Reservation Cost (\$)	Execution Cost (\$)	Time (minutes)	Birth/Death	Repay?
BaseCall	≤ 0 ≥ 10	$\mathcal{U}_c(0.2, 0.5)$ $\mathcal{U}_c(0.1, 0.2)$	$\mathcal{U}_c(0.5, 1.0)$ $\mathcal{U}_c(0.5, 1.0)$	$\mathcal{U}_c(0, 0)$ $\mathcal{U}_c(0, 0)$	$\mathcal{U}_c(6, 10)$ $\mathcal{U}_c(1, 2)$	1/3 0.1/0	no
GeneAsemble	≤ 5 ≥ 10	$\mathcal{U}_c(0.1, 0.2)$ $\mathcal{U}_c(0, 0.1)$	$\mathcal{U}_c(2, 5)$ $\mathcal{U}_c(1, 2)$	$\mathcal{U}_c(2, 5)$ $\mathcal{U}_c(0, 0)$	$\mathcal{U}_c(10, 30)$ $\mathcal{U}_c(5, 10)$	1/0.5 0.5/1	yes
Blast	≤ 5 ≥ 15	$\mathcal{U}_c(0.5, 1)$ $\mathcal{U}_c(0, 0.1)$	$\mathcal{U}_c(2, 3)$ $\mathcal{U}_c(1, 5)$	$\mathcal{U}_c(1, 5)$ $\mathcal{U}_c(2, 3)$	$\mathcal{U}_c(20, 40)$ $\mathcal{U}_c(5, 10)$	1/2 0.5/0.5	no
LookUp	≤ 0 ≥ 1	$\mathcal{U}_c(0.5, 0.7)$ $\mathcal{U}_c(0.5, 0.7)$	$\mathcal{U}_c(0, 0)$ $\mathcal{U}_c(0, 0)$	$\mathcal{U}_c(4, 10)$ $\mathcal{U}_c(4, 10)$	$\mathcal{U}_c(2, 8)$ $\mathcal{U}_c(2, 8)$	0.5/0.25 0.5/0.25	yes
Render	≤ 15 ≥ 30	$\mathcal{U}_c(0.2, 0.3)$ $\mathcal{U}_c(0, 0)$	$\mathcal{U}_c(5, 10)$ $\mathcal{U}_c(5, 10)$	$\mathcal{U}_c(10, 15)$ $\mathcal{U}_c(5, 10)$	$\mathcal{U}_c(150, 240)$ $\mathcal{U}_c(80, 120)$	0.5/1 1/1	no
Translate	≤ 10 ≥ 30	$\mathcal{U}_c(0.7, 1.0)$ $\mathcal{U}_c(0.3, 0.4)$	$\mathcal{U}_c(2, 5)$ $\mathcal{U}_c(0.1, 0.5)$	$\mathcal{U}_c(1, 2)$ $\mathcal{U}_c(0.25, 0.5)$	$\mathcal{U}_c(10, 40)$ $\mathcal{U}_c(5, 10)$	0.5/1 0.5/0.5	no
Fold	≤ 15 ≥ 45	$\mathcal{U}_c(0.25, 0.75)$ $\mathcal{U}_c(0, 0.05)$	$\mathcal{U}_c(5, 20)$ $\mathcal{U}_c(10, 20)$	$\mathcal{U}_c(5, 20)$ $\mathcal{U}_c(20, 30)$	$\mathcal{U}_c(80, 400)$ $\mathcal{U}_c(20, 30)$	3/5 1/1	yes
Print	≤ 0 ≥ 1	$\mathcal{U}_c(0.4, 0.8)$ $\mathcal{U}_c(0.4, 0.8)$	$\mathcal{U}_c(1, 2)$ $\mathcal{U}_c(1, 2)$	$\mathcal{U}_c(0, 0)$ $\mathcal{U}_c(0, 0)$	$\mathcal{U}_c(8, 12)$ $\mathcal{U}_c(8, 12)$	2/2 2/2	yes

TABLE 1
Distributions used to generate random offers for bioinformatics services.

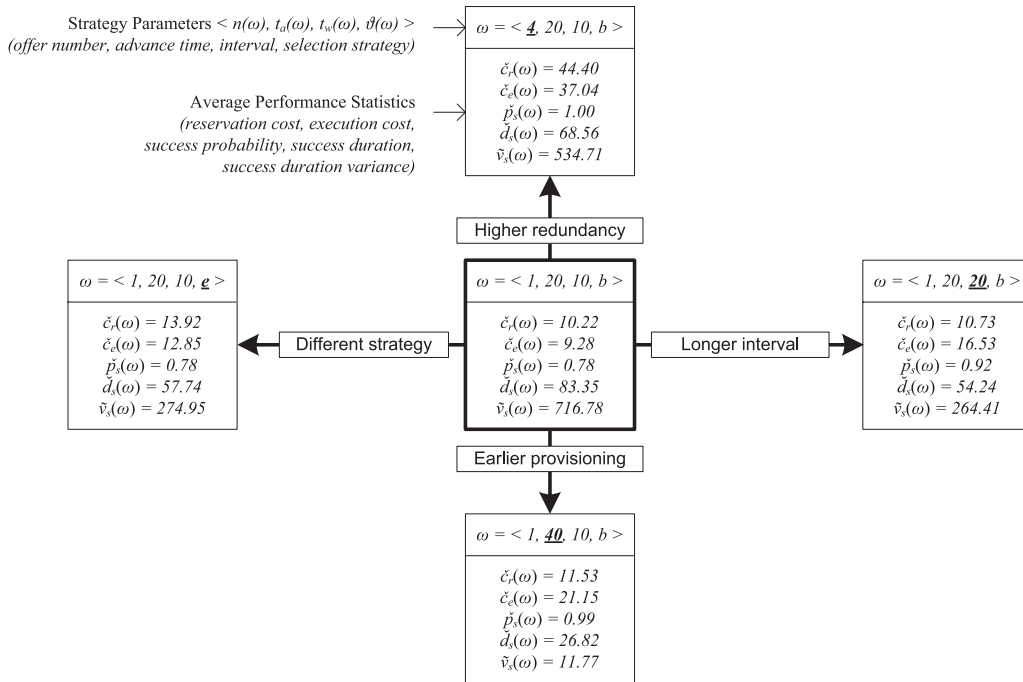


Fig. 2. Example high-level strategies for the *Fold* task.

($\vartheta(\omega) = \text{balanced}$)³. We note that this strategy is relatively cheap, has a success probability of only 78% and takes a long and highly uncertain time to complete. The remaining strategies shown in Figure 2 demonstrate the impact of slightly altering this reservation strategy — for example, when increasing the advance notice period from 20 to 40 time steps, the cost rises, but the success probability also increases to 99%, while the duration and its variance drop significantly (these trends all emerge from the distributions given in Table 1).

Now, these high-level strategies form the basic decisions available to our algorithm. In the following, we

3. Here, and in the remainder of the section, we abbreviate each selection strategy in $\{\text{cost, unreliability, end_time, balanced}\}$ with its first letter.

demonstrate how they are used in the context of the entire workflow.

5 NON-URGENT SCENARIO

First, we consider a scenario where the workflow has a relatively low value and a long deadline. Hence, it may represent a routine task for a scientist, which is not particularly critical to her work. More specifically, we assume the workflow has a maximum utility $u_{\max} = 150$, deadline $t_{\max} = 240$ (4 hours) and penalty $\delta = 1$ (i.e., no more utility is gained after 6.5 hours).

Figure 3 now shows the initial high-level decisions that our agent takes for the workflow. Here, the agent generally attempts to spend as little on services as possible, preferring to wait longer for completion. Thus, the

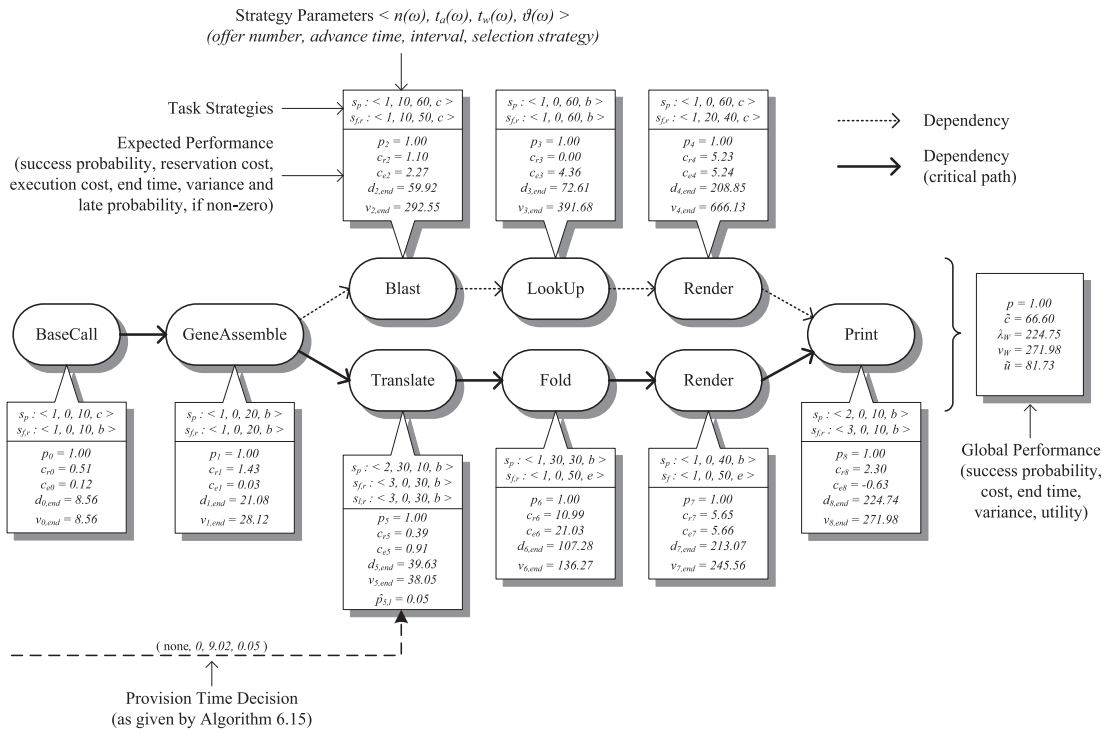


Fig. 3. Workflow with initial high-level decisions.

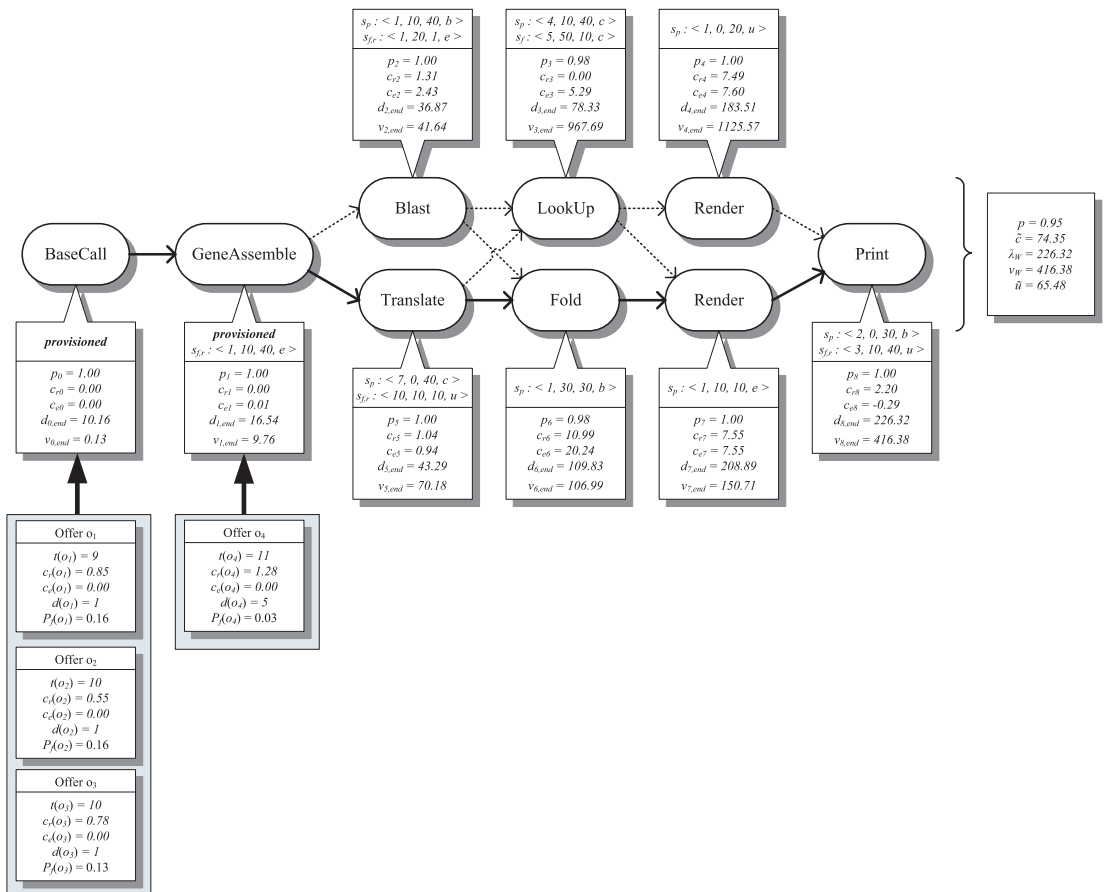


Fig. 4. Workflow after reserving initial offers.

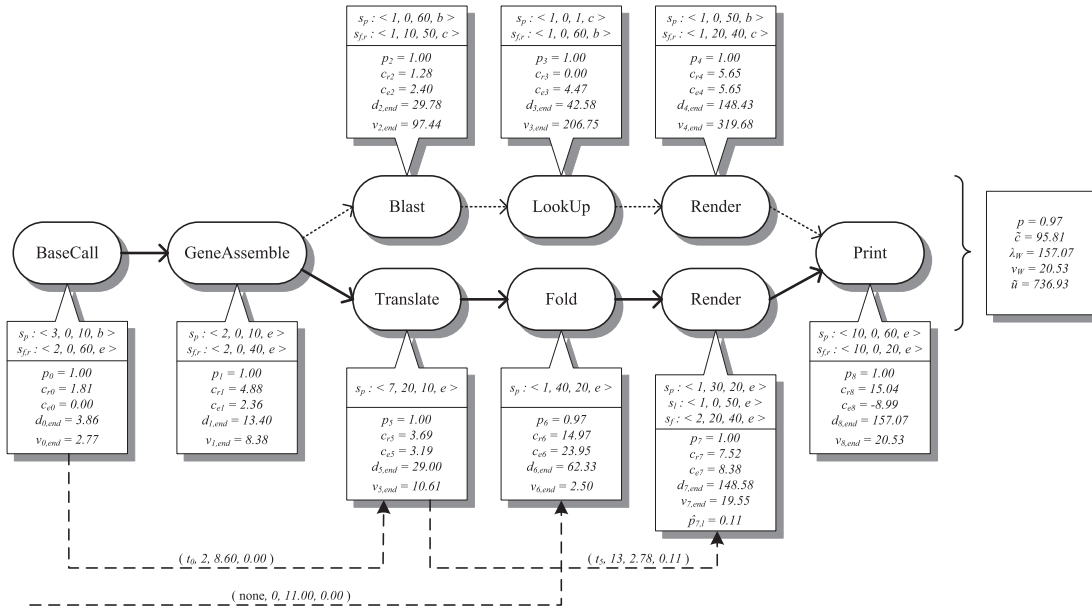


Fig. 5. Urgent workflow with initial high-level decisions.

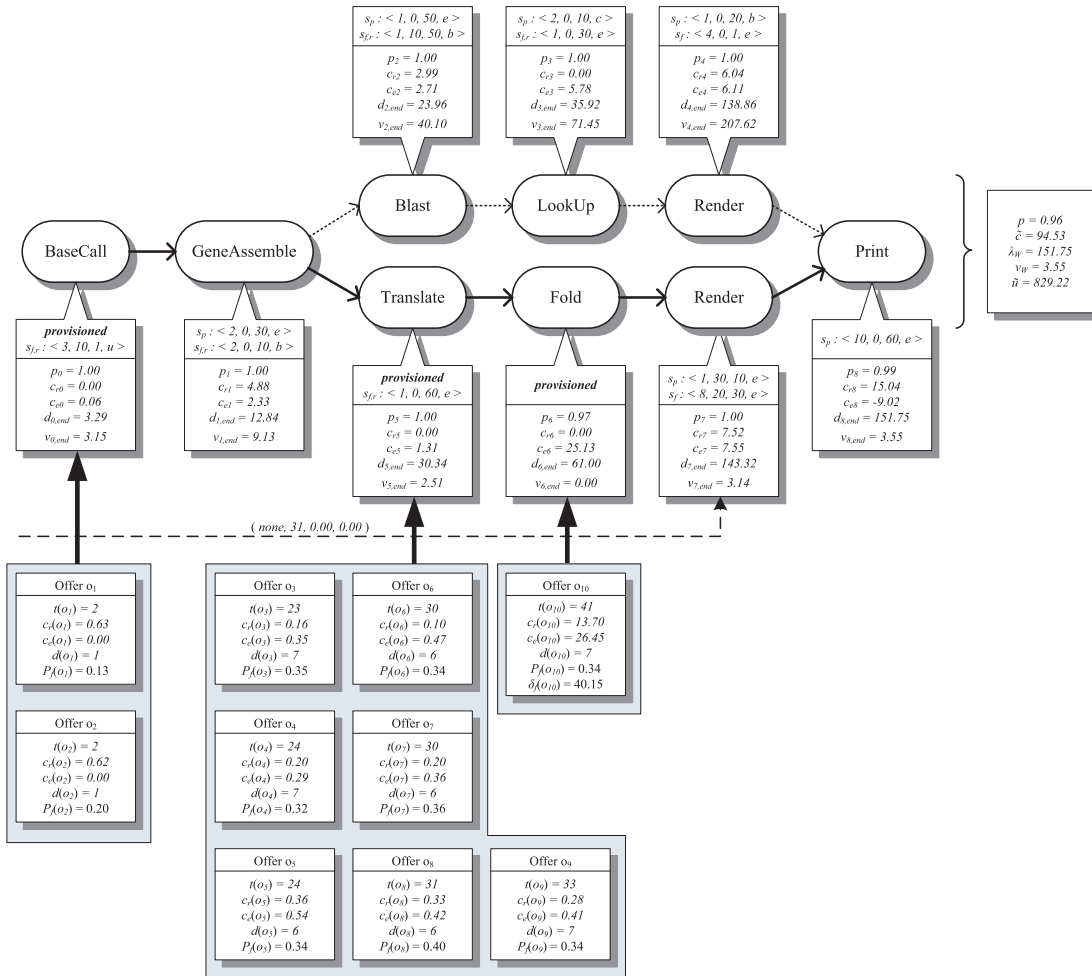


Fig. 6. Urgent workflow reserving initial offers.

algorithm decides to reserve only single offers for most tasks and relies on cautious contingency plans, where more single offers are reserved gradually (and repeatedly) in case of failure. The only parallel redundancy in the plan is used for tasks t_5 (*Translate*) and t_8 (*Print*), which are relatively cheap. Due to the longer deadline in this case, the consumer also decides to include few task overlaps in the workflow and instead prefers to leave the reservation of each task until all predecessors are complete. The only exception to this is t_5 (*Translate*), which the consumer chooses to reserve while its predecessors are still executing. Using Algorithm 1 in the main paper, the strategy here determines that the task should be reserved immediately when the workflow is started (for time step 30) and that this will result in a 5% probability of losing the reserved offers later on and an additional delay of 9.02 time steps. These figures are based on the uncertain duration of its predecessors (*BaseCall* and *GeneAssemble*), which are expected to complete by $t = 21.08$. Overall, the consumer expects the workflow to finish just before the deadline, after 224.75 time steps (but with considerable variance) and expects to spend \$66.60, thus achieving an expected utility of $\tilde{u} = 81.73$.

Next, Figure 4 shows the same workflow after the first offers have been reserved (during the initial time step). Here, the agent has consulted the market and followed its high-level strategies in reserving offers for some of the workflow tasks. In particular, the consumer has now reserved three offers for t_0 (*BaseCall*). In this case, it is different from the initial decision of reserving a single offer, as the agent immediately revises and improves its decisions as it observes the actual offers available on the market. Given these three offers, the task parameters are updated to reflect their terms (hence, the task end time is now almost certain). As is evident in the remainder of the workflow, the consumer has also now adapted its high-level strategies based on the new information. In particular, knowing that t_0 is almost certain to complete by time step $t = 11$, it has decided to reserve offers for task t_1 (*GeneAssemble*) earlier than originally planned. On the other hand, it also delayed reserving offers for task t_5 (*Translate*) to a later time. Finally, the consumer has introduced a number of additional edges into the workflow. Some of these have no impact on the estimated workflow utility, but the additional edge between t_3 (*LookUp*) and t_7 (*Render*) ensures that the former, slightly uncertain task completes before the expensive *Render* task is started.

In the next section, we contrast these decisions with a case where the workflow has a higher priority.

6 URGENT SCENARIO

In order to investigate how the agent's reservation decisions change as the workflow becomes more urgent and valuable, we now consider a second scenario. For this, we assume a maximum utility $u_{\max} = 1000$, a deadline

$t_{\max} = 150$ and penalty $\delta = 20$. This could occur in the case where the results are critical to the scientist and when they are needed within a short amount of time (for example, for an impending project meeting or a publication deadline). Figure 5 illustrates the initial high-level reservation decisions, clearly highlighting an increasing reliance on redundant services and also on advance reservations, which here allows the agent to obtain better services and decrease the overall execution time (as services are reserved before their predecessors are completed). In more detail, the agent here decides to reserve some tasks immediately (such as *Fold*), but leaves the reservation of others until later (such as the lower *Render* task), according to the advance notice periods required and the expected completion times of their predecessors.

Next, Figure 6 shows the workflow after the first offers have been reserved. Here, the strategy has mostly followed its initial decisions. However, based on the offers reserved for task t_0 (*BaseCall*), it also immediately reserves offers for t_5 (*Translate*), rather than waiting an additional two time steps. As the offers were generally as expected, most remaining high-level decisions are unchanged and the overall expected utility has risen slightly, due to an earlier estimated completion time.

7 CONCLUSIONS

In this online appendix to [1], we have applied our algorithm to a real bioinformatics workflow. This work demonstrates how our approach dynamically decides when to reserve services in advance and how its reservation decisions are driven by the overall utility function, implicitly balancing the criticality and uncertainty of completing workflow tasks with the additional cost of introducing redundancy or selecting better service offers. Thus, our approach allows service consumers to execute complex workflows in a dependable manner, even when service providers are highly unreliable.

ACKNOWLEDGMENTS

This work was funded by BAE Systems and the EPSRC.

REFERENCES

- [1] S. Stein, T. R. Payne, and N. R. Jennings, "Robust execution of service workflows using redundancy and advance reservations," *IEEE Transactions on Services Computing*, 2010.
- [2] —, "Flexible provisioning of web service workflows," *ACM Transactions on Internet Technology*, vol. 9, no. 1, pp. 2:1–2:45, 2009.
- [3] E. Deelman, J. Blythe, Y. Gil, C. Kesselman, S. Koranda, A. Lazarini, G. Mehta, M. A. Papa, and K. Vahi, "Pegasus and the Pulsar Search: From Metadata to Execution on the Grid," in *Parallel Processing and Applied Mathematics: 5th International Conference, PPAM 2003, Czestochowa, Poland, September 7-10, 2003*, vol. 3019 / 2004, September 2003.
- [4] T. Oinn, M. Greenwood, M. Addis, M. N. Alpdemir, J. Ferris, K. Glover, C. Goble, A. Goderis, D. Hull, D. Marvin, P. Li, P. Lord, M. R. Pocock, M. Senger, R. Stevens, A. Wipat, and C. Wroe, "Taverna: lessons in creating a workflow environment for the life sciences," *Concurrency and Computation: Practice and Experience*, vol. 18, no. 10, pp. 1067–1100, 2006.

- [5] S. Coles, J. G. Frey, M. B. Hursthouse, M. E. Light, K. E. Meacham, D. J. Marvin, and M. SurrIDGE, "ECSES - examining crystal structures using 'e-science': a demonstrator employing web and grid services to enhance user participation in crystallographic experiments," *Journal of Applied Crystallography*, vol. 38, no. 5, pp. 819–826, 2005.
- [6] Y. Gil, V. Ratnakar, E. Deelman, G. Mehta, and J. Kim, "Wings for pegasus: Creating large-scale scientific applications using semantic representations of computational workflows," in *Proc. 22nd Conf. on AI, Canada*, 2007, pp. 1767–1774.
- [7] T. M. Smith, C. Abajian, and L. Hood, "Hopper: software for automating data tracking and flow in DNA sequencing," *Comput. Appl. Biosci.*, vol. 13, no. 2, pp. 175–182, 1997.
- [8] K. Kochut, J. Arnold, A. Sheeth, J. Miller, E. Kraemer, B. Arpinar, and J. Cardoso, "IntelliGEN: A distributed workflow system for discovering protein-protein interactions," *Distributed and Parallel Databases*, vol. 13, no. 1, pp. 43–72, 2003.
- [9] A. O'Brien, S. Newhouse, and J. Darlington, "Mapping of scientific workflow within the e-protein project to distributed resources," in *Proceedings of UK e-science All Hands Meeting (AHM 2004)*, Nottingham, UK. EPSRC, 2004, pp. 404–409.
- [10] B. Ewing, L. Hillier, M. C. Wendl, and P. Green, "Base-calling of automated sequencer traces using phred. i. accuracy assessment," *Genome Research*, vol. 8, no. 3, pp. 175–185, 1998.
- [11] M. H. DeGroot and M. J. Shervish, *Probability and Statistics*, third edition ed. Addison-Wesley, 2002.