

Adaptive distributed replica–exchange simulations

BY ANDRE LUCKOW¹, SHANTENU JHA^{2,3,4,*}, JOOHYUN KIM²,
ANDRE MERZKY² AND BETTINA SCHNOR¹

¹*Institute of Computer Science, Potsdam University, 14482 Potsdam, Germany*

²*Center for Computation and Technology, and* ³*Department of Computer Science, Louisiana State University, Baton Rouge, LA 70803, USA*

⁴*e-Science Institute, Edinburgh EH8 9AA, UK*

Owing to the loose coupling between replicas, the replica–exchange (RE) class of algorithms should be able to benefit greatly from using as many resources as available. However, the ability to effectively use multiple distributed resources to reduce the time to completion remains a challenge at many levels. Additionally, an implementation of a *pleasingly distributed* algorithm such as replica–exchange, which is independent of infrastructural details, does not exist. This paper proposes an extensible and scalable framework based on Simple API for Grid Applications that provides a general-purpose, opportunistic mechanism to effectively use multiple resources in an infrastructure-independent way. By analysing the requirements of the RE algorithm and the challenges of implementing it on real production systems, we propose a new abstraction (BIGJOB), which forms the basis of the adaptive redistribution and effective scheduling of replicas.

Keywords: replica–exchange; Simple API for Grid Applications; Migol; adaptive; fault tolerance

1. Introduction

Several classes of applications are well suited for distributed environments. Probably the best known and most powerful examples are those that involve an ensemble of decoupled tasks, such as simple parameter sweep applications (Casanova *et al.* 2000). A slightly more complicated and challenging class of distributed applications are those that have a degree of coupling between individual subtasks. An interesting example of such applications is those based upon the *replica–exchange* (RE; Hansmann 1997; Sugita & Okamoto 1999) algorithm. RE simulations are used to understand physical phenomena—ranging from protein folding dynamics to binding affinity calculations.

The RE method involves the concurrent execution of multiple similar simulations—the *replicas*. The coupling between the replicas occurs via periodic exchange attempts between paired replicas. The exchange is typically infrequent compared with the run-time of each replica, and is small in terms of

* Author for correspondence (sjha@cct.isu.edu).

One contribution of 16 to a Theme Issue ‘Crossing boundaries: computational science, e-Science and global e-Infrastructure I. Selected papers from the UK e-Science All Hands Meeting 2008’.

communication bandwidth requirements. Thus, RE is *prima facie* a perfect algorithm to exploit distributed resources. We label such a class of algorithms as *pleasingly distributed*.

Most RE implementations are either infrastructure specific (Woods *et al.* 2005) or, if using multiple distributed resources, they require prior co-scheduling (Manos *et al.* 2008). Manos *et al.* (2008) is an important example of a first-generation Grid application, wherein the effectiveness of coupling multiple distributed resources for scientific problems has been demonstrated. The real power of distributed systems, however, arises from adaptive algorithms and implementations that provide applications with an agile execution model, and thus the ability to use resources dynamically as opposed to a static execution model inherited from parallel and cluster computing. Unfortunately, the barrier to the development of such adaptive applications is high and the infrastructure support is poor. Specifically, there is no implementation of an adaptive RE algorithm, which is able to both effectively and reliably use multiple distributed resources without prior scheduling as well as being independent of any specific infrastructure.

In this paper, we address some of the challenges and performance bottlenecks encountered when performing RE simulations over multiple distributed resources, such as the overall slowdown due to synchronization arising from the light coupling and the lack of co-scheduled resources. The unique contribution of this paper is the implementation of a RE framework that overcomes the described limitations by being able to adapt at run-time to a change in the availability of resources and application resource requirements. The framework builds upon preliminary work of integrating Simple API for Grid Applications (SAGA) and MIGOL to provide fault tolerance. While SAGA represents a well-defined, standardized interface for writing Grid applications, MIGOL provides the underlying middleware services to guarantee the correct and reliable execution of applications even in the presence of failures (Lucknow *et al.* 2008).

We provide evidence that, as more resources become available, our framework can opportunistically use these resources, leading to a reduction in the time to completion of the scientific problem. The remainder of the paper is structured as follows: in §2, we provide the basic ideas and advantages of using RE simulations to understand physical properties of a RNA system. The replica–exchange molecular dynamics (REMD) framework architecture and implementation are presented in §3. Section 4 discusses the new BIGJOB abstraction and the SAGA GLIDE-IN framework. In §5, we describe the deployment and results of different experiments using the SAGA-based RE framework on the TeraGrid (TG), and, in §6, we present data establishing the advantages of REMD for the physical system under study.

2. Hepatitis C virus (RNA) using replica–exchange

In molecular dynamics (MD) approaches, sufficient sampling of configurations is an important requirement for connecting atomistic results to macroscopic or thermodynamic quantities available from experiments. This provides an important motivation for researching ways to accelerate sampling and to enhance the ‘effective’ time scales studied. Generalized ensemble approaches—of

which REMD (Sugita & Okamoto 1999) is a prominent example—represent a promising attempt to overcome the general limitations of insufficient time scales, as well as specific limitations of inadequate conformational sampling arising from kinetic trappings. The fact that one single long-running simulation can be substituted for an ensemble of loosely coupled shorter-running simulations make these ideal candidates for distributed environments.

RE simulations consist of two distinct components: the underlying simulation engine used for each replica, and the coupling mechanism between the individual replicas. The degree and frequency of coupling and exchange can be either regular (Sugita & Okamoto 1999) or irregular (Shirts & Pande 2001). An example of the latter—parallel replica dynamics as implemented in Folding@home (Folding at home. See <http://folding.stanford.edu/>) involves coordination between replicas only when an ‘event’ occurs. By contrast, for regular RE applications, attempts to exchange states between certain pairs occur at fixed intervals.

The hepatitis C virus (HCV) internal ribosome entry site (IRES) is recognized specifically by the small ribosomal subunit and eukaryotic initiation factor 3 before viral translation initiation. This makes it a good candidate for new drugs targeting HCV. Our aim is to use REMD to enhance the sampling of the conformational flexibility of the internal loop referred to as *HCV IRES IIIb CA variant* (Collier *et al.* 2002) as well as the equilibrium energetics. The model of the physical system under investigation in this work comprises an RNA system of nucleotides; the total number of atoms in the simulating box is 21 887—including the RNA system, explicit water molecules and ions for neutralization of the system. The initial conformation of the RNA is taken from the NMR structure (PDB ID: 1PK7).

3. Implementing distributed replica–exchange using SAGA/MIGOL

(a) Replica–exchange manager architecture

The framework comprises three components, the RE-MANAGER, the REPLICAGENT and the MIGOL infrastructure. The *RE-MANAGER*, also referred to as task manager, is deployed on the user’s desktop and provides the interface to the overall RE run. It orchestrates all replicas, which involves file staging, job spawning and the conduction of the exchange attempts, using the SAGA APIs.

The second element is the task agent, the *REPLICAGENT*, which resides on the machines where the RE replicas are executed. The *REPLICAGENT* is launched using SAGA CPR and MIGOL. Nanoscale molecular dynamics (NAMD) (Phillips *et al.* 2005), a highly scalable, parallel MD code, is used to carry out the MD simulation corresponding to each replica run. It is important to mention that any other MD or Monte Carlo code could be used just as simply and effectively. Finally, MIGOL handles the reliable execution of the *REPLICAGENT* and the replicas, i.e. the submission, the monitoring and, if required, the recovery of replicas or the application itself.

(b) Replica–exchange logic

RE simulations involve the running of multiple replica jobs. Each replica is assigned a different temperature. Depending on the number of processes n , the RE-MANAGER creates $n/2$ pairs of replicas. Before launching a job, the RE-MANAGER ensures that all required input files are transferred to the respective

resource. For this purpose, the SAGA File API and the GRIDFTP adaptor are used. The replica jobs are then submitted to the resource using the SAGA CPR API and the MIGOL/GRAM middleware.

When all replicas reach a pre-determined state (e.g. the NAMD job finishes after a fixed number of steps), the decision as to whether to pairwise exchange temperatures between neighbouring replicas is determined by the METROPOLIS scheme. The run of an ensemble of replicas in parallel and the subsequent pairwise exchange attempt are referred to as *generation*. No two replicas can belong to different generations. If the exchange attempt is successful, parameters such as the temperature are swapped. Both jobs are then relaunched. Often the METROPOLIS scheme returns a negative result, and an exchange is not carried out; thus, it is difficult to respond to a possible exchange speculatively.

(c) *Deploying on production environments*

The RE framework has been successfully deployed on Louisiana optical network initiative (LONI) and TG production resources (Lucknow *et al.* 2008). In these environments, a significant slowdown was observable, in particular when running a larger number of replicas. A major reason for this slowdown was the fact that the restarted replicas are required to queue again at the local scheduler. In pathological cases, the complete system came to a halt solely due to a single-crowded or slow resource.

To avoid such bottlenecks, the multiple subtasks that constitute distributed applications need to avoid requeuing at the system batch queue level. Additionally, distributed applications that are decomposable into subtasks should be able to respond to the dynamic availability of resources. Unfortunately, current infrastructures do not support such dynamic scheduling directly. To provide this capability to applications, we need (i) abstractions that enable agile execution models via application-level allocation of resources, and (ii) different adaptivity strategies that determine how resources are efficiently used. Section 4 describes the extensions to the simple RE framework that enables efficient scheduling of subtasks and supports adaptive applications.

4. Adaptive replica–exchange: abstractions and implementation

As mentioned before, the use of multiple *simple* Grid jobs to execute many replicas has a severe limitation: all simple jobs must queue at the resource management system, i.e. a single delayed job can cause an overall slowdown. We overcome this issue by using an efficient dispatching scheme, which builds upon the ability to cluster replicas using the novel BIGJOB abstraction before submission. Based on this abstraction, we propose different strategies that address the dynamic conditions of distributed environments.

(a) *Abstractions*

A common approach to avoid queuing delays is the use of place-holder jobs, which are able to dispatch several subjobs without each subjob needing to queue at the local scheduler. A specific mechanism to support this pattern is the *GLIDE-IN* abstraction, in reference to the Condor *GLIDE-IN* system (Frey *et al.* 2002),

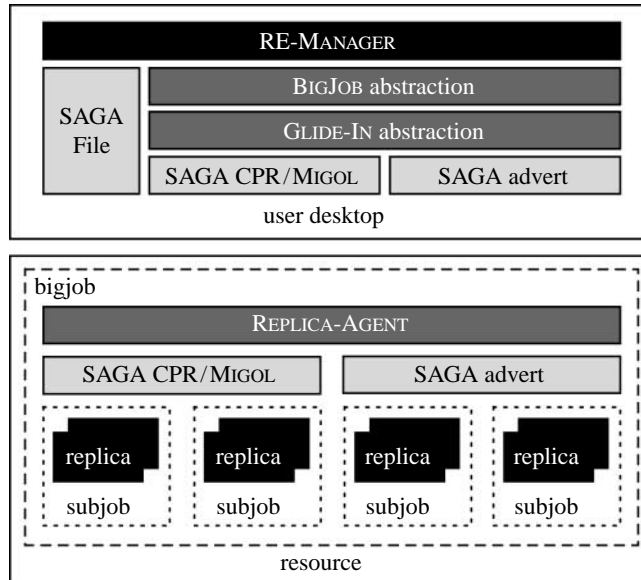


Figure 1. RE-MANAGER abstractions. The BIGJOB abstraction provides the capability to cluster subjobs into a larger bigjob, and is implemented on top of the GLIDE-IN abstraction. (Black rectangles, RE framework; dark grey rectangles, SAGA GLIDE-IN framework; light grey rectangles, SAGA.)

which pioneered this idea. A GLIDE-IN job requests a sufficiently large chunk of resources; smaller subjobs can then rapidly be executed through the GLIDE-IN job. By avoiding the high initial costs for queuing each individual replica job, the time to completion can be dramatically reduced.

Figure 1 summarizes the abstractions developed and used in this work to support the clustering of subjobs into larger bigjobs and the effective dispatching of the subjobs. The specific capability to cluster subjobs is provided to the application via the *BIGJOB* abstraction. The SAGA GLIDE-IN abstraction is used to support the commonly occurring place-holder job pattern. The BIGJOB abstraction defines a `big_job` and `sub_job` object; for each `big_job` object, a GLIDE-IN job with the desired number of resources is started, and `sub_job` objects, which correspond to individual replicas, are mapped to a `big_job` using the `jobID` as reference. It is helpful to reiterate that, although there is a `big_job` object, it is submitted as a GLIDE-IN job. Also, the BIGJOB abstraction in turn uses the GLIDE-IN abstraction to map the individual bigjob and subjobs to physical resources.

(b) Implementation

The RE framework has been extended to support the BIGJOB and GLIDE-IN abstractions. BIGJOB provides the ability to cluster subjobs; GLIDE-IN allows the effective scheduling of the subjobs. As illustrated in figure 2, the RE-MANAGER uses the `big_job` and `sub_job` objects as replacements for the `job` object defined by SAGA CPR. The `big_job` and `sub_job` objects behave similarly to regular SAGA CPR job objects. Thus, the RE-MANAGER, which is the application in this case, does not require any extensive modification; all that the RE-MANAGER has to do is to provide a mapping from a `sub_job` to a suitable `big_job` via a `jobID`.

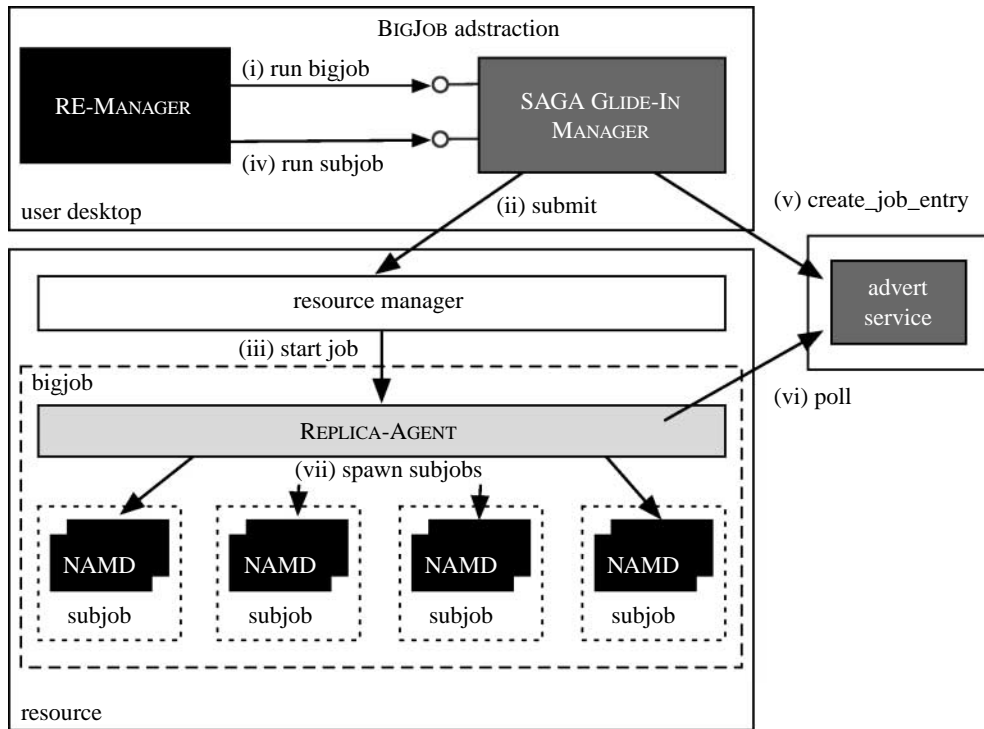


Figure 2. RE-MANAGER and SAGA GLIDE-IN framework. The GLIDE-IN job (REPLICA-AGENT) is used as place-holder job for all replica subjobs running on a single cluster. The RE-MANAGER controls both the REPLICA-AGENT and the replica jobs using the BIGJOB abstraction. (Black rectangles, RE framework; dark grey rectangles, SAGA GLIDE-IN framework; light grey rectangles, SAGA.)

The SAGA GLIDE-IN implementation comprises three components: (i) the *GLIDE-IN MANAGER* that provides the GLIDE-IN abstraction and manages the orchestration and scheduling of GLIDE-IN jobs (which in turn allows the management of both bigjob objects and subjobs) (ii) the *REPLICA-AGENT* that represents the GLIDE-IN job and thus, the application-level resource manager on the respective resource, and (iii) the *advert service* that is used for communication between the GLIDE-IN MANAGER and REPLICA-AGENT.

Before running regular jobs, an application, in this case the RE-MANAGER, must initialize a `big_job` object. The GLIDE-IN MANAGER then queues a GLIDE-IN job, which actually runs a REPLICA-AGENT on the respective resource. For this `big_job` instance, the specified number of resources is requested. Subsequently, `sub_job` objects can be submitted through the GLIDE-IN MANAGER using the jobID of the `big_job` as reference. The GLIDE-IN MANAGER ensures that the subjobs are launched onto the correct resource based upon the specified jobID using the right number of processes.

Communication between the REPLICA-AGENT and GLIDE-IN MANAGER is carried out using the SAGA advert service, a central key/value store. For each new job, an advert entry is created by the GLIDE-IN MANAGER. The REPLICA-AGENT periodically polls for new jobs. If a new job is found and resources are available, the job is dispatched, otherwise it is queued. Furthermore, the agent encapsulates

local machine-specific settings. The REPLICIA-AGENT ensures that the right combination of compiler, message passing interface (MPI) library and NAMD executable is used.

(c) Adaptive replica scheduling

Distributed applications including RE simulations must be able to deal with time-varying resource availabilities. An application is referred to as *dynamic* when either its resource requirements or the availability and usage of resources change during its run-time. *Adaptivity* is a mechanism to respond to dynamic changes; a dynamic application may deploy multiple adaptive strategies or choose between competing adaptive strategies. For an application to be adaptive, it is necessary for it to be able to effectively use an expanded or reduced set of resources; additionally, for an adaptive application to be scalable, it must also be able to determine which resources to use efficiently. For resource determination, our framework currently relies on a static, user-defined mapping of replicas and resources. In the remainder of this paper, we will focus on dynamic resource usage (and not on dynamic resource determination or optimization).

For jobs that want to maximize their throughput, the ability to adapt to dynamically changing resource loads is critical. It is equally important for long-running applications to be able to support an agile execution model allowing the effective usage of resources as they become available. Specifically, there are different ways a RE simulation can respond to a change in the number of resources required/available:

- *Scenario A*. By increasing the number of processes assigned to each replica, the time-to-completion can be reduced. In addition, resources can be partitioned in a way that balances the different speeds of resources. For example, by adding processes to a delayed replica, bottlenecks due to synchronization of replicas can be avoided.
- *Scenario B*. As resources become available, the number of replicas can be adjusted. Depending on the underlying physics model, the additional replicas can be used to either refine the temperature range (adaptive sampling) or to extend the temperature range (enhanced dynamics). This REMD approach is also referred to as *cool walking* (Brown & Head-Gordon 2003). Our framework supports both adaptive strategies.

5. Distributed replica-exchange on the TeraGrid

To evaluate the performance of the RE manager, several experiments have been conducted on TG and LONI resources. The resources used are: Ranger (TG), Abe (TG) and QueenBee (QB; both TG and LONI). The RE-MANAGER was configured to run a parallel NAMD simulation with up to 16 replicas sampling a temperature range between 300 and 450 K. Replica-exchanges are carried out between pairs of replicas. Thus, there are up to eight exchanges in each generation. Each test run comprises up to 64 attempted exchanges; each replica can use up to 24 MPI processes and runs for 500 time steps between exchange attempts. The metric used is the time to completion for 64 attempted exchanges.




Figure 3. Replica size adaptivity (scenario A). As the number of resources available increases, the number of cores assigned to a replica (a NAMD job) is dynamically adjusted, while keeping the number of replicas constant at 16. There are four GLIDE-INS with 32 cores each. As more resources become available, more cores are assigned to each replica, which leads to a reduction in T_c .

Initially, we investigated the effect of the GLIDE-IN framework on the time to completion (T_c) using 16 replicas on QB. Using the GLIDE-IN framework, there was a reduction in T_c from 52 to 26 min on average, which corresponds to a decrease of 50 per cent. In the best case, improvements of up to 70 per cent were observed. This effect is attributed to the elimination of the queuing times for every subjob. Once the REPLIC-AGENTS become active, replicas can be dispatched without requiring interactions with the local scheduler.

Further, we performed tests for the two adaptive scenarios. In scenario A, the number of replicas was kept constant (conventional REMD) and the *replica size*, i.e. the number of MPI processes assigned to each replica, was varied as more resources became available. In scenario B, the *replica number*, i.e. the number of replicas participating in a generation, was varied (cool walking). We compare T_c for 64 attempted exchanges on different sets of distributed resources and GLIDE-IN configurations.

In scenario A, up to three different resources are used; the number of GLIDE-INS varies from 4 up to 12, while the number of replicas used is fixed at 16. Thus, the size of the individual replicas is varied. When a resource is being used, it runs four GLIDE-INS, and each GLIDE-IN job has a constant size of 32 cores. Although, the number of GLIDE-IN jobs on a resource is fixed at 4 for reasons of simplicity, our results will hold for general values. If all resources (Abe, QB and Ranger) were being used, there would be 12 GLIDE-IN jobs with 32 cores each, i.e. a total of 384 cores would be available. GLIDE-INS are submitted to 1, 2 or 3 statically configured resources. After submission, the GLIDE-INS are subject to different queuing delays at the local schedulers. To reflect these different loads, the number of MPI processes assigned to each replica is dynamically increased as new resources become available. Depending on the number of available resources, between 8 (for 4 GLIDE-INS) and 24, MPI processes (for 12 GLIDE-INS) are used for each replica.

Figure 3 shows the results of the distributed run. In spite of the overhead for migrating replicas to newly available resources, a notable decrease in T_c of up




Figure 4. Replica number adaptivity (scenario B): in this scenario, the framework dynamically adjusts the number of replicas. A constant number of four GLIDE-INS with 64 cores each are distributed across one, two or three machines; the size of each replica is kept fixed at 16 cores. Once again, the greater the number of distributed resources that can be used, the smaller the T_c .

to 15 min can be observed as the number of resources increases from one to three. Although, the efficiency, defined as run-time on one resource divided by the run-time on multiple resources scaled by the number of resources, is only approximately 0.5, this is a limitation of the used set-up rather than a general scalability barrier. The set-up comprises rather short NAMD jobs; in particular, during the initial phase, jobs are often migrated to other resources, which mainly cause this overhead. During longer runs, this overhead is negligible. What is also very important to note is the reduced fluctuation in the T_c when multiple resources are used. This is indicative of the fact that there is a reduction in the sensitivity to queue loads—something that applications on real-production environments have to battle with.

In scenario B, the capability of the RE-MANAGER to adaptively adjust the number of replicas by varying either the range of temperature or the specific temperatures simulated is evaluated. For this scenario, four GLIDE-INS with 64 cores each are distributed across one, two or three different distributed resources. This means that the total of 256 cores is allocated on (i) Abe only, (ii) Abe and Ranger, and (iii) Abe, Ranger and QB. Each replica has a constant number of 16 MPI processes, and thus any GLIDE-INS can run exactly four replicas when active.

At the beginning of the experiment, all four GLIDE-INS are submitted to one, two or three resources (statically configured). Similar to scenario A, different queueing delays usually occur. Consequently, not all GLIDE-INS start simultaneously. Using the adaptive temperature sampling algorithm, the number of replicas is dynamically varied as the number of active GLIDE-INS increases. Depending on the number of running GLIDE-INS, the ensemble consists of 4–16 replicas (increase in steps of 4).

As shown in figure 4, T_c decreases with the number of resources used. With the simulation distributed onto a greater number of resources, the probability that a single heavily used resource delays the overall progress of the simulation is




Figure 5. The plots show the time series of the average times (solid curve) between exchange attempts and the number of active GLIDE-INS (dotted curve) over a 6 hour run on the TG.

reduced. The results clearly demonstrate the benefits of the adaptive replica scheme—it is favourable to instantly use resources as they become active, instead of waiting for the complete set of nodes to become available.

6. Results: enhanced sampling of hepatitis C virus

In §4, we discussed the design and performance of the framework and showed how it enables the effective usage of multiple resources. In §2, we provided motivation for why the RE approach is required to understand the energetics and conformational flexibility of the internal loop of the HCV. The effectiveness of the RE approach in increasing the rate of convergence to equilibrium (Boltzmann distribution) or enhancing the sampling can be measured by the rate of attempted exchanges (Lei & Duan 2007), or equally by the inverse of the *time* between attempted exchanges.

REMD simulations (scenario A) were performed for HCV IRES IIIb CA variant using the model described in §2. The replicas covered a temperature range from 300 to 450 K, thus fixing the number of replicas to 16. To demonstrate the effectiveness of our RE framework, we analyse a typical time series of the number of resources used (available) and thus, the number of active GLIDE-INS during a 6 hour run on a production infrastructure. Each GLIDE-IN has 32 cores; thus, the number of cores for each replica is determined by the numbers of active GLIDE-IN jobs. Figure 5 shows how the number of active GLIDE-INS changed during the 6 hour interval. To start with, there were only enough processors to support one GLIDE-IN job, but after approximately 2 hours, there were enough to activate two GLIDE-INS; after 3 hours, three GLIDE-INS are running. Eight GLIDE-INS were activated before the 6 hour time limit. As the number of GLIDE-INS increases, the number of processors assigned to each replica increases, with the physically important consequence that there is a concomitant decrease in the average time between exchange attempts (solid curve in figure 5).




Figure 6. The upper plot illustrates how the average time (solid curve) between exchange attempts decreases as the number of GLIDE-Ins (dotted curve) increases. At the same time, the speed-up grows.

The average time between exchange attempts as a function of the number of active GLIDE-Ins and the speed-up—measured as the inverse of the average time normalized by the time taken with one GLIDE-IN (thus the value of 1 for the one GLIDE-IN case)—is shown in figure 6.

7. Conclusion and future work

In summary, SAGA provides a well-defined and sufficiently powerful interface to develop the required abstractions to support adaptive distributed RE applications. SAGA allows the simple decoupling of the RE orchestration logic from the underlying distributed infrastructure. The SAGA GLIDE-IN framework represents the first known instance of creating a run-time, system-level abstraction for distributed systems from basic programming interfaces. Both SAGA and the RE framework are general purpose components, while remaining extensible. The adaptive RE framework has been successfully deployed on TG and LONI resources. Using the BIGJOB abstraction, RE subjobs can efficiently be dispatched reducing the time to completion by up to 70 per cent. The use of different adaptivity strategies to dynamically use additional resources led to a further reduction in the time to completion. Importantly, we have shown how our REMD framework was used to enhance the sampling and rate of convergence for a biological system, the inner loop of the HCV IRES IIIb CA variant, which is an important drug-delivery target.

In the future, we will refine our RE framework by making it more adaptive towards dynamic environments, e.g. by deploying an asynchronous RE scheme as described by Gallicchio *et al.* (2007). At the same time, we will improve our RE infrastructure to support further adaptive strategies for resource determination and usage. While it has been shown that resources can efficiently be allocated with the BIGJOB abstraction, a mechanism for dynamic resource discovery and for intelligent placements of jobs will be beneficial to further

decrease the time to completion. Various approaches for the resources determination have been proposed, e.g. batch queue prediction (Nurmi *et al.* 2007; Chakraborty *et al.* 2009) and advance reservation-based schemes (Jeske *et al.* 2007).

This work would not have been possible without the support of the wider SAGA team. Important funding for SAGA specification and development has been provided by the UK EPSRC grant no. GR/D0766171/1 (via OMII). S.J. acknowledges the e-Science Institute, Edinburgh, for supporting the research theme, ‘Distributed Programming Abstractions’. We would also like to thank Yaakoub el-Khamra for useful discussions. This work has also been made possible thanks to computer resources provided by the TG and LONI.

References

- Brown, S. & Head-Gordon, T. 2003 Cool walking: a new Markov chain Monte Carlo sampling method. *J. Comput. Chem.* **24**, 68–76. (doi:10.1002/jcc.10181)
- Casanova, H., Obertelli, G., Berman, F. & Wolski, R. 2000 The AppLeS parameter sweep template: user-level middleware for the grid. *Sci. Program.* **8**, 111–126.
- Chakraborty, P., Jha, S. & Katz, D. S. 2009 Novel submission modes for tightly coupled jobs across distributed resources for reduced time-to-solution. *Phil. Trans. R. Soc. A* **367**, 2545–2556. (doi:10.1098/rsta.2009.0054)
- Collier, A., Gallego, J., Klinck, R., Cole, P., Harris, S., Harrison, G., Aboul-ela, F., Varani, G. & Walker, S. 2002 A conserved RNA structure within the HCV IRES eIF3-binding site. *Nat. Struct. Biol.* **9**, 375–380.
- Frey, J., Tannenbaum, T., Livny, M., Foster, I. & Tuecke, S. 2002 Condor-G: a computation management agent for multi-institutional grids. *Cluster Comput.* **5**, 237–246. (doi:10.1023/A:1015617019423)
- Gallicchio, E., Levy, R. & Parashar, M. 2007 Asynchronous replica exchange for molecular simulations. *J. Comput. Chem.* **29**, 788–794. (doi:10.1002/jcc.20839)
- Hansmann, U. 1997 Parallel tempering algorithm for conformational studies of biological molecules. *Chem. Phys. Lett.* **281**, 140–150. (doi:10.1016/S0009-2614(97)01198-6)
- Jeske, J., Luckow, A. & Schnor, B. 2007 Reservation-based resource-brokering for grid computing. In *Proc. German e-Science Conf., Baden-Baden, Germany*.
- Lei, H. X. & Duan, Y. 2007 Improved sampling methods for molecular simulation. *Curr. Opin. Struct. Biol.* **17**, 187–191. (doi:10.1016/j.sbi.2007.03.003)
- Luckow, A., Jha, S., Kim, J., Merzky, A. & Schnor, B. 2008 Distributed replica-exchange simulations on production environments using SAGA and Migol. In *Proc. 4th IEEE Int. Conf. on e-Science, Indianapolis, IN*.
- Manos, S., Mazzeo, M., Kenway, O., Coveney, P. V., Karonis, N. T. & Toonen, B. 2008 Distributed MPI cross-site run performance using MPIg. In *HPDC’08: Proc. 17th Int. Symp. on High Performance Distributed Computing*, pp. 229–230. New York, NY: ACM.
- Nurmi, D., Brevik, J. & Wolski, R. 2007 QBETS: queue bounds estimation from time series. *SIGMETRICS Perform. Eval. Rev.* **35**, 379–380. (doi:10.1145/1269899.1254939)
- Phillips, J. *et al.* 2005 Scalable molecular dynamics with NAMD. *J. Comput. Chem.* **26**, 1781–1802. (doi:10.1002/jcc.20289)
- Shirts, M. & Pande, S. 2001 Mathematical analysis of coupled parallel simulations. *Phys. Rev. Lett.* **86**, 4983–4987. (doi:10.1103/PhysRevLett.86.4983)
- Sugita, Y. & Okamoto, Y. 1999 Replica-exchange molecular dynamics method for protein folding. *Chem. Phys. Lett.* **314**, 141–151. (doi:10.1016/S0009-2614(99)01123-9)
- Woods, C. J. *et al.* 2005 Grid computing and biomolecular simulation. *Phil. Trans. R. Soc. A* **363**, 2017–2035. (doi:10.1098/rsta.2005.1626)