SmartGridRPC: A New RPC Model for High Performance Grid Computing and its Implementation in SmartGridSolve

Thomas Brady
B.Sc. Hons., M.Sc.

The thesis is submitted to University College Dublin
for the degree of Doctor of Philosophy

Head of Department: Joe Carthy
Supervisor of Research: Dr. Alexey Lastovetsky

School of Computer Science and Informatics,
College of Engineering, Matematical and Physical Sciences,
University College Dublin.

June 2009

I would like to dedicate this thesis to my mother for her help, support and encouragement throughout my whole studies, without which this thesis would have never been written.

Acknowledgements

I would like to thank my supervisor Alexey Lastovetsky for his guidance, advice, support and enthusiasm throughout my research. He was available at all times to offer valuable discussion and guidance which helped shape the research presented in this thesis. I would also like to thank Jack Dongarra and Keith Seymour for their valuable discussions, colloboration and support in this project.

I would like to thank the other postgraduates in the HCL laboratory, Brett, Robert, Michele, Maureen, Ravi, Xin and Vladimir. I would like to thank Robert and Michele for their discussions and advice during the implementation of SmartGridSolve. I would also like to thank Robert for always being willing to proof read my papers and I would like to thank Brett for his work setting up the cluster.

I would like to thank Science Foundation Ireland for their financial support with my grant funding.

Finally I would like to thank all my friends and family, in particular my mother, for their constant support and motivation.

Abstract

This thesis presents the SmartGridRPC model, an extension of the GridRPC model, which aims to achieve higher performance.

The traditional GridRPC model provides an API and model for mapping individual tasks of an application in a distributed Grid environment which is based on the client-server model or star network topology.

The SmartGridRPC model provides an API and model for mapping a group of tasks of an application in a distributed Grid environment which is based on the fully connected network topology.

The SmartGridRPC programming model/API, its implementation and its performance advantages over the GridRPC model are outlined in this thesis. In addition experimental results using a real-world application are also presented.

Contents

1 Introduction **1**

2 Motivation **8**

2.1 Motivation: GridRPC model 8

2.2 Motivation: SmartGridRPC model 10

3 Related research **14**

4 GridRPC programming model and API **21**

4.1 Design of the GridRPC programming model 22

4.2 GridRPC : API and semantics 23

4.3 GridRPC : A GridRPC application 24

5 SmartGridRPC programming model and API **27**

5.1 Design of the SmartGridRPC programming model 28

5.2 SmartGridRPC: API and semantics 33

5.3 SmartGridRPC: A SmartGridRPC application 35

6 GridSolve: Implementation of the GridRPC model **38**

6.1 GridSolve: Agent discovery 38

6.2 Run-time GridRPC task call 39

6.3 GridSolve: Performance models 41

6.4 GridSolve: Mapping heuristic 46

7 SmartGridSolve: Implementation of the SmartGridRPC model **48**

7.1 SmartGridSolve: Agent discovery 48

7.2 Run-time of client application 49
7.3 SmartGridSolve performance models 52
7.3.1 Network graph 52
7.3.2 Task graph 53
7.4 Mapping Solution Graph 56
7.5 Communication model 62
7.6 Fault tolerance 65

8 SmartGridRPC benchmark application: The evolution of a cluster of galaxies 66
8.1 GridRPC implementation of Hydropad 70
8.2 SmartGridRPC implementation of Hydropad 71

9 Experimental results 77
9.1 Experiments with the GridSolve version of Hydropad 78
9.2 Experiments with the SmartGridSolve version of Hydropad 80

10 Conclusion 88
10.1 Contributions 92

A Appdx A - SmartGridSolve Manual 96
A.1 Using SmartGridSolve 96
A.2 SmartGridSolve API 97
A.2.1 The grpc_map() function 97
A.2.2 The grpc_map_ft() function 98
A.2.3 The grpc_local() function 100

B Appdx B - Hydropad Manual 103
B.1 Introduction to Hydropad 103
B.2 Installing Hydropad 104

References 116

List of Figures

4.1 Overview GridRPC model and API 21

5.1 Overview SmartGridRPC model and API 28

6.1 GridSolve - Agent discovery 40

7.1 SmartGridSolve - Agent discovery 50
7.2 SmartGridSolve - The network graph 53
7.3 SmartGridSolve - The task graph 55
7.4 SmartGridSolve - The mapping solution graph 57
7.5 Communication transactions of task 0 of the mapping solution in figure 7.4, which happen prior to the execution of task 0. 64
7.6 Communication transactions of task 0 of the mapping solution in figure 7.4, which happen subsequent to the execution of task 0. . . 64

8.1 Example of Hydropad output 67
8.2 Overview of the Hydropad application 69
8.3 Task graph for two evolution steps 74

9.1 Evolution time step of the local and GridSolve computation on client C100-256 81
9.2 Execution times of the GridSolve and SmartGridSolve version of Hydropad on client C100-256 81
9.3 Execution times of the GridSolve and SmartGridSolve version of Hydropad when the client machine C1-256 has 256MB of memory 86

9.4 Execution times of the GridSolve and SmartGridSolve version of Hydropad when the client machines are C1-1 and C100-256 87

10.1 SmartGridSolve extensions 94

Chapter 1

Introduction

The idea of distributed computing has been around for years. What distinguishes Grid computing from typical distributed computing or cluster computing is that Grids tend to be more loosely coupled, heterogeneous, and geographically dispersed [1].

A Grid programming model can be present in many different forms, for example: a language, a library API, or a tool with extensible functionality [2].
The following are common Grid programming models:

- RPC model [3] (GridRPC [4], Java RMI [5]).
- Shared state models (javaspaces [6], publish/subscribe [7]).
- Message Passing Models (MPI [8]).
- Hybrid Models (OpenMP/MPI [9], OmniRPC [10], MPI).
- Peer-to-peer (JXTA [11]).
- Grid API (Globus Toolkit [12], GAT [13], SAGA [14]).
- Application frameworks (Cactus [15], GridSphere [16]).
- Component models (Corba [17], CoG [18], Legion [19]).
- Web service model (OGSA [20]).

The remote procedure call (RPC) programming model [21] is the premier model for executing scientific applications in a distributed environment. Improving the performance of this programming model in the Grid environment is the focus of the thesis.

RPC provides a straightforward procedure for executing parts of an application on a remote computer. To execute a RPC, the application programmer does not need to learn a new programming language but merely uses the RPC API. Using the API the application programmer specifies the remote task to be performed, the server to execute the task, the location of the input data on the user's computer required by the task and the location on the user's computer where the results will be stored. The execution of the remote call involves transferring input data from the user's computer to the remote computer, executing the task on the remote server and delivering output data from the remote computer to the user's one.

GridRPC is a standard promoted by the Open Grid Forum which extends the traditional RPC for the Grid environment. A number of Grid middleware systems are GridRPC compliant including GridSolve [22], Ninf-G [23] and DIET [24].

A GridRPC system processes each GridRPC task call by first performing dynamic resource and task discovery, then mapping the task to a server and then executing the task on the mapped server. Since each GridRPC task call consists of these operations (discovery, mapping and execution) and each GridRPC task is processed individually, the GridRPC model imposes the restriction that these three operations are atomic and cannot be separated. As a result, each task has to be mapped separately and independently of other tasks of the application. Given that the model has the restriction that tasks have to be mapped individually and independently, the model can only support the minimization of the execution time of each individual task of the application rather than the minimization of the execution time of the whole application.

Another important aspect of the GridRPC model is its communication model. The communication model of GridRPC is based on the client-server model or star network topology. This means that tasks can be executed on any of the servers and inputs/outputs can only traverse the client-server links.

Mapping tasks individually on to star network results in mapping solutions that are far from optimal. If tasks are mapped individually, the mapping heuristic is unable to take into account any of the tasks that follow the task being mapped. Consequently, the mapping heuristic does not have the ability to optimally balance the load of computation and communication. Another consequence of mapping tasks in this way is that dependencies between tasks are not known at the time of mapping. Therefore this approach forces bridge communication. Bridge communication occurs when the output of one task is required as an input to another task. In this case, using the traditional GridRPC model, the output of the first task must be sent back to the client and the client then subsequently sends it to the server executing the second task when it is called.

Also, since dependencies are not known and the network is based on the client-server model, it is impossible to employ any parallelism of communication between the tasks in the group. For example, this could be implemented if there is a dependency between two tasks and the destination task is not executed in parallel or immediately after the source task. In theory, this dependent data could be sent to the destination task in parallel with any computation or communication on any other machine (client or other servers) which happens in the intervening time. But since tasks are mapped individually on to a star network, this parallelism of communication cannot be realized using the GridRPC model.

In this thesis, we propose an enhancement of the traditional GridRPC model which would allow a group of tasks to be mapped collectively on to a fully connected network. This would remove each of the limitations of the GridRPC model already described. The SmartGridRPC model has extended the GridRPC model to support collective mapping of a group of tasks by separating the mapping of tasks from their execution. This allows the group of tasks to be mapped collectively and then executed collectively.

In addition the traditional client-server model of GridRPC has been extended so that the group of tasks can be collectively executed on a network topology which is fully connected. This is a network topology where all servers can communicate directly or servers can cache their outputs locally.

There are a number of advantages of mapping tasks collectively on to a fully connected network. When mapping tasks individually, the communication load

and computation load of a single task are only considered. However when tasks are mapped collectively the communication load and computation load of multiple tasks can be considered together and therefore this load can be better distributed over the fully connected network. In addition the relationships between tasks can also be considered such as the data dependencies between tasks. This allows bridge communication to be eliminated by mapping these dependencies on to virtual links connecting servers. As a result, servers can send data directly to other servers and therefore do not need to send it via the client. Eliminating bridge communication can significantly decrease the overall communication time of an application and hence improve the overall performance of the application because communication is often the more time consuming phase in an RPC context.

In addition, this may also eliminate memory paging on the client which would otherwise occur when bridge communication is forced and intermediate results are stored on the client. Also since dependencies between tasks are known, it means that remote communication of one task can be parallelized with other computation and communication in the group.
Consequently, the mapping heuristic can improve the performance of that group by:

- more effectively balancing the load of computation of the group of tasks;
- more effectively balancing the load of communication of the group of tasks;
- reducing the overall volume of communication of the group by eliminating bridge communication either by caching or direct data transfers between servers;
- reducing memory usage and paging;
- increasing the parallelism of communication.

The client-server model of GridRPC results in a communication network which has a star topology. Therefore in this case for any given mapping of a group of tasks to remote servers there will be only one communication path between any pair of servers that should be considered when mapping. This path consists of two communication links connecting the servers with the client machine. Any

other path connecting the two servers obviously results in higher communication cost.

However, if the communication network is fully connected, then there will be multiple independent paths connecting the servers and each of these paths can be considered when mapping. In other words, for each mapping of a group of tasks to remote servers in a star communication network there is only one fixed communication scheme that can be employed. However, when a group of tasks is mapped on a fully connected network there are many communication schemes to choose from. These communication schemes may employ direct communication, server broadcast, client broadcast or caching. Therefore the mapping of a group of tasks on a fully connected network not only involves the mapping of tasks to servers but also the mapping of the dependencies between tasks on to the communication paths of the network. This increases the mapping solution space and allows for further optimization to be achieved by choosing the optimal paths for data to traverse between servers. This increase of the solution space means that the mapping heuristics implemented in the SmartGridRPC model have more potential of finding a closer to optimal solution than the mapping heuristics inherent in implementations using the standard GridRPC model.

GridSolve, is a middleware system that implements the GridRPC model. It enables users to solve complex scientific tasks remotely on distributed resources. GridSolve emphasises ease-of-use for the user and includes resource monitoring, mapping and service-level fault tolerance. In addition to providing Fortran and C clients, GridSolve enables Scientific Computing Enironments (SCEs) such as Matlab to be used as clients, so domain scientists can use Grid resources from within their preferred environments.

SmartGridSolve [25] is an extension of GridSolve, which makes the GridSolve middleware compliant with the SmartGridRPC model. SmartNetSolve [26] was previously implemented to make the NetSolve [27] middleware, which was the predecessor of GridSolve, compliant with the SmartGridRPC model.

The SmartGridSolve extension is interoperable with GridSolve. Therefore if GridSolve is installed with the SmartGridSolve extension, the user can choose whether to implement an application using the standard GridRPC model or the extended SmartGridRPC model. In addition SmartGridSolve is incremental to

the GridSolve system. Therefore if the SmartGridSolve extension is installed only on the client side, the system will be extended to allow for collective mapping. If SmartGridSolve is installed on the client side and on only some of the servers in the network, the system will be extended to allow for collective mapping on a partially connected network. If it is installed on all servers, the system will be extended to allow for collective mapping on the fully connected network.

The high level design of the SmartGridRPC model was inspired by both the design of mpC [28] and HeteroMPI [29]. SmartGridRPC uses the same approach of using the performance model of the application and the performance model of the executing heterogeneous network to optimally map the application on to the underlying heterogeneous network.

The contributions of this work include:

The design and concept of the SmartGridRPC model and API: The goal of the SmartGridRPC API was to allow an application programmer to easily trasform their existing GridRPC enabled applications into SmartGridRPC enabled application. Therefore an application programmer with a few simple changes can easliy benefit from all the performance enhancements of the SmartGridRPC model.

The SmartGridRPC model was designed so that it is both incremental and interoperable with the underlying GridRPC middleware. It is incremental meaning that if the exension is installed, the application programmer has the choice whether to implement their application using the SmartGridRPC model or the GridRPC model. It is interoperable in that a SmartGridRPC enabled client application can be executed on a network consisting of both standard GridRPC servers and SmartGridRPC enabled servers.

The design and implementation of SmartGridSolve: The GridSolve middle was extended to be SmartGridRPC compliant and this extension is called SmartGridSolve. The implementation of SmartGridSoilve requires extensions to the following aspects of GridSolve:

- Network Discovery.
- Task Discovery.

- Mapping Heuristics.
- Execution model.

Details of this implementation will be described in this thesis to demonstrate how a developer of a GridRPC middleware can extend their middleware so that it is SmartGridRPC enabled.

The thesis is outlined as follows. Chapter 2 gives the motivation of GridRPC and SmartGridRPC model. Chapter 3 outlines research papers which are related to the SmartGridRPC model. The GridRPC programming model and API is described in chapter 4. The SmartGridRPC model and API is described in chapter 5. Chapter 6, describes the implementation of the GridRPC model in GridSolve and chapter 7 describes the implementation of the SmartGridRPC model in SmartGridSolve. Chapter 8 outlines the Hydropad application, which is an astrophysics application used to benchmark the GridRPC model and the SmartGridRPC model. Chapter 9 gives experimental results which compare the GridRPC model with the SmartGridRPC model using the Hydropad application as a benchmark. The thesis will conclude with chapter 10.

Chapter 2

Motivation

2.1 Motivation: GridRPC model

The following are some of the key benefits of implementing GridRPC enabled applications:

- Improved performance of applications.
- Solution of large scale applications.
- More control over applications.
- Solution of hardware specific applications.
- Portability.
- Easy and powerful development of applications.

Improved performance of applications: The performance related benefits include the potential for faster solution of a problem of a given size and solution of problems of larger sizes . There are two main reasons for this. Firstly, if parts of the code can be executed in parallel on remote servers then the GridRPC model allows one to implement their parallel execution on remote servers. This parallelisation will decrease the computation time of the application. Secondly, if the Grid environment contains machines that are more powerful than the client machine, then remote execution of the tasks of this application on these more

powerful machines will also decrease the computation time of the application. However, this decrease of the computation time is not without cost. The application will pay the communication cost due to remote execution of the tasks. If the communication links connecting the client machine to the server machines are relatively slow, then the communication cost may be more than the increased performance of computation, resulting in the total execution time of the application being higher than its sequential execution on the client machine.

Solution of large scale applications: The GridRPC model provides a solution for applications which cannot be executed on a client machine due to their strong demands on the resources (memory, disk space etc.). In this case GridRPC provides a means to allocate these demands to remote servers in the Grid environment. For example, the remote execution of parts of a memory intensive application on remote servers could eliminate heavy paging that would otherwise occur on the client machine.

More control over applications: In some cases, applications that could be executed in a Grid environment could potentially be executed in a high performance computer (HPC) system. Unfortunately in a HPC system, where applications are executed in batch mode, the user will not have much control over the execution. Grid-enabled applications allow the user to have a high control over its execution because, although the tasks are being computed in remote servers, the main component of the application is running on the client machine. This can be important for applications that need a direct interaction with the data produced. For a given application it would be possible for a user to see intermediate results of the application. Furthermore, while the user/client is checking these results, they could decide to change some parameters of the application or restart the application.

Solution of hardware-specific applications: Some applications have a task that is inherently remote. For example, a task could be a proprietary precompiled binary which has been compiled for a specific architecture or a task may be tuned or tweaked to execute more efficiently on a specific type of hardware such as FPGA. Furthermore a task could require interaction with a resource that can only interface with a particular machine such as a telescope, video camera,

microscope etc. In such cases, an environment that allows the resources (including software) to be used on a particular computer is needed.

Portability: Since a GridRPC enabled application comprises of a client application and server-side compiled executables, the client application can be easily ported, compiled and executed on a new machine in the Grid environment. This does not require the recompilation of server-side task executables, which could make up a large proportion of the application.

An easy and powerful development paradigm: Any task which has been developed for remote execution for one GridRPC enabled application can be easily reused for other Grid applications. This situation can reduce the programmers effort on developing a Grid application. For example the programmer can use already existing tasks that they would not have the time or skill to write.

2.2 Motivation: SmartGridRPC model

The following are some of the key benefits of SmartGridRPC enabled applications over GridRPC enabled applications:

- Improved balancing of computation load.
- Reduced volume of communication.
- Improved balancing of communication load.
- Increased parallelism of communication.
- Reduced memory usage and paging.

Improved balancing of computation load: In GridRPC tasks get mapped individually on to a client server network. This could result in poor load balancing of computation. Since tasks get mapped individually, it is impossible to balance the load of computation of a group of tasks which are executing in parallel. If tasks are mapped individually, each task will be mapped without the knowledge of any of the subsequent parallel tasks. This means that if a large task follows a smaller task, the mapping heuristic will give the smaller task priority over

the larger task. This is because when the smaller task is mapped, the mapping heuristic cannot take into account that a larger task will be executing in parallel. Therefore it maps the smaller task to the faster server as this will yield the lowest execution time for this individual task. And when the mapping heuristic maps the larger task, it will assign it to the next fastest server as the fastest server is busy executing the previous task. This is poor load balancing of computation. However if you implement the collective mapping of the SmartGridRPC model, then the computation load can be better distributed over the network. In this case, if both tasks can be mapped collectively then the larger task would be mapped to the faster server and the smaller to the slower server. This improved balancing of the computation load will increase the performance of both tasks executing in parallel.

Reduced volume of communication: Since the GridRPC model maps tasks individually on to a client-server network, the model forces bridge communication between tasks. This occurs because dependencies are not known between tasks and data can only traverse the client-server links. As a result the source task can only send the dependent data to the destination task via the client. This requires two communication steps, the first from the source task to the client and the second from the client to destination task. However this can be eliminated with the SmartGridRPC model, where tasks can be mapped collectively on to a network which is fully connected. Since tasks are mapped collectively, dependencies between tasks are known. These dependencies can then be mapped on to virtual links connecting the source server to the destination server, which is only one communication step. Therefore the overall volume of communication required to be sent over the network will be reduced which would result in improved application performance. Moreover if the source task and destination task are both executing on the same server then this output could be cached to the local file system or cached in memory which would further reduce the overall communication on the network and increase the performance of the application.

Improved balancing of communication load: Since the GridRPC model is based on the client-server model, communication can only be mapped to client-server links. This may result in the client-server links becoming heavily loaded.

SmartGridRPC can increase the performance of an application by better distribution of the communication load over the network. When tasks are mapped collectively, the volume of communication of each task in the group of tasks is known. Since the sizes of inputs and outputs of each task are known and this communication is mapped onto a network which is fully connected, this communication can be better distributed over the fully connected network. This improved load balancing of communication will result in improved overall communication times and hence improved application performance.

Increased parallelism of communication: In much the same way that the GridRPC model improves on the RPC model with the parallelism of computation; the SmartGridRPC improves on the GridRPC model with the parallelism of communication. With the GridRPC model, the parallelism of communication is limited to the sending of inputs to a non-blocking task which is an asynchronous operation. With the SmartGridRPC model any communication on one machine can be done in parallel with computation or communication on any other. This asynchronous communication is only achievable since the dependencies between tasks are known prior to the execution of the group due to the collective mapping. This parallelism of communication can be advantageous if a task executing on one server has a dependency on another task which will be executed on another server and the destination task is not executed immediately after the source task. In this case, this communication can be done asynchronously. This means that the server initiates the communication but does not wait for it to finish. Therefore this communication can be done in parallel with any other computation or any communication on any other machine (client or servers) which happens in the intervening time. In addition this parallelism of communication can be beneficial if the client broadcasts an argument to more than one task which is to be executed on different machines. If any of the tasks are not executing immediately after the communication, then the communication to these tasks can be done in parallel with any computation on the client machine and computation or communication on any other server which happens in the intervening time. The same is true for server broadcast communication. This parallelism of communication reduces overall communication times and thus improves the overall performance of the group of tasks executing on the fully connected network.

Reduced memory usage and paging: The direct communication between servers and the data caching that the SmartGridRPC model implements means that intermediate results do not have to be sent back and stored on the client. This minimizes the amount of memory used on the client and the volume of communication necessary between client and server. This could eliminate paging on the client that would otherwise occur. This elimination of paging would considerably increase the performance of an application.

Chapter 3

Related research

This section examines those systems which implement the GridRPC model (GridSolve, Ninf-G, DIET) and their predecessors (i.e. NetSolve and Ninf) and will focus on the papers which mostly relate to our research. Specifically those papers which fall into the following categories:

- papers presenting extensions to the client-server model which implement direct communication between servers or data persistence.
- papers presenting extensions which extend the system so that a group of tasks can be collectively mapped.

These papers will be presented in chronological order and we will outline the limitations of each approach in comparison with the SmartGridRPC model.

Both NetSolve [30] and Ninf [31], the predecessor to the GridSolve and Ninf-G system, were started at roughly the same time. The projects were both started in 1994 and were first released in 1995. These systems were designed to resolve the difficulty of performing computational science problems over loosely connected geographically disperse networks. The computational libraries that the most common computational problems use may be highly optimized for only certain platforms and do not provide a convenient interface to other computer systems. Other libraries demand considerable programming effort from the user, who may not have the time to learn the required programming techniques. The resolution of these issues was the motivation behind both projects. These systems were

called Network Enabled Server (NES) or Problem Solving Environment (PSE) systems and employed a RPC-style model to perform remote computations.

In 1999, task farming [32] was introduced to NetSolve. The farming feature of NetSolve allowed a certain class of tasks called farming jobs to be processed collectively. A farming job fell into the class of embarrassingly parallel programs for which it is very clear how to partition the jobs for parallel programming environments. While these tasks were processed collectively they were not mapped collectively. Each task was individually mapped but computation loads of subsequent tasks were dynamically adjusted at run-time based on previous task response times.
The limitations of task farming are:

- It can only be implemented for a certain class of application where tasks do not have dependencies and can run independently (e.g. embarassingly parallel applications).
- Tasks are mapped individually and therefore the mapping heuristic cannot take advantage of characteristics of the group such as data dependencies.
- Conditional statements cannot exist in the scope of the task farming job.
- Client computation cannot exist in the scope of the task farming job.
- The group of tasks is called as one atomic call, therefore intermediate results cannot be viewed or analysed.

In 2000, task sequencing [33] was introduced to NetSolve. Using the task sequencing API a group of tasks could be processed collectively so that data dependencies could be analysed. This group of tasks is subsequently mapped on to a single server and if any data dependencies exist, the data would be stored locally and not sent back to the client. Therefore, using this API data persistency could be implemented and therefore if dependencies exist, bridge communication could be eliminated.
The limitations of task sequencing are:

- The group of tasks can only be mapped to a single server.

  - This computation load could be better distributed over a number of servers.
  - There may not be a server in the environment that can execute all tasks.

- Conditional statements such as for, if and while are forbidden between tasks.
- Client computation cannot exist in the scope of the task sequencing job.

In 2002, data transfers between servers were introduced to NetSolve [34]. This was achieved with an added function to the API that allowed the user to explicitly outline data dependencies. If there are two tasks which have a data dependency and are executing on different servers, this data would be stored in the source server when it finished execution and then the destination server would pull the argument from that server when it is called for execution.
The limitations to this approach are:

- Tasks are mapped individually.
- Push communication cannot be implemented when tasks are mapped individually. A push transfer is a transfer of data where the source server initiates the communication from the source server to the destination server. This can be done in parallel with other computation or communication in the group of tasks. This is because it can be transferred asynchronously as the source server and destination server are not required to wait for the transaction to complete before doing other work. In addition the client is not required to wait and therefore can perform local computation, client-server communication or initate other remote computation or communication. This can significantly increases the amount of parallelism in a given group of tasks.

  A pull data transfer is a transfer of data where the destination server requests the data and it is then sent from the source server to the destination server. This cannot be done asynchronously as the destination task has to wait to receive the data before it can begin execution. Also if the destination task is a blocking task, client or remote computation and communication

cannot be parallelised with this data transfer. However using push communication they can be overlapped.

- The user has to explicitly specify dependencies.
  - More labour intensive.
  - More prone to error.

This feature was later implemented in the GridRPC model in the GridSolve system [35], DIET [36] and NINF-G [37] and had the same limitations.

In July 2002, the DIET system was launched which implemented an architecture where the scheduler/agent is scattered across a hierarchy of Local Agents and Master Agents. The motivation for this architecture was it was more scalable and solved the problem of bottlenecks in a centralised agent/scheduler when many clients try to access several servers. In addition the DIET system employed direct communication between servers and data persistency. Where a dependency existed between tasks this output would remain on the source server. When the destination task is called for execution this data would be pulled from the source server. If the source server is the same as the destination server this output would be stored and retrieved locally (data persistency).
The limitations of this approach are:

- Tasks are mapped individually.
- Push communication cannot be implemented when tasks are mapped individually
- Increased communication times since communication cannot be done in parallel with computation or other communication.

In 2002, Distributed Storage Infrastructure (DSI) [38] was implemented in NetSolve. The DSI was another feature that attempted to minimize data movement in the NetSolve middleware. With DSI, data could be stored in storage depots which are close to servers which require the data. Instead of having multiple transmissions of the same data, DSI allows the transfer of data once from the client to a storage depot. A data handle is then used to retrieve only the relevant

portions of the stored data when running computations. This reduced communication times but again did not change how tasks are fundamentally mapped.

Also in 2002 the Global Grid Forum (now known as the Open Grid Forum) standardized the RPC mechanism for Grid computing with the GridRPC programming model and API [4]. This was implemented in NetSolve [39] and Ninf-G [23]. Ninf-G is the second generation of Ninf and was implemented on top of the Globus toolkit [40]. The Globus toolkit provides a reference implementation of standard protocols and it deploys Globus Security Infrastructure so that all the components of Ninf-G are protected properly. In 2003 GridSolve which is the second generation of NetSolve was released and provided full support for the GridRPC model.

In 2005, the GridRPC model was implemented in DIET [24]. This paper also introduced the Data Tree Manager (DTM). The DTM allows data to be left on a server after computation and then retrieved by another server during its computation. This paper described how JUXMEM (Juxtaposed memory) could be used in the DIET system to allow servers to share memory data. Both the DTM and JUXMEM avoided multiple transmissions of the same data from a client to a server but again tasks could only be processed and mapped individually. The limitations to this approach are:

- Tasks are mapped individually.
- Push communication is not implemented.

SmartNetSolve was designed in 2004, implemented in 2005 and was first presented in [26] in April 2006. SmartNetSolve is the predecessor to SmartGridSolve. SmartNetSolve allowed a group of tasks to be collectively mapped and collectively executed on a fully connected network. The initial design allowed the user to give a description of the group of tasks and then at run-time this description would be used to generate a task graph. This task graph and a graph of the network were used to generate a mapping solution which was then used to execute the tasks on the fully connected network. Initially, the description of the task graph was given using an XML file which was read at run-time. A new language, Application Definition Language (ADL) [41], was also designed to make this more user friendly.

In May 2006, a non-intrusive and incremental approach for enabling direct communication in NetSolve was introduced [42]. The software component was non-intrusive which meant that the software component is supplementary, working on top of the original system and therefore the original code base was not altered. It was incrementental which meant that the software component does not have to be installed on all computers to enable applications with the new feature.
The limitations to this approach are:

- Tasks are mapped individually.
- Push communication cannot be implemented when tasks are mapped individually.
- Increased communication times since communication cannot be done in parallel with computation or other communication.
- The user has to explicitly specify dependencies.
  - More labour intensive.
  - More prone to error.

In September 2006, distributed task sequencing was developed for the GridRPC model [43]. A new function was introduced that allows direct data transfer between servers when executing a task sequencing job in a Grid environment. This meant that multiple servers could be used and not just a single server as was originally a restriction of task sequencing.
The limitations of this approach are:

- Tasks were not mapped collectively.
- Conditional statements cannot exist in the scope of the task sequencing job.
- Client computation cannot exist in the scope of the task sequencing job.
- Push communication is not implemented.

In October 2006, the special agent called MADAG was implemented in DIET which handled workflow submissions [44]. The user gives a description of this workflow using an XML file including the values of any arguments (i.e. element values of vectors, matrices, etc). Using the DIET API, the user references the file which has the DAG description. This is used to create a DAG or task graph which is submitted to the MADAG agent which is responsible for scheduling the DAG. This is implemented for the DIET API and has not yet been implemented for the GridRPC model and API. This implementation does not follow the RPC style of calling each task in the application. Instead, the application calls a function that submits the entire task graph as a single entity.
The limitations to this approach are:

- Since this approach does not follow the GridRPC model, intermediate results cannot be sent back to the client.
- The task graph has to be known at compile time. Therefore no conditional statements can exist and initial values of input matrices, vectors etc. have to be known before run-time.
- Client computation cannot exist between tasks.
- It is not user friendly as it can be difficult and time consuming to write the XML description of a task graph.
- Writing XML files to generate task graphs is more prone to error than if the task graphs were automatically generated by the system.

Since this initial design of MADAG system, a GUI has been developed which has made it more user friendly [45]. However this is also prone to error due to the fact that the application programmer has to outline the task graph and in addition it is still more labour intensive than if the task graph was automatically generated. In late 2006, work began on SmartGridSolve and the SmartGridRPC model to address the limitations described above. This was presented in 2008 [25].

Chapter 4

GridRPC programming model and API

The aim of the GridRPC model is to provide a standardized, portable and simple programming interface for remote procedure call (figure 4.1). It intends to unify client access to existing Grid computing systems (such as GridSolve, Ninf-G and DIET).

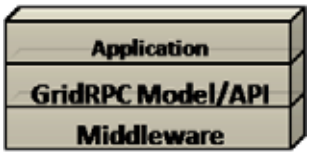

Figure 4.1. Overview GridRPC model and API

This standardisation provides portability of the programmers' source code across all GridRPC implemented platforms. Since the GridRPC model specifies the API and the programming model but does not dictate the implementation details of the servers which will execute the remote procedure call, there may be multiple different middleware implementations of the GridRPC model on which the source code could be executed.

4.1 Design of the GridRPC programming model

The functions presented in this section are shared by all the implementations of the GridRPC model. However the mechanics of these functions differ in each implementation.

Registry discovery: The servers of the Grid environment register the tasks which they can execute with a *registry*. This involves sending information such as how the client should interface with the task and what type of arguments the server expects when the task is called (the calling sequence). In this thesis the registry will be an abstract term for the entity/entities which store the information about the registered tasks and the underlying network. This may be a single entity such as the Agent in GridSolve or several entities such as the MDS [46] (or LDIF [47]), running on servers in Ninf-G or the Global Agents and Local Agents running in the DIET system.

Run-time of client application: When the GridRPC call is invoked, the client contacts the registry to look-up a desired task and receives a handle which is used by the client to interface with the remote task. A task handle is a small structure that describes various aspects of the task and its arguments such as:

- The task name.
- The object types of the arguments (scalars, vectors, matrices etc.).
- The data type of the arguments (integer, float, double, complex etc).
- Whether the arguments are inputs or outputs.

The client then uses the handle to call the task which eventually returns the results. Each GridRPC call gets processed individually, where each task is discovered (task look-up) and executed separately from all the other tasks in the application.

Currently a task is discovered by explicitly asking the registry for a known function through a string look-up. For applications which are run using the GridSolve middleware, the discovery mechanism is done via the GridSolve agent. In Ninf -G, discovery is done via the Globus MDS which runs on each server and

in DIET discovery is done via the Global Agent. The GridRPC model does not dictate the mechanics of resource discovery since different underlying GridRPC implementations may use vastly different protocols.

GridSolve and DIET are GridRPC systems that can perform dynamic mapping of tasks. Discovery for dynamic mapping also involves discovery of performance models which are used by the mapping heuristics. The performance models for DIET are the FAST prediction tool [24], CORI [24] and NWS [48]. The performance models for GridSolve are described in section 6.3.

4.2 GridRPC : API and semantics

We introduce the fundamental objects and functions of the GridRPC API and explain their syntax and semantics.

There are two fundamental objects in the GridRPC model, the task handles and the session IDs. The task handle represents a mapping from a task name to an instance of that task on a particular server.

Once a particular task-to-server mapping has been established by initializing a task handle, all GridRPC calls using that task handle will be executed on the server specified in that binding. In GridRPC systems which perform dynamic resource discovery and mapping, it is possible to delay the selection of the server until the task is called. In this case, resource discovery and mapping is done when the GridRPC task call is invoked with this initialized handle. In theory there is more chance to choose a "better" server in this way, since at the time of invocation more information regarding the task and network is known, such as the size of input/outputs, complexity of task and dynamic performance of client-server links.

There are two types of GridRPC task call functions: blocking and non-blocking calls. The *grpc_call()* function makes a blocking remote procedure call with a variable number of arguments. This means the function does not return until the task has completed and the client has received all outputs from the server.

The *grpc_call_async()* function makes a non-blocking remote procedure call with a variable number of arguments. When this call is invoked, the remote task

and data transfer of the input is initiated and the function returns. This means that either the client computation or server computation can be done in parallel with the *grpc_call_async()* call.

The *grpc_wait()* function waits for the result of the asynchronous call with the supplied session ID. The *grpc_wait_all()* function waits for all preceding asynchronous calls.

4.3 GridRPC : A GridRPC application

Table 4.1 is a simple application which uses the GridRPC API. It comprises of three task handles and three corresponding remote calls. The task handles are set up so that the remote call is bound to a server at call time by passing "bind_server_at_call_time"[1] as a parameter. This string could be substituted with a server host name or the user could assign it to the default server by calling *grpc_function_handle_default()*.

The task "mmul" takes four arguments, the size of the matrices, the two input matrices and the one output matrix. In this application the size of the matrices are not known prior to run-time as they can only be established by executing the local functions (*initMatA* and *initMatB*). Therefore, it is impossible for a user to decide which servers to assign which tasks since the size of inputs and outputs and complexity is not known until the application is run. This is a difficult decision even if the size of the matrices are known before run-time as the performance of underlying networks is dynamic and difficult to predict in Grid environments.

It is also impossible for a dynamic GridRPC system such as GridSolve, which can discover resources and map tasks at run-time, to optimally map the tasks in this application. This is due to the current GridRPC model only permitting a single task to be processed at any time. Therefore, when the system maps the GridRPC task call executing handle *h1* it has no knowledge of what tasks are executing in parallel with this task and the computation load of the tasks executing in parallel.

[1] *This special string is a GridSolve-specific workaround to enable lazy binding in GridRPC*

Table 4.1. Hydropad evolve loop

```
main()
{
    int N;
    int M;
    double A[N*N], B[N*N], C[N*N];
    double D[M*M], E[M*M], F[M*M], G[M*M];

    grpc_function_handle_t h1, h2, h3;
    grpc_session_t s1, s2;
    grpc_initialize(argv[1]);

    /* initialize */
    char * hndl_str= "bind_server_at_call_time";

    grpc_function_handle_init(&h1, hndl_str,"mmul/mmul");
    grpc_function_handle_init(&h2, hndl_str, "mmul/mmul");
    grpc_function_handle_init(&h3, hndl_str, "mmul/mmul");

    N=getNSize();
    initMatA(N, A);  initMatB(N, B);
    if(grpc_call_asnc(&h1,&s1, N, A, B, C)!= GRPC_NO_ERROR) {
        fprintf(stderr, "Error in grpc_call\n");
        exit(1);
    }

    M=getNSize();
    initMatD(M, D);  initMatD(M, E);
    if(grpc_call_async(&h2, &s2, M, D, E, F)!=GRPC_NO_ERROR){
        fprintf(stderr, "Error in grpc_call\n");
        exit(1);
    }
    grpc_wait(s1);
    grpc_wait(s2);

    if (grpc_call(&h3, M, C , F, G) != GRPC_NO_ERROR) {
        fprintf(stderr, "Error in grpc_call");
        exit(1);
    }
    grpc_function_handle_destruct(&h1);
    grpc_function_handle_destruct(&h2);
    grpc_function_handle_destruct(&h3);
    ...
    grpc_finalize();
}
```

Consider the following scenario - M is initialized to 1000 and N is initialized to 100. Therefore the computational load of the first task will be far less than that of the second task. In this circumstance when the system maps the function handle *h1* it will map this to the fastest server as this will yield the lowest execution time for this task. Then when the system maps the function executing handle *h2* it will map it to the second fastest server as the fastest server is currently heavily loaded with the first task. This is poor load balancing of computation and will affect the overall performance of both tasks executing in parallel.

In addition since tasks are processed individually in the GridRPC model, it

is impossible for systems which implement this model to know the dependencies between tasks. Since dependencies between tasks are not known and the communication model of the GridRPC model is based on the client-server model, bridge communication between remote tasks is forced. With the GridRPC model this dependent argument would have to be sent from the source task to the destination task via the client which is two communication steps. This necessity for the client to buffer intermediate data may also cause memory paging on the client. In this application, the third task *h3* is dependent on argument F from the second task *h2* and argument C from task *h1*. In this case, the only way to send F from the server executing *h2* and C from the server executing *h1* to the server executing *h3* is via the client, which is two communication steps. Mapping tasks individually in this application has forced bridge communication. Since intermediate results are sent back to the client, this will also increase the amount of memory used on the client. As a result, the overall volume of communication will increase and there may be paging on the client, which would significantly affect the performance of the application. In addition, since tasks are mapped individually on to a star network, parallelism of remote communication cannot be employed. In this case if dependencies were known, argument C could be sent from the server executing *h1* to the server executing *h3* in parallel with computation and communication of task *h2* (permitting that task *h2* has been assigned a different server than *h3*).

From this application it is evident that the potential for higher performance applications could be increased if we can map tasks collectively as a group on to a network which is fully connected. This is the premise of the SmartGridRPC model.

Chapter 5

SmartGridRPC programming model and API

The aim of the SmartGridRPC model is to enhance the GridRPC model by providing functionality for collective mapping of a group of tasks on a fully connected network.

The SmartGridRPC programming model is designed so that when it is implemented it is interoperable with the existing GridRPC implementation (figure 5.1). Therefore, if any middleware has been extended to be made SmartGridRPC compliant, the application programmer has the option whether their application is implemented for the SmartGridRPC model, where tasks are mapped collectively on to a fully connected network or for the standard GridRPC model, where tasks are mapped individually on to a client-server star network.

In addition, the SmartGridRPC model is designed so that when it is implemented it is incremental to the GridRPC system. Therefore, if the SmartGridRPC model is installed only on the client side, the system will be extended to allow for collective mapping. If the SmartGridRPC model is installed on the client side and on only some of the servers in the network, the system will be extended to allow for collective mapping on a partially connected network. If it is installed on all servers, the system will be extended to allow for collective mapping on the fully connected network.

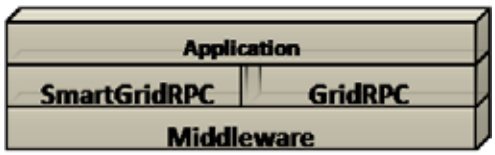

Figure 5.1. Overview SmartGridRPC model and API

5.1 Design of the SmartGridRPC programming model

The SmartGridRPC model provides an API, which allows the application programmer to specify a block of code, in which a group of GridRPC task calls should be mapped collectively. Then, when the application is run, the specified group of tasks in this block of code is processed collectively and each operation in the GridRPC call is separated and done collectively for all tasks in the group. Namely, all tasks in the group are discovered collectively, mapped collectively and executed collectively on the fully connected network. In the discovery phase, performance models are generated for estimating the execution time of the group of tasks on the fully connected network. In the mapping phase, the performance models are used by the mapping heuristic to generate a mapping solution for the group of tasks. In the execution phase, the group of tasks is executed on the fully connected network according to the mapping solution generated.

In the context of this thesis, a performance model is any structure, function, parameter etc., which is used to estimate the execution time of tasks in the distributed environment. The SmartGridRPC performance model refers to performance models, which are used to estimate the time of executing a group of tasks on the fully connected network. The GridRPC performance model refers to performance models, which are used to estimate the execution time of an individual task on a star network. A mapping heuristic is an algorithm, which aims to generate a mapping solution that satisfies a certain criterion, for example, minimum completion time. The SmartGridRPC mapping heuristics refer to mapping heuristics that map a group of tasks on to a fully connected network. The GridRPC mapping heuristics refer to mapping heuristics which map an in-

dividual task on to a client-server network. Furthermore, a mapping solution is a structure, which outlines how tasks should be executed on the distributed network. The SmartGridRPC mapping solution outlines both a task-to-server mapping of each task in the group to a server in the network and the communication operations between the tasks in the group. The GridRPC mapping solution outlines the server list, which specifies where the called task should be executed, and the backup servers which should execute the task should the execution fail.

The job of generating the performance models is divided between the different components of the GridRPC architecture (i.e. client, server and registry). The components may only be capable of constructing part of the performance model required to estimate the groups execution time. Therefore, the registry accumulates these parts from the different components and generates the required performance models.

There are numerous methods for estimating the execution time of the group of tasks on a fully connected network so the implemented performance models are not specified in the SmartGridRPC model. Examples of performance models are the ones currently implemented in SmartGridSolve, which have extended the performance models used in GridSolve (section 6.3). In the future, SmartGridSolve will implement performance models such as the Functional Performance Model, which is described in [49] [50]. Other possible implementations could include the Network Weather Service [51], the MDS directories (Globus, Ninf) [23] and the Historical Trace Manager (GridSolve) [52]. In general, in the SmartGridRPC model, the performance models are used to estimate:

- The execution time of a task on a server.
- The execution time of multiple tasks on a server and the affect the execution of each task has on the other (perturbation).
- The communication time of sending inputs and outputs between client and server.
- The communication time of sending inputs and outputs between different servers.

Mapping heuristics implement a certain methodology that uses these performance models to generate a mapping solution, which satisfies a certain criterion. Examples of mapping solutions include the greedy mapping heuristic and the exhaustive mapping heuristics, which have been currently implemented in SmartGridSolve. There has been extensive research done in the area of mapping heuristics [53] so this is not the focus of our study.

The following sections describe the programming model of SmartGridRPC in the circumstance where the performance models are generated on the registry and the group of tasks is mapped by a mapping heuristic on the registry. However, the SmartGridRPC model could have an alternative implementation. These performance models could be generated on the client and the group of tasks could also be mapped by a mapping heuristic on the client. This may be a more suitable model for systems, such as Ninf-G, which have no central daemon like the GridSolve Agent or the DIET Global Agent.

Registry discovery: The servers provide the part of the performance model, which would facilitate the estimation of the execution time of its available tasks on the underlying network. This partial model can either be automatically generated by the server or has to be explicitly specified or both. This partial model will be referred to as the *server PM*. This part of the performance model is server specific such as the performance of the server and its communication links and a partial model of its available tasks.

As previously mentioned, the SmartGridRPC model does not specify how to implement the *server PM* as there are many possible implementations. Exactly when the *server PM* is sent to the registry is also not specified by the SmartGridRPC model as this would depend on the type of performance model implemented.

But for example, the *server PM* could be sent to the registry upon registration and then updated after a certain event has occurred (i.e. when the CPU load or communication load has changed beyond a certain threshold) or when a certain time interval has elapsed. Or it may be updated during the run-time of the application when actual running times of tasks are used to build the performance model. Suffice it to say that the *server PM* is updated on the registry and is stored there until it is required during the run-time of a client application.

Client application run-time: The client also provides a part of the performance model, which is sent to the registry during the run-time of the client application. This will be referred to as the *client PM*. This part of the performance model is application-specific such as the list of tasks in the group, their order, the dependencies between tasks and the values of the arguments in the calling sequences. In addition, the *client PM* specifies the performance of the client-server links.

In order to determine the part of the performance model that is application-specific, each task that has been requested to be mapped collectively will be iterated through twice. On the first iteration, each GridRPC task call is discovered but not executed. This is the discovery phase. After all tasks in the group are discovered, the client determines the performance of the client-server links and sends the *client PM* to the registry. The registry then generates the performance models based on the stored *server PM* and the *client PM*. Based on these performance models, the mapping heuristic generates a mapping solution. This is the mapping phase. On the second iteration through the group of tasks, each task is then executed according to the mapping solution generated. This is the execution phase. This approach of iterating twice through the group tasks to separate the discovery, mapping and execution of tasks into three distinct phases is the basis that allows the SmartGridRPC model to collectively map and then collectively execute a group of tasks.

The run-time map function, *grpc_map()*, is part of the SmartGridRPC API and allows the application programmer to specify a group of GridRPC calls to map collectively.

This is done by using a set of parenthesis, which follows the map function, to specify a block of code, which consists of the group of GridRPC task calls that should be mapped collectively.

```
grpc_map(char * mapping_heuristic_name){
  ...
  //group of GridRPC calls to map collectively
  ...
}
```

When this function is called, the code and GridRPC task calls within the parenthesis of the function are iterated through twice as previously described.

Discovery phase: On the first iteration through the group of tasks, each GridRPC task call within the parenthesis is discovered but not executed so all tasks in the group can be discovered collectively. This is different to the GridRPC model, which only allows a single task to be discovered at any one time. The client can therefore look up and retrieve handles for all tasks in the group at the same time. In addition to sending the handles, the registry also sends back a list of all the servers that can execute each task. The client then determines the performance of the client-server links to the servers in the list. The client may only determine the performance of some of these links, depending on how many servers are in this list, or may not determine the performance of any of the links if the arguments being sent over the links are small. Exactly how the client determines the performance of these links is not specified by the SmartGridRPC model. This could be implemented using NWS sensors, ping-pong benchmarks, MDS directory or any other conceivable method for determining the performance of communication links.

The client now sends the *client PM* to the registry. The *client PM* specifies the order of tasks in the group, their dependencies and the scalar values of each argument in the calling sequence of each task and the performance of the client-server links. This does not involve sending non-scalar arguments, such as matrices or vectors, but just the pointer value as this will be used to determine the dependencies between tasks. The registry then uses the *server PM* and *client PM* to generate the performance models for estimating the time of executing a group of tasks on the fully connected network. These performance models are then used in the mapping phase to generate a mapping solution.

Mapping Phase: Based on the performance models, the mapping heuristic produces a mapping solution, which satisfies a certain criterion, for example, minimizing the execution time of tasks. The mapping heuristic is specified by the application programmer using the SmartGridRPC API.

There is an extensive number of possible mapping heuristics that could be implemented and therefore the mapping heuristics implemented are not bound by

the SmartGridRPC model. However, the SmartGridRPC framework allows different mapping heuristics to be added and therefore provides an ideal framework for testing and evaluating these mapping heuristics.

Execution Phase: The execution phase occurs on the second iteration through the group of tasks. In this phase, each GridRPC call is executed according to the mapping solution generated by the mapping heuristic on the previous iteration. The mapping solution not only outlines the task-to-server mapping but also the remote communication operations between the tasks in the group.

5.2 SmartGridRPC: API and semantics

The SmartGridRPC API allows a user to specify a group of tasks that should be mapped collectively on a fully connected network. The SmartGridRPC map function is used for specifying the block of code, which consists of the group of GridRPC task calls that is to be mapped collectively. When the *grpc_map()* function is called, the code within its parenthesis will be iterated through twice as previously described in section 5.1. After the first iteration through the group of tasks, the mapping heuristic specified by the parameter *mapping_heuristic_name* of the *grpc_map()* function generates a mapping solution.

The mapping solution outlines a task to server mapping and also the communication operations between tasks. These communication operations include:

- Client-server communication.
  - Standard GridRPC communication.
- Server-server communication.
  - Server sends a single argument to another server.
- Client broadcasting.
  - Client sends a single argument to multiple servers.
- Server broadcasting.
  - Server sends a single argument to multiple servers.

- Server caching.
  - Server stores an argument locally for future tasks.

As a result, the network may have:

- A fully connected topology - where all the servers are SmartGridSolve enabled servers (*SmartServers*), which can communicate directly with each other.
- A partially connected topology - where only some of the servers are *SmartServers*, which can communicate directly. The standard servers can only communicate with each other via the client.
- A star connected topology - where all servers are standard servers and they can only communicate with each other via the client.

During the second iteration through the code, the tasks will be executed according to the generated mapping solution. The SmartGridRPC model also requires a method for identifying code that will be executed on the client. There are many possible approaches, which could be implemented to identify client code. For example, a preprocessor approach could be used to identify the client code transparently. Where the client code cannot be identified, we provide a *grpc_local()* function call, which the application programmer can use to explicitly specify client computation.

```
grpc_map(char * mapping_heuristic_name){

   //reset variables which have been updated
   // during the discovery phase

   grpc_local(list of arguments){
    //code to ignore when generating task graph
   }
   ...
   // group of tasks to map collectively
```

```
    ...
}
```

The *grpc_local()* function is used to specify the code block that should be ignored during the first iteration through the scope of *grpc_map()*. The function is also used to specify remote arguments that are required locally. This information is used to determine when arguments will be sent back to the client and also facilitates the generation of the task graph.

Any segment of client code that is not part of the GridRPC API should be identified using this function. There is one exception to this rule, when the client code directly affects any aspect of the task graph. For example, if a variable is updated on the client that determines which remote tasks get executed or determines the size of inputs/outputs of any task, then the operations on this variable should not be enclosed by the *grpc_local()* function. If any variables or structures are updated during the task discovery cycle then they should be restored to their original values before the execution cycle begins.

5.3 SmartGridRPC: A SmartGridRPC application

Table 5.1 is the SmartGridRPC implementation of the GridRPC application in section 4.3. There is only one extra call required to make this application SmartGridRPC enabled, which is the *grpc_map()* function. In this example, the user has specified that all three tasks should be mapped collectively using the greedy mapping heuristic.

Let us consider the same simple scenario as in section 4.3, where task *h2* has a larger computational load than *h1* and the underlying network consists of two servers, which have different performances. In this case, since all tasks are mapped together, the SmartGridRPC model will improve the load balancing of computation by assigning task *h2* to the faster server and *h1* to the slower.

In addition, task *h3* has a dependency on the argument *F*, which is an output of task *h2*, and argument *C*, which is an output of task *h1*. Since the tasks are mapped as a group and therefore dependencies can be considered, this dependency

Table 5.1. SmartGridRPC model - Example application

```
main()
{
    int N=getNSize();
    int M=getMSize();

    double A[N*N], B[N*N], C[N*N];
    double D[M*M], E[M*M], F[M*M], G[M*M];
    grpc_function_handle_t h1, h2, h3;
    grpc_session_t s1, s2;
    grpc_initialize(argv[1]);

    /* initialize */
    initMatA(N, A);  initMatB(N, B);
    initMatD(M, D);  initMatD(M, E);


    grpc_function_handle_default(&h1, "mmul/mmul");
    grpc_function_handle_default(&h2, "mmul/mmul");
    grpc_function_handle_default(&h3, "mmul/mmul");

    grpc_map("greedy_map"){
      if(grpc_call_asnc(&h1,&s1,N,A,B,C)!= GRPC_NO_ERROR) {
        fprintf(stderr, "Error in grpc_call\n");
        exit(1);
      }
      if(grpc_call_async(&h2, &s2,M,D,E,F)!=GRPC_NO_ERROR){
        fprintf(stderr, "Error in grpc_call\n");
        exit(1);
      }
      grpc_wait(s1);
      grpc_wait(s2);

      if (grpc_call(&h3,M,C ,F,G) != GRPC_NO_ERROR){
        fprintf(stderr, "Error in grpc_call\n");
        exit(1);
      }
    }

    grpc_function_handle_destruct(&h1);
    grpc_function_handle_destruct(&h2);
    grpc_function_handle_destruct(&h3);
    ...
    grpc_finalize();
}
```

can be mapped on to the virtual link connecting the servers executing both tasks, which will reduce the communication load. Or if the tasks are executing on the same server, then the output can be cached and retrieved from the same server, which would further reduce the communication load and further increase the overall performance of the group of tasks.

Also, since no intermediate results are sent back to the client, the amount of memory utilised on the client will be reduced and this will reduce the risk of paging on the client. This prevention of paging could also considerably reduce the overall execution time of the group of tasks.

In addition, since dependencies are known and the network is fully connected, the remote communication of argument C from the server executing task *h1* to the server executing task *h2*, could be done in parallel with the communication and computation of *h2*.

Chapter 6

GridSolve: Implementation of the GridRPC model

The GridSolve agent, which is the focal point of the GridSolve system, has the responsibility of performing discovery and mapping of tasks. The GridSolve agent is an implementation of the registry entity, which was outlined in section 4 of this thesis.

In order to map a task on the client-server network, the agent must discover performance models, which can be used to estimate the execution time of individual tasks on different servers on the network. These performance models include functions for each task, which calculate the computation and communication load of tasks, and parameters, which specify the dynamic performance of the network. These performance models are sent from each server in the network to the agent before run-time of the client application (Agent discovery).

6.1 GridSolve: Agent discovery

This section outlines the GridSolve implementation of the "Registry discover" part of the GridRPC model outlined in section 4.1. The agent maintains a list of all available servers and their registered tasks. This list is incremented when each new server registers with the agent. In addition, the agent stores performance models required to estimate the execution time of its available tasks on the servers. This includes the dynamic performance of each server and functions/parameters,

which are used to calculate the computational and communication load of tasks. These performance models are implemented by executing the LINPACK benchmark [54] on each server when they are started, running the CPU load monitor on the server and using descriptions of the task provided by the person that installed the task to generate functions for calculating the computation and communication load of the tasks (figure 6.1).

The LINPACK benchmark is a measure of a system's floating point computing power. It measures how fast a computer solves a dense N by N system of linear equations $Ax = b$, which is a common task in engineering. The solution is obtained by Gaussian elimination with partial pivoting, with *2/3N3 + 2N2* floating point operations. The result is reported in millions of floating point operations per second (MFLOP/s).

The CPU load monitors determines the CPU load of the server machine via the UNIX *uptime* utility. The CPU load monitor updates the dynamic performance of the server whenever the CPU load significantly changes or when certain time interval elapses [55].

At run-time of the client application, when each GridRPC task call is invoked, these performance models are used to estimate the execution time of the called task on each server.

6.2 Run-time GridRPC task call

In practice, from the user's perspective the mechanism employed by GridSolve makes the GridRPC task call fairly transparent. However, behind the scenes a typical GridRPC task call involves the following operations:

- The discovery phase.
- The mapping phase.
- The execution phase.

The discovery phase: When the GridRPC call is invoked, the client queries the agent for an appropriate server that can execute the desired function. The agent returns a list of available servers, ranked in order of suitability.

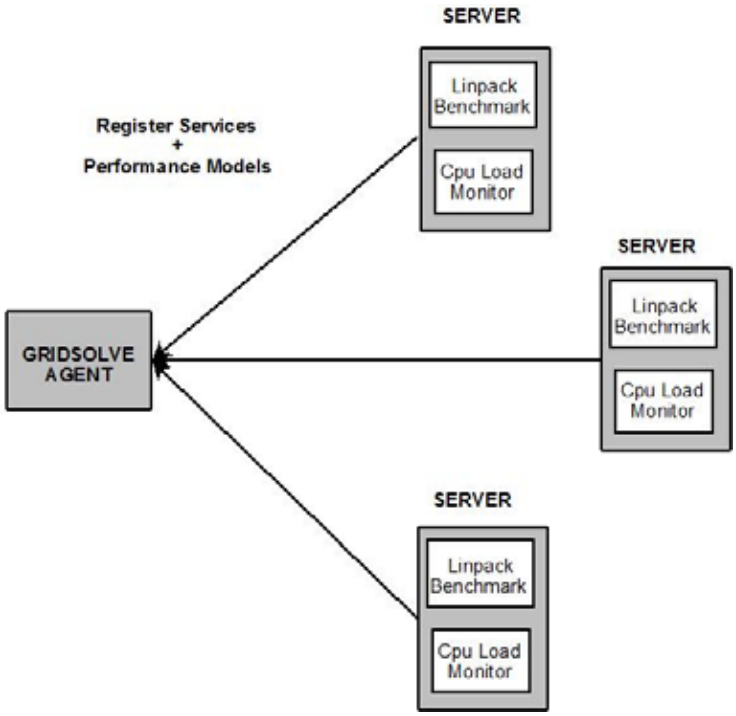

Figure 6.1. GridSolve - Agent discovery

This ranked list is sorted based only on task computation times. Normally, the client would simply submit the service request to the first server on the list, but if specified by the user it is resorted according to its computation and communication time. If this is specified, the bandwidth from the client to the top few servers is measured. This is done using a simple 32KB ping-pong benchmark. The time required to do the measurement will depend on the number of servers, which have the requested task, and the bandwidth and latency from the client to those servers. When the data is relatively small, the measurements are not performed because it would take less time to just send the data than it would take to do the measurements. Also, since a given service may be available on many servers, the cost of measuring network speed to all of them could be prohibitive. Therefore, the number of servers to be measured is limited to those with the highest computational performance.

The mapping phase: As previously described, the agent sends a server list, which is ordered according to their estimated computation time.

In GridSolve, there is a number of mapping heuristics, which can be employed to generate the mapping solution. Among the mapping heuristics is the minimum completion time (MCT) mapping heuristic, which bases its execution time on the performance models and the dynamic network performance of each server outlined in section 6.3. Also included are a set of mapping heuristics that rely on the other performance model in GridSolve called the Historical Trace Manager (HTM) [**?**].

The execution phase: The client attempts to contact the first server from the list. It sends the input data to the server, the server then executes the task on behalf of the client and returns the results. If at any point the execution fails, the client automatically moves down the list of servers.

6.3 GridSolve: Performance models

The performance models of GridSolve are used by the mapping heuristics to estimate the execution time of individual tasks on a client-server network. In GridSolve, the performance models can be used to estimate:

- The execution time of a task on a server.

- The communication time of sending inputs and outputs between client and server.
- The perturbation that one task has on another.

The mapping heuristics use these performance models to estimate the time of different possible task-to-server mapping solutions and choose the mapping solution, which most satisfies a certain criterion.

The performance models of GridSolve specify the dynamic performance of each server and the dynamic performance of the client-server links. They also specify the computation load and communication load of the called task. When a task is called for execution, the computation load and dynamic performance of a server is used to estimate the task's computation time. The communication load and dynamic performance of the client-server links are used to estimate the task's communication time.

The dynamic performance of a server is parameterized by the number of floating point operations per second (flop/sec) that the server can perform. It is obtained by first determining the static performance of each server by running a sequential benchmark on each server (figure 6.1). This sequential benchmark is the LINPACK benchmark, which is executed on each server when it is started. The benchmark times the execution of a routine, which solves a dense system of linear equations. This benchmark is close to the peak performance rate of the server (P). There is also a "CPU load monitor" on each server, which continually monitors the CPU load (w). When this CPU load changes beyond a certain threshold or if a certain time interval has elapsed (approx. 5mins), then this CPU load is sent to the agent. To get the dynamic performance (p) of the server, the agent uses this updated CPU load to scale the value for the servers peak performance (P). The dynamic performance of a server is calculated as follows:

$$p = \frac{P \times n}{\frac{w}{100} + 1} \tag{6.1}$$

- where p is the dynamic performance of the server, w is the current CPU load, P is the peak performance (benchmark) of the server and n is the number of processors on the server.

The dynamic performance of a client-server link is parameterized by its bandwidth (bw), which is the number of bytes per second (bytes/sec) that can traverse the link. It is obtained using the ping-pong benchmarks. Using these benchmarks the bandwidth of the link can be calculated as follows:

$$bw = \frac{PING\_PACKET\_SIZE}{PING\_TIME} \tag{6.2}$$

- where the PING_PACKET_SIZE is 32 KB and PING_TIME is the time it takes to send the packet between the client and server.

In addition, the performance model includes functions for calculating computation and communication load of tasks. These functions are generated from the task description of each task, which is provided by the person who writes or installs a task. They are written in a language called the GridSolve Interface Definition Language (IDL). With this language, the task writer/installer provides a specification of the calling sequence of the task. This specification describes the data type of each argument (integer, float, double etc.), the object type of each argument (scalar, vector, or matrix) and whether each argument is an input, an output or an input-output. Table 6.1 shows the IDL description of the DGESV task, which is a LAPACK [56] routine that solves:

$$A * X = B \tag{6.3}$$

The IDL description specifies that the first two arguments of the calling sequence are input scalar integers. The third argument of the calling sequence is an input-output matrix of doubles. The fourth argument is an output scalar argument. The fifth argument is an output vector of integers and the sixth is an input-output matrix of doubles. The seventh is a scalar integer input and the eight is a scalar output integer. This specification of the calling sequence is used to generate functions for calculating both the computation load and the communication load of the task.

Included in the IDL description of the task is a "string" formula that is used in conjunction with the specification of the calling sequence to generate a function for calculating the computation load of a task. This formula is denoted in the

Table 6.1. IDL Description of DGESV task

```
SUBROUTINE dgesv(
   IN int N,
   IN int NRHS,
   INOUT double A[LDA][N],
   INOUT int LDA,
   OUT int IPIV[N],
   INOUT double B[LDB][NRHS],
   IN int LDB,
   OUT int INFO)


"This solves Ax=b using LAPACK"
LANGUAGE="FORTRAN"
LIBS="$(LAPACK_LIBS) $(BLAS_LIBS)"
COMPLEXITY= "2.0*pow(N,3.0)*(double)NRHS"
MAJOR= "COLUMN"
```

IDL file as the "COMPLEXITY" parameter. The string formula for the DGESV task is:

$$dgesv_{flop} = 2 * N^3 * NRHS \tag{6.4}$$

This formula in conjunction with the specified calling sequence in table 6.1 generates a function that describes the computational load as a multiplication of 2 by the first argument cubed multiplied by the second argument of the calling sequence.

At run-time, when the values of these arguments are known, this function can be used to calculate the computation load of task. The computation load is measured in the number of floating operations (flop), which the task will execute. An example of the calling sequence of the DGESV task could be as follows:

$$grpc\_call(\&degesv\_handle, 400, 100, A, 800, IPIV, B, 400, INFO) \tag{6.5}$$

For this calling sequence the computation load would be:

$$dgesv_{flop} = 2 * 400^3 * 100 = 12800 * 10^6 Flop = 12800 MFlop \tag{6.6}$$

The communication load of a task can be calculated using the following formulas in conjunction with the specified calling sequence in the IDL specification:

$$arg\_size(matrix) = rows * cols * get\_elem\_size(arch, DATA\_TYPE) \quad (6.7)$$

$$arg\_size(vector) = rows * get\_elem\_size(arch, DATA\_TYPE) \quad (6.8)$$

In these formulas, the DATA_TYPE variable specifies whether the argument type is a double, integer, float etc. And the rows and cols variables are the dimensions of the matrix/vector. The *get_elem_size()* function returns the size of bytes of the specified DATA_TYPE (double, integer etc.).

The formula for calculating matrix argument size in conjunction with the specification of the calling sequence in table 6.1 would generate a function that outlines that the communication load of argument A of the DGESV task can be calculated by multiplying the fourth argument (LDA) in the calling sequence by the first argument (N) in the calling sequence by the size of a double (e.g. 8 bytes). At run-time, when the values of these arguments are known, this function can be used to calculate the communication load of argument A of the DGESV task.

However, in some instances, the platform of the sending machine must be known to determine this communication load. One of the problems with C and C++ is that the built in data types such as *int* and *long int* are platform dependent. There is nothing in the standard to say how many bytes each data type occupies beyond some basic ordering. For example, *long int* must use at least as many bytes as *int* (but could be the same). Table 6.2 outlines the number of bytes of different data types on different platforms.

For this reason the *get_elem_size()* function also takes an architecture identifier as a parameter. For the calling sequence outlined in equation 6.5 for the DGESV task, the communication load of non-scalar arguments on a 32bit Intel machine running LINUX OS would be:

$$A_{bytes} = LDA * N * get\_elem\_size(LINUX\_86, DOUBLE) \quad (6.9)$$

Table 6.2. Size of datatypes on different platforms

OS	Arch	Size of int	Size of long int	Size of double
LINUX	x86	4 bytes	4 bytes	8 bytes
LINUX	x86-64	4 bytes	8 bytes	8 bytes
Windows	x86	4 bytes	4 bytes	8 bytes
Windows	x86-64	4 bytes	4 bytes	8 bytes
MAC OS X	x86	4 bytes	4 bytes	8 bytes
MAC OS X	x86-64	4 bytes	8 bytes	8 bytes

$$A_{bytes} = 800 * 400 * 8 = 256 * 10^4 Bytes = 2.44 MBytes \tag{6.10}$$

$$IPIV_{bytes} = N * get\_elem\_size(LINUX\_86, INTEGER) \tag{6.11}$$

$$IPIV_{bytes} = 400 * 4 Bytes = 1600 Bytes = 0.00153 MBytes \tag{6.12}$$

$$B_{bytes} = LDB * NRHS * get\_elem\_size(LINUX\_86, DOUBLE) \tag{6.13}$$

$$B_{bytes} = 400 * 100 * 8 = 320000 Bytes = 0.30518 MBytes \tag{6.14}$$

From these performance models, it is possible to estimate both the communication time and computation time of individual tasks on the client-server network. These performance models are used by the mapping heuristics to generate a mapping solution.

6.4 GridSolve: Mapping heuristic

There have been several mapping heuristics implemented in GridSolve. Each task is mapped when it is called for execution and therefore each task is mapped individually on the client-server network. The following mapping heuristics have been implemented:

- Minimum Completion Time.

- HTM - Minimum Completion Time.
- HTM - Minimum Perturbation.
- HTM - Minimum Sum Flow.
- HTM - Minimum Length.

The Minimum Completion Time (MCT) maps the individual task based on the performance models described in 6.3.

All the HTM mapping heuristics generate mapping solutions based on the Historical Trace Manager (HTM) performance model. When a new task arrives, the HTM simulates the execution of the task on each server. Using the HTM information, the heuristic has an estimation of the finishing time of each task running on each server. This is used to consider the perturbation that tasks induce on each other and compute the best server according to the main objective of that heuristic.

When a task has completed the server sends a message to the agent that the task has completed and this information is used by the HTM to correct what has been simulated and improve the quality of future predictions.

Chapter 7

SmartGridSolve: Implementation of the SmartGridRPC model

7.1 SmartGridSolve: Agent discovery

This section presents the SmartGridSolve implementation of the "registry discovery" part of the SmartGridRPC model outlined in section 5.1. In addition to registering services, the servers also send the *server PM*. The *server PM* makes up part of the performance model used for estimating the execution time of the server's available tasks on the fully connected network. This along with the *client PM* is used to generate a performance model, which is used by the mapping heuristics to produce mapping solutions.

Currently, the *server PM* of SmartGridSolve extends that of GridSolve, which comprises of functions for calculating the computation load and communication load and parameters for calculating the dynamic performance of the servers and client-server links. This is described in section 6.3.

However, the network discovery of GridSolve is extended to also discover the dynamic performance of each link connecting *SmartServers*. These are those servers, which can communicate directly with each other and store/receive data in their local cache. The dynamic performance of the server-server links are taken periodically using the same 32KB ping-pong technique used by GridSolve (figure 7.1).

To achieve backward compatibility and to give the server administrators full control over how the server operates, a *SmartServer* may be also started as a standard GridSolve server.

As a result, the network may have:

- A fully connected topology.
- A partially connected topology.
- A star connected topology.

Also to minimize the volume of data transferred around the network, each *SmartServer* is given an ID. Each *SmartServer* only sends ping-pong messages to those *SmartServers* that have an id that is less than their own. This prevents the performance of the same communication link being measured twice. Once determined, these values are sent to the agent to update the *server PM*. The *server PM* is stored on the registry and updated either periodically (every 5 minutes) or when the CPU load monitor records a change, which exceeds a certain threshold. This *server PM* is then used to generate the performance models during the run-time of a client application.

7.2 Run-time of client application

This section presents the SmartGridSolve implementation of the "Client application run-time" part of SmartGridRPC model outlined in section 5.1. Each phase of the SmartGridRPC run-time map function (*grpc_map()*) will be described.

Discovery phase: On the first iteration through the group of tasks, each GridRPC task call (*grpc_call()*) within the parenthesis is discovered but not executed. This involves discovering the name of each task and the calling sequence of each task, which involves discovering the pointers to the non-scalar arguments (such as matrices, vectors etc.) and the values of the scalar arguments.

After the first iteration through the group, the client contacts the agent and looks up the group of tasks, which involves sending the agent a list of the task names.

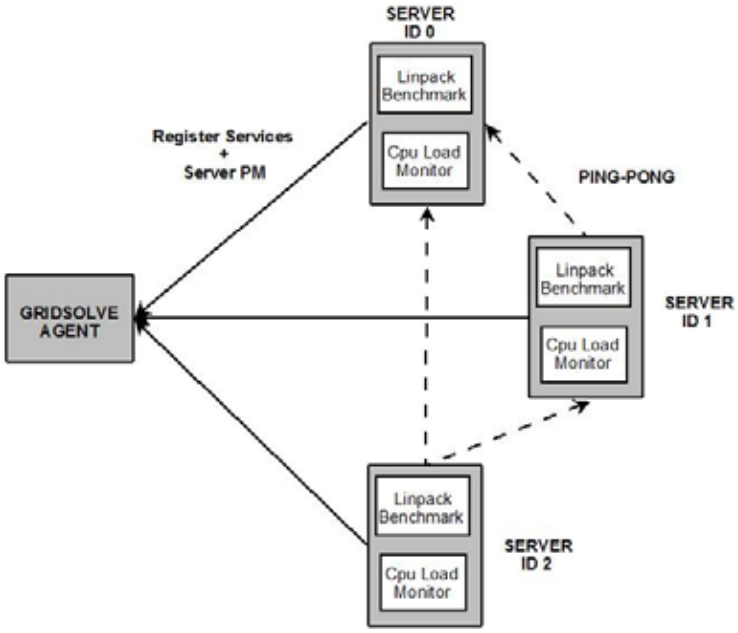

Figure 7.1. SmartGridSolve - Agent discovery

The agent then creates a handle for each instance of a task. The agent sends back the group of handles, one for each task. In addition, for each handle it sends a list of servers, which can execute each task.

The client then uses the list of servers to perform the ping-pong benchmark on each of the links from the client to each server that can execute a task in the group of tasks (figure 7.1). Subsequent to this, the client sends the *client PM*, which is a structure that specifies application-specific information such as the list of tasks, the calling sequence and the dependencies between the tasks. In addition it specifies the performance of each client-server link.

The agent can now generate all the performance models necessary for estimating the execution time of the group of tasks on the fully or partially connected network. In SmartGridSolve, these performance models consist of a task graph, a network graph and functions for estimating computation and communication times.

The task graph specifies the order of tasks, their synchronisation (whether they are executed in sequence or parallel), the dependencies between tasks, the load of computation and communication of each task in the group.

The network graph specifies the performance of each server in the network and the communication links of the fully connected, partially connected or star network. These performance models will be used by the mapping heuristics in the mapping phase to generate a mapping solution for the group of tasks.

Mapping Phase: The mapping heuristic produces a mapping solution graph based on the task graph, the network graph and the functions for estimating computation and communication time. The mapping heuristics currently implemented in SmartGridSolve are:

- Exhaustive mapping heuristic.
- Random walk mapping heuristic.
- Greedy mapping heuristic

The mapping solution generated by these heuristics is then used in the execution phase to determine how the group of tasks should be executed on the network. Mapping heuristics are not the focus of this thesis as there has been extensive research done in this area. However, the above heuristics have been implemented in SmartGridSolve and it is possible to choose which mapping heuristic you want to use to generate a mapping solution for a given application. It is also possible to use other existing or your own mapping heuristic by implementing them into in the SmartGridSolve system.

Execution Phase: This execution phase occurs on the second iteration through the group of tasks. In this phase, each GridRPC call is executed according to the mapping solution generated by the mapping heuristic. The mapping solution not only specifies the task-to-server mapping but also the communication operations between the tasks in the group. In addition to the standard GridRPC communication, the mapping solution can use the following communication operations:

- Server-server communication.
- Client broadcasting.
- Server broadcasting.
- Server caching

7.3 SmartGridSolve performance models

This section presents the performance models, which are currently implemented in SmartGridSolve. The performance models are used by the mapping heuristics to estimate the execution times of different mappings of the group of tasks on the network. This involves both estimating the computation time of tasks of the application on the servers of the network and also estimating the communication time of sending inputs and outputs over the network. The accuracy of these performance models affects the ability of the mapping heuristics to generate optimal mapping solutions.

7.3.1 Network graph

The network graph is a representation of the performance of the servers and communication links of the fully connected, partially connected or star network. If SmartGridSolve is installed only on the client side, this structure will represent a star network where no servers can communicate directly. With this network topology, the application programmer may only benefit from improved mapping of tasks to servers. If some of the servers are *SmartServers* then this structure will represent a partially connected network. With this network topology, the application programmer may also benefit from improved mapping of communication. If all servers are *SmartServers*, the network will be fully connected. With this network topology, the application programmer will benefit from the full potential of improved mapping of communication.

The graph specifies the performance of each server and also the performance of each link connecting it with the client. Where there are two or more servers in the network that are *SmartServers*, the graph will include links which specify the performance of the link between these servers.

Figure 7.2 illustrates a network graph, which represents three *SmartServers* and one standard server. Each circle node in the graph represents a server and is weighted by its dynamic performance, which is measured in floating point operations per second (flop/sec). The single diamond shaped node represents the client. Each link connecting nodes represents a "virtual communication link" and is weighted by its dynamic performance (bandwidth), which is measured in the

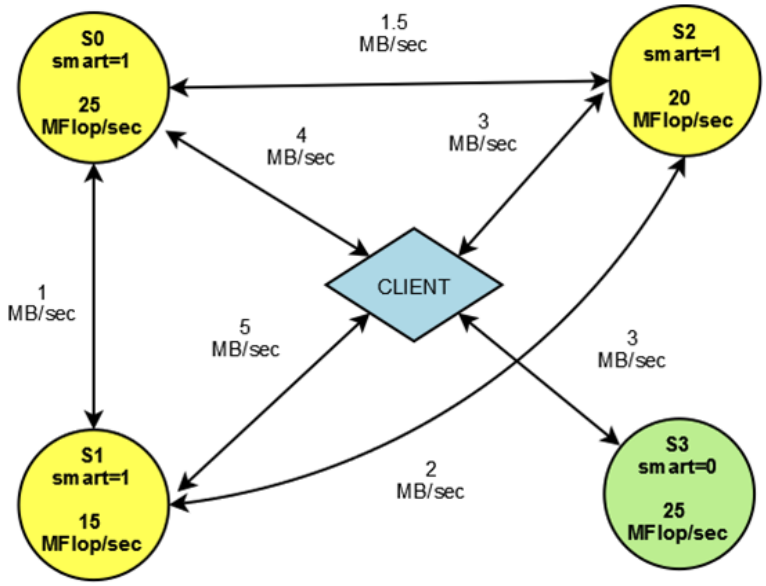

Figure 7.2. SmartGridSolve - The network graph

number of bytes per second (bytes/sec), which can traverse the link. All four servers have links connecting them to the client but only *SmartServers* have links connecting them to other *SmartServers*. In figure 7.2, the servers S0, S1 and S2 are *SmartServers* and therefore have links connecting them with each other and also to the client. Server S3 is a server, which was started without direct communication enabled. If GridSolve has been compiled with the SmartGridSolve extension, a server administrator has the option whether the server is started with direct communication enabled or disabled.

The performance of the servers and communication links are calculated using the equations outlined in section 6.3.

7.3.2 Task graph

The task graph is the representation of the mapped group of tasks. The task graph specifies the order of tasks, their synchronisation (whether they are executed in sequence or parallel), the dependencies between tasks, the load of computation and communication of each task in the group.

Figure 7.3 illustrates a task graph, which represents 5 tasks, where task 0 and

task 1 are executed in sequence and then task 2, task 3 and task 4 are executed in parallel. The task graph has three sets of nodes, the task nodes, the client node and the argument nodes. Each task node is represented by a rectangle node and is weighted by its computation load (flop). Each input and output non-scalar argument (matrix, vector etc.) is represented by a circle shaped node and is weighted by its communication load (bytes).

The functions for calculating the computation load of each task in the group are generated using the formulas specified by the person that wrote or installed the task in conjunction with the specification of the tasks calling sequence in the IDL description (table 6.1). This is part of the *server PM*, which is sent from each server to the agent prior to the execution of the client application. Then at run-time, the calling sequence of each task in the group is discovered collectively and these calling sequences are sent as part of the *client PM* to the agent. The functions of the *server PM* and the calling sequences of the *client PM* can be used to determine the computation load of each task in the group.

The functions for calculating the communication load of each non-scalar argument in the group are generated using the functions for calculating argument sizes in section 6.3 in conjunction with information on the tasks calling sequence in the IDL description (table 6.1). This is part of the *server PM*. At run-time, the calling sequences of the tasks in the group are discovered collectively and are sent to the agent as part of the *client PM*. Then, the communication load functions of the *server PM* and the calling sequence of the *client PM* are used to determine the communication load of each non-scalar argument in the group.

The dependencies between tasks are determined by examining the pointers of non-scalar arguments of the calling sequence of each task (which is specified in the *client PM*) and using the IDL description (which is specified in the *server PM*) to determine whether they are inputs or outputs. The links in the graph represent the data-flow between tasks. There are two types of data-flow dependencies, the input data-flow dependency and output data-flow dependency.

Input dependencies occur when a task has a dependency on an input of another task. This is specified in the task graph by a link from an input argument node of one task pointing to another task node. If multiple tasks require the same input argument then a link will emanate from this argument node to each dependent

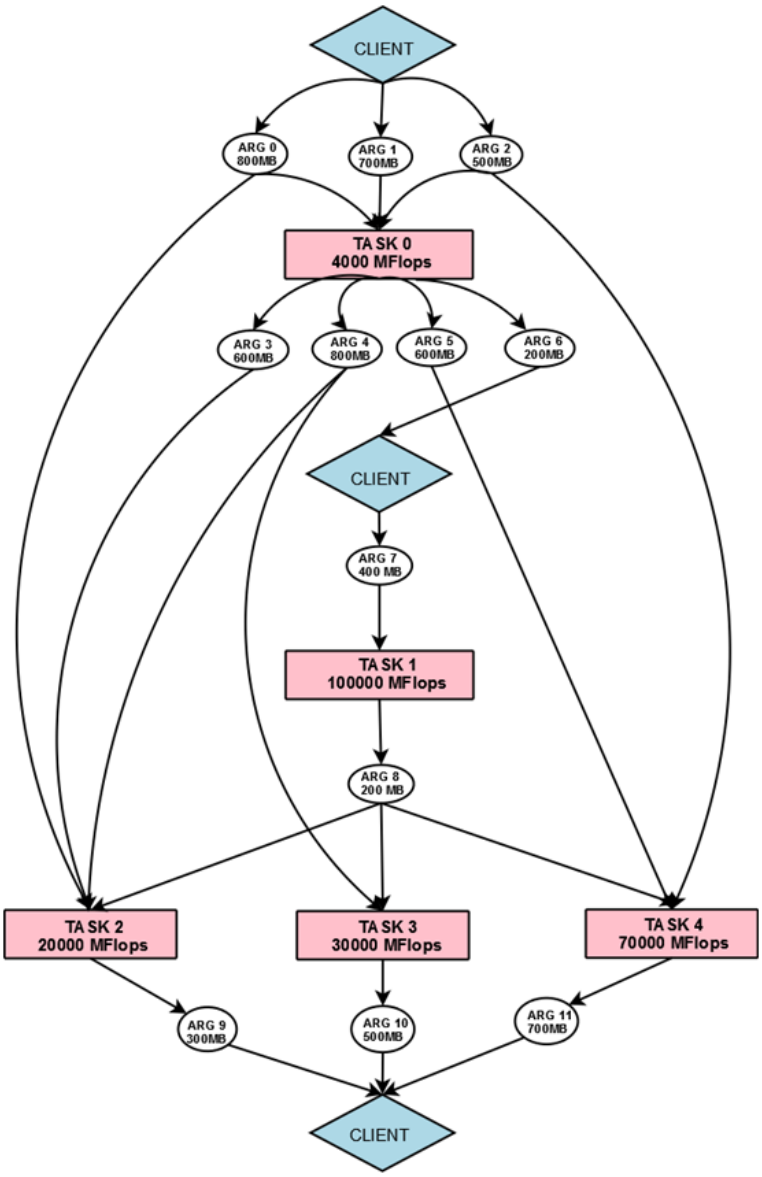

Figure 7.3. SmartGridSolve - The task graph

task node. In this case the mapping heuristics can choose a mapping solution, which broadcasts the input argument of the source task from the client to each of the servers of the destination tasks.

Output dependencies occur when a task has a dependency on an output of another task. This is specified in the task graph by a link from an output argument node of one task pointing to another task node. If multiple tasks require the same output argument, then a link will emanate from this argument to each dependent task. In this case, mapping heuristics can choose a mapping solution, which broadcasts the argument from the source task to each of the servers of the destination tasks.

7.4 Mapping Solution Graph

The mapping heuristics of SmartGridSolve generate a mapping solution based on the task graph and the performance model of the fully connected network. A mapping solution graph is a structure, which outlines both the task-to-server mapping of the group of tasks and the communication scheme between the tasks in the group. In addition, the mapping solution graph outlines the estimated computation time of each task on their assigned server and the estimated communication time of each task dependency on their assigned communication path.

Figure 7.4 shows an example of mapping solution graph that could be generated based on the task graph in figure 7.3 and the network graph in figure 7.2.

This communication scheme in this mapping solution implements each type of communication transaction, which can be employed in the SmartGridRPC model:

- Direct server-server communication.
- Client broadcasting.
- Server broadcasting.
- Server caching of inputs.
- Server caching of outputs.

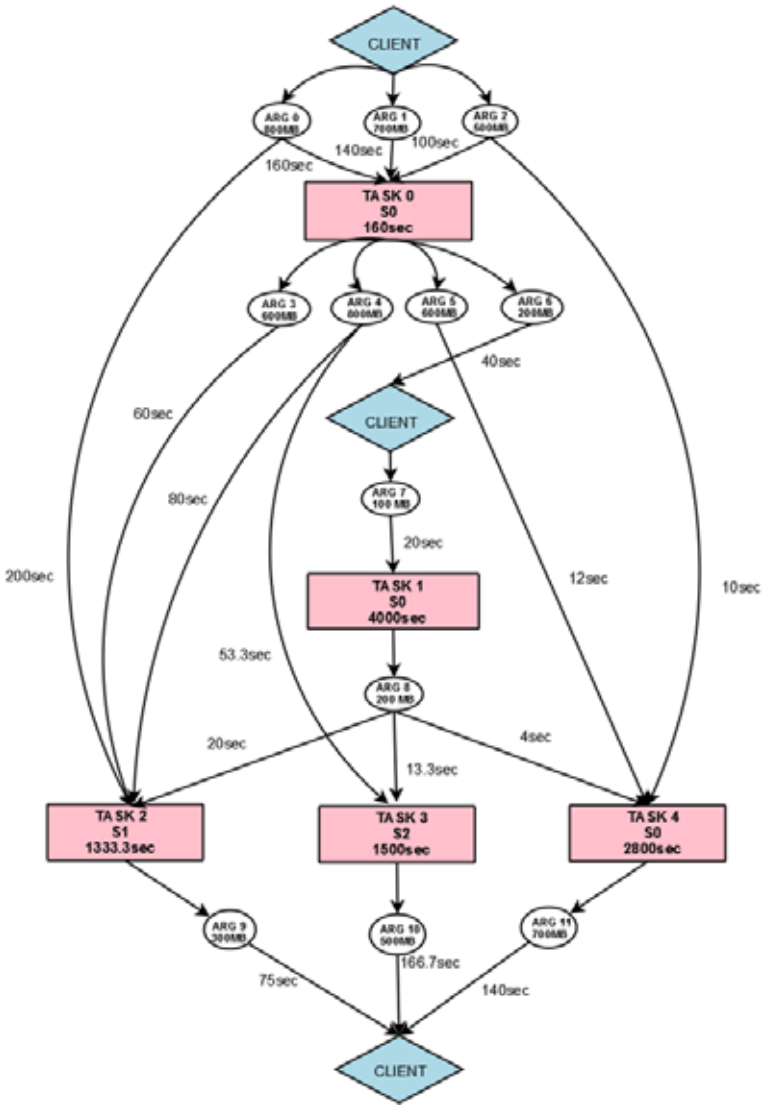

Figure 7.4. SmartGridSolve - The mapping solution graph

The mapping solution outlines **direct server-server communication** of *argument 3* from *server 0* to *server 1* after *task 0* has executed. This argument is subsequently used on *server 1* for the execution of *task 2*. It outlines **server broadcast communication** of *argument 4* from *server 0* to *server 1* and *server 2* after *task 0* has executed. This argument is subsequently used on *server 1* for execution of *task 2* and on *server 2* for the execution of *task 3*. It outlines **client broadcast communication** of *argument 0* from the client to *server 0* and *server 1* before the execution of *task 0*. It outlines the **server caching of input** *argument 2* on *server 0* before the execution of *task 0*. This argument is subsequently used on the same server by *task 4*. It outlines the **server caching of output** *argument 5* on *server 0* after the execution of *task 0*. This argument is subsequently used on the same server by *task 4*.

The estimated time of each of these remote communication transactions is calculated by dividing the communication load of the argument outlined in the task graph in figure 7.3 and the bandwidth of the communication link outlined in the network graph in figure 7.2. For example the direct server-server communication of *argument 3* is estimated to take 60 seconds, which is calculated by dividing the communication load of 600MB by the link speed, which is 10MB/sec.

The estimated time of the caching transactions are based on a naïve assumption that the disk speed is 50MB/sec. For example, the caching of input *argument 2* takes 10 seconds, which is calculated by dividing the argument size of this argument which is outlined in the task graph which is 500MB by the disk speed which is 50MB/sec. In the future, benchmarks could be used to determine a more accurate disk speed for each machine.

The estimated time for computation is calculated by dividing the computation load of the task, outlined in the task graph, and the server performance speed, outlined in the network graph. For example, the computation time for *task 0* is 160 seconds, which is calculated by dividing the computation load of 4000MFlops by the server speed, which is 25MFlops/second.

However, not every task in the group contributes to the overall execution time of the group of tasks. Parallelism of computation has been employed between *tasks 2*, *task 3* and *task 4* and therefore only the task that takes the longest time

of all three contributes to the total execution time of all three. In this mapping solution, the time saved due to parallelism of computation is:

$$par.comp.time = t(t2, s1) + t(t3, s2) = 1333.3 + 1500 = 2833.3 seconds \quad (7.1)$$

The SmartGridRPC model also permits the parallelism of communication. Any communication transaction may be done in parallel with other computation/communication in the group. This is advantageous when there is a dependency between two tasks and the destination task is not executed in parallel or immediately after the source tasks. In this case, the dependent data can be sent to the destination task in parallel with any computation or communication on any other machine (client and servers), which happens in the intervening time.

For example, each of the communication transactions from *server 0* after the execution of *task 0* can be done in parallel with other computation and communication. This is because the tasks that require the arguments, *task 2*, *task 3* and *task 4* - are not executed in parallel or immediately after *task 0*. Therefore, these communication transactions can be done in parallel with the computation of *task 1* on *server 0* or any computation on the client.

Moreover, broadcast communication from the client can also be done in parallel with other computation/communication that happens in the intervening time. The sending of *argument 0* from the client to *server 1* can be done in parallel with:

- The computation of *task 0* on *server 0*.
- All the communication transactions from *server 0* that happen after *task 0* has executed.
- The computation of *task 1* on *server 0*.
- The broadcast of *argument 6* after the execution of *task 1*.
- Any computation on the client which happens in the intervening time.

Therefore, the time saved due parallelism of communication will be:

$$\begin{aligned} par.comm.time &= t(cl \rightarrow a0 \rightarrow s1) + t(s0 \rightarrow a2 \rightarrow s0) + t(s0 \rightarrow a0 \rightarrow s1) \\ &+ t(s0 \rightarrow a4 \rightarrow s1) + t(s0 \rightarrow a4 \rightarrow s2) + t(s0 \rightarrow a5 \rightarrow s0) \\ &= 200 + 10 + 60 + 80 + 53.3 + 12 \\ &= 415.3 seconds \end{aligned} \tag{7.2}$$

In addition to specifying a more advantageous communication scheme, the mapping solution will outline a more advantageous computation scheme (i.e. task-to-server mapping). Since the mapping heuristics can consider all tasks in the group collectively, it can better distribute the load of parallel computation over the servers. Because the computations of all tasks in the group are considered collectively, the mapping heuristic is able to balance the load of the computation of the three parallel tasks. It therefore assigns *task 4*, which has the highest computation load, to the fastest sever (*server 0*) and *task 2*, which has the lowest computation load, to the second slowest server (*server 1*). If these were mapped individually in the GridRPC model, they could be mapped in reverse order as individual mapping gives priority to tasks in the order they are mapped.

In this example, the amount of time saved by employing parallelism of computation is 2833.3 seconds and the amount of time saved by employing parallelism of communication is 415.3 seconds.

The time saved due to this parallelism of computation and communication does not contribute to the overall group time and therefore would not be included in the calculation for the total execution time for the group.

$$
\begin{aligned}
group.time &= t(cl \rightarrow a0 \rightarrow s0) + t(cl \rightarrow a1 \rightarrow s0) + t(cl \rightarrow a2 \rightarrow s0) \\
&+ t(t0, s0) + t(s0 \rightarrow a6 \rightarrow cl) + t(cl \rightarrow a7 \rightarrow s0) \\
&+ t(t1, s0) + t(s0 \rightarrow a8 \rightarrow sl) + t(s0 \rightarrow a8 \rightarrow s2) \\
&+ t(t4, s0) + t(s1 \rightarrow a9 \rightarrow cl) \\
&= 160 + 140 + 100 + 160 + 40 + 20 + 4000 + 20 + 13.3 + 4 \\
&+ 2800 + 75 + 166.7 + 140 \\
&= 7839 seconds
\end{aligned} \tag{7.3}
$$

This example has outlined that the mapping heuristics of the SmartGridRPC model have more potential of finding a better mapping solution due to collective mapping and employing the SmartGridRPC communication model, which permits parallel remote communication.

7.5 Communication model

The communication model of SmartGridSolve is based on the fully connected network. This extends the GridRPC communication model, so that in addition to client-server communication, the following communication transactions can be employed:

- Server-server communication.
- Server broadcasting.
- Client broadcasting.
- Server caching.

To apply a communication scheme, which employs any of these transactions, the client and the servers must be able to identify where to send and receive arguments of each task. To achieve this functionality, the communication scheme of the group of tasks is stored in communication structures, which specify the communication required for each task in the group. When each task is called for execution, a communication structure is created for that task, which is based on the mapping solution outlined by the mapping heuristic. They are subsequently used by the client to determine where to send the inputs of each task and where to receive the outputs of each task. In addition, the client sends the communication structure to other servers if they are involved in any remote communication. The servers use the structure to determine whether to send their inputs/outputs to remote destinations, to cache them locally or to send them back to the client. If arguments are sent remotely, the structure specifies which servers to send it to and the filename of where arguments should be stored. If the argument is received remotely it specifies the filename where the argument should be read from. These filenames are unique for each argument that is sent remotely.

To demonstrate how the communication model of SmartGridSolve operates using these communication structures, we will consider what communication operations occur if *task 0* of the mapping solution in figure 7.4 was called for execution.

Firstly, the communication transactions, which are initiated before the execution of *task 0*, will be described. These communication transactions are illustrated in figure 7.5. When *task 0* is called for execution, the client generates a communication structure for this task based on the mapping solution. The client then interprets this communication structure, which specifies that *argument 0*, *argument 1* and *argument 2* should be sent to *server 0*. In addition to sending these arguments to the server, the client also sends the communication structure.

The server interprets it and deciphers what to do with the input arguments prior to the execution of the task and the output arguments after the execution of the task. In this case, it outlines that input *argument 2* should be cached locally in a specified file as it is required by *task 4*. This operation is done asynchronously, which means that remote computation/communication (on the client or other server) can be done in parallel with this operation.

The communication structure also outlines that *argument 0* should be sent to *server 1*. This is again done asynchronously and so computation on the client or communication/computation on any other server can be done in parallel with this communication. In addition to sending the argument, the client also sends the communication structure for this argument, which outlines that the argument should be stored locally in the file specified. It should be noted that arguments can also be stored to memory and it is the server administrator's responsibility to choose which method of storage is implemented on the server.

After task 0 has finished executing, the communication structure is used by the server to determine what should be done with the output arguments (figure 7.6).

The communication structure specifies that *argument 5* should be cached locally in a specified file as it is required by *task 4*. In addition, *argument 3* and *argument 4* are sent to *server 1* and *argument 4* is sent to *server 2*. In these transactions, the communication structure is also sent so the destination servers know the files names to store the arguments.

Once again, these remote transactions happen asynchronously and therefore if there is any computation/communication on any machine (client or server), then this will be done in parallel with this communication.

When the destination tasks, which require these remote arguments, are called for execution, the client will send a communication structure outlining that ar-

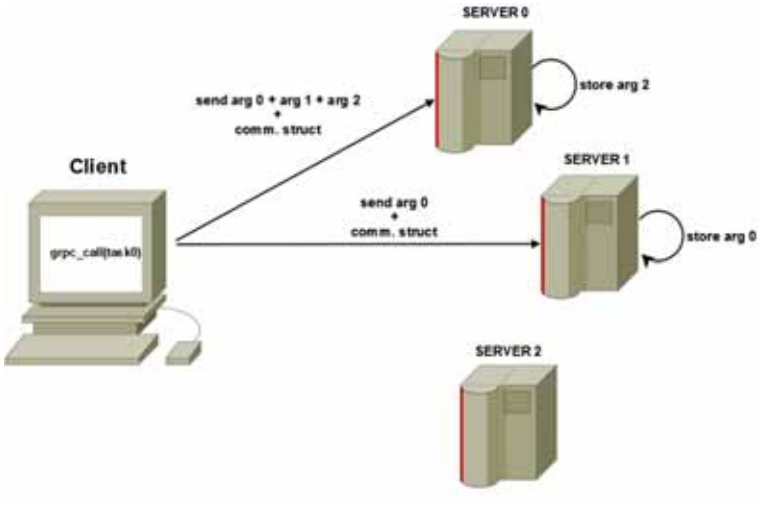

Figure 7.5. Communication transactions of task 0 of the mapping solution in figure 7.4, which happen prior to the execution of task 0.

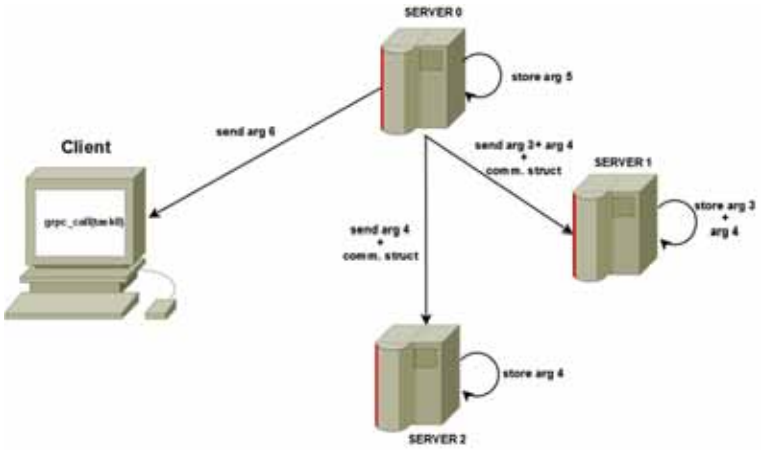

Figure 7.6. Communication transactions of task 0 of the mapping solution in figure 7.4, which happen subsequent to the execution of task 0.

guments should be received locally from a file and it will specify the file name where the argument is stored.

7.6 Fault tolerance

The *grpc_map_ft()* function in SmartGridSolve is a fault tolerant version of the *grpc_map()* function:

```
grpc_map_ft(char * mapping_heuristic_name){
  ...
  //group of GridRPC calls to map collectively
  ...
}
```

This is the same as *grpc_map()* function, except that the mapping solution generated does not implement server-server communication. The mapping solution specifies a task to server mapping and a communication scheme, which only implements communication between client and server. The communication scheme may implement:

- Client-server communication
  - Standard GridRPC communication.
- Client broadcasting
  - Client sends a single argument to multiple servers.

If any server that is part of the mapping solution fails, the tasks mapped to that server will be mapped to the next server, which is estimated to give lowest execution time for that task.

Although, no direct server communication or caching is implemented when this function is called, the performance of a group of tasks can be increased due to improved load balancing of computation and client broadcasting.

In the future, we plan to introduce a fault-tolerant method of mapping a group of tasks, which will remove this restriction on remote communication.

Chapter 8

SmartGridRPC benchmark application: The evolution of a cluster of galaxies

A typical numerical simulation needs a lot of computational power and memory footprint to solve a physical problem with a high accuracy. A single hardware platform that has enough computational power and memory to handle problems of high complexity is not easy to access. Grid computing provides an easy way to gather computational resources, whether local or geographically separated, that can be pooled together to solve large problems.

A scientific application that obviously benefits from the use of GridRPC consists of tasks with high computational loads and low communication loads. These applications, which are the best suited to run on a Grid environment, are not representative of many real-life scientific applications. Unfortunately they are typically chosen, or artificially created, to test and show the performance of GridRPC middleware systems. We believe that to justify the use of GridRPC for a wide range of applications, we should not use an extremely suitable application as a benchmark but a real life application that shows the eventual limits and benefits of the GridRPC middleware systems tested.

In this chapter, we present Hydropad [57] [58], a real-life astrophysics application that simulates the evolution of clusters of galaxies in the Universe. This application is composed of tasks that have a balanced ratio between computation

and communication. Hydropad requires high processing resources because it has to simulate an area comparable to the dimension of the Universe.

The cosmological model, which this application is based on, has the assumption that the Universe is composed of two different kinds of matter. The first is baryonic matter, which is directly observed and forms all bright objects. The second is dark matter, which is theorised to account for most of the gravitational mass in the Universe. The evolution of this system can only be described by treating both components at the same time, looking at all of their internal processes, while their mutual interaction is regulated by a gravitational component. Figure 8.1 shows an example of a typical output generated by Hydropad.

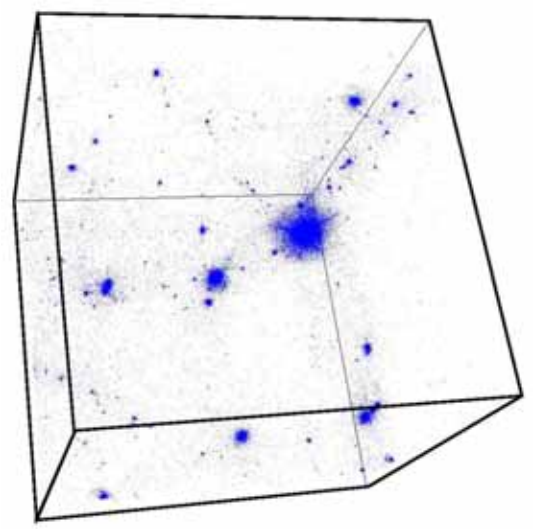

Figure 8.1. Example of Hydropad output

The dark matter computation can be simulated using N-Body methods [59]. This method utilises the interactions between a large number, Np, of collisionless particles. These particles, subjected to gravitational forces, can simulate the process of the formation of galaxies. The accuracy of this simulation depends on the quantity of particles used. Hydropad utilises a Particle-Mesh (PM) N-Body algorithm, which has a linear computational cost and depends on the number of particles, O(Np). In the first part this method transforms the particles, through an interpolation, into a grid of density values. Afterwards the gravitational po-

tential is calculated from this density grid. In the last part the particles are moved depending on the gravitational forces of the cell where they were located.

The baryonic matter computation utilises a Piecewise-Parabolic-Method (PPM) Hydrodynamic algorithm [60]. This is a high order method for solving partial differential equations. PPM reproduces the formation of pressure forces and the heating and cooling processes generated by the baryonic component during the formation of galaxies. For each time step of the evolution, the fluid quantities of the baryonic matter are estimated over the cells of the grid by using the gravitational potential. The density of this matter is then retrieved and used to calculate the gravitational forces for the next time step. The accuracy of this method depends on the number of cells of the grid used, Ng, and its computational cost is linear O(Ng). The application computes the gravitational forces, needed in the two previous algorithms, by using the Fast-Fourier-Transform (FFT) method to solve the Poisson equation. This method has a computational cost of O(Ng logNg). All the data, used by the different components in Hydropad, are stored and manipulated in three-dimensional grid-like structures. In the application, the uniformity of these base structures permits easy interaction between the different methods.

Figure 8.2 shows the workflow of the Hydropad application. It is composed of two parts: the initialisation of the data and the main computation. The main computation of the application consists of a number of iterations that simulate the discrete time steps used to represent the evolution of the Universe from the Big Bang to present time. This part consists of three tasks: the gravitational task (FFT method), the dark matter task (PM method) and the baryonic matter task (PPM method). For every time step in the evolution of the Universe, the gravitational task generates the gravitational field using the density of the two matters calculated in the previous time step. Hence the dark and baryonic tasks use the newly produced gravitational forces to calculate the movement of the matter that happens during this time step. Then the new density is generated and the lapse of time in the next time step is calculated from it. It is possible to see in figure 8.2 that the dark matter task and baryonic matter task are independent of each other.

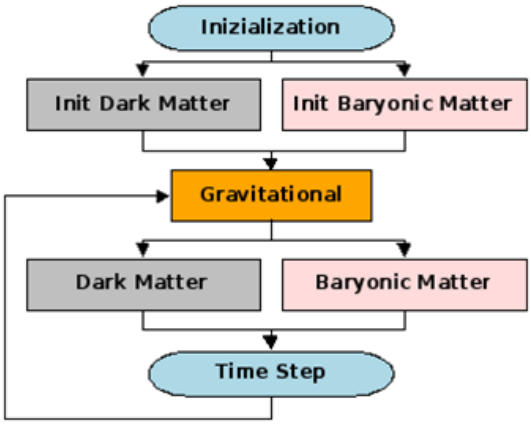

Figure 8.2. Overview of the Hydropad application

The initialisation part is also divided into two independent tasks. The main characteristic of dark matter initialisation is that the output data is generated by the external application, a module of the package COSMICS [61]. Given the initial parameters as an input, this module generates the position and velocity of the particles that will be used in the N-Body method. The output data is stored in two files where information has to be read by the application during the initialisation part. Like the main application, grafic has a high memory footprint.

An important characteristic of Hydropad is the difference in computational and memory load of its tasks. Despite both algorithms being linear, the computational load of the baryonic matter component is far greater than the dark matter one, $C_{bm} >> C_{dm}$, when the number of particles is equal to the number of cells in the grid, $N_p = N_g$. Furthermore, the quantity of data used by the dark matter computation is greater than the baryonic matter one, $D_{dm} >> D_{bm}$.

As previously indicated, Hydropad utilises three dimensional grid structures to represent the data. In the application code, these grids are represented as vectors. In the case of the dark matter component, the application stores the position and velocity in three vectors for each particle, one for each dimension. The size of these vectors depends on the number of particles, N_p, chosen to run on the simulation. For the gravitational and baryonic components, the different physical variables, such as force or pressure, are stored in vectors, with the size depending on the given number of grid cells, Ng. In a typical simulation the

number of particles is of the order of billions, while the number of cells in a grid can be over 1024 for each grid side. Given that for the values of $Ng = 128^3$ and $Np = 10^3$ the total amount of memory used in the application is roughly 500MB, the memory demand to run a typical simulation is very high.

8.1 GridRPC implementation of Hydropad

Hydropad was originally a sequential Fortran code, which was upgraded to take advantage of the GridRPC API and to work with the GridSolve middleware. Table 8.1 shows the original Hydropad code of the main loop, written in the C language. Three functions, *grav*, *dark*, and *bary*, are called in this loop to perform the three main tasks of the application. In addition, at the first iteration of this loop, a special task, *initvel* is called to initialise the velocities of the particles. The dark and baryonic tasks compute the general velocities of the respective matter. At each iteration, these velocities are used by a local function, *timestep*, to calculate the next time step of the simulation. The simulation will continue until this time becomes equal to the present time of the Universe, tsim = tuniv .

```
t_sim=0
while(t_sim<t_univ) {
  grav(phi,phiold,rhoddm,rhobm,...);
  if(t_sim==0){
    initvel(phi,...);
  }
  dark(xdm,vdm,...,veldm);..
  bary(nes,phi,...,velbm);
  timestep(veldm,velbm,...,t_step);
  t_sim+=t_step;
}
```

Table 8.1. Hydropad evolve loop

The GridRPC implementation of Hydropad application uses the APIs *grpc_call()* and *grpc_call_async()* to execute respectively a blocking and an asynchronous remote call of the Fortran functions. The first argument of both APIs is the handler of the task executed; the second is the session ID of the remote call while the following arguments are the parameters of the task. Furthermore, the code uses the method *grpc_wait()* to block the execution until the chosen, previously issued, asynchronous request has completed. When the program runs, the GridSolve middleware maps each *grpc_call()* and *grpc_call_async()* functions singularly to a remote server. Then, the middleware communicates the data from the client computer to the chosen server and then executes the task remotely. At the end of the task execution, the data is communicated back to the client. In the blocking call method, the client cannot continue the execution until the task is finished and all the outputs have been returned. Instead, in the asynchronous method, the client does not wait for the task to finish and proceeds immediately to execute the next code. The output of the remote task is retrieved when the respective wait call function is executed.

Table 8.2 outlines the GridRPC implementation of the main loop of Hydropad that simulates the evolution of the Universe. At each iteration of the loop, the first *grpc_call()* results in the gravitational task being mapped and then executed. When this task is completed, the client proceeds to the next call, which is a non-blocking call of the dark matter task. This call returns after the task is mapped and its execution is initiated. Then, the baryonic matter call is executed in the same way. Therefore, the baryonic and dark matter tasks are executed in parallel. After this, the client waits for the outputs of both these parallel tasks using the *grpc_wait()* calls.

8.2 SmartGridRPC implementation of Hydropad

The code in table 8.3 shows the modifications required to implement Hydropad for the SmartGridRPC model. The only minor difference between the GridRPC code in table 8.2 and the SmartGridRPC code in table 8.3 is the addition of: the *grpc_map()* block and *grpc_local()* condition. These belong to the SmartGridRPC API.

```
t_sim=0;
while(t_sim<t_total) {
  grpc_call(grav_hndl,phiold,...);
  if(t_sim==0){
    grpc_call(initvel_hndl,phi,...);
  }
  grpc_call_async(dark_hndl,&sid_dark,x1,..);
  grpc_call_async(bary_hndl,&sid_bary,n,...);

  /* wait for non blocking calls */
  grpc_wait(sid_dark);
  grpc_wait(sid_bary); /* to finish */
  timestep(t_step,...);
  t_sim+=t_step;
}
```

Table 8.2. Hydropad implementation in GridRPC

On the first iteration through the map block each task is discovered. The specified mapping heuristic, in this case the greedy mapping heuristic, generates a mapping solution for this group of tasks. On the second iteration through the group of tasks the group is executed according to the mapping solutions generated.

The *grpc_local()* function is used by the application programmer to indicate when a local computation is executed. At run time, on the first discovery iteration, the code within this conditional statement is not executed. This is to avoid computing local executions when generating a performance model for the group of remote tasks. However, if a local computation directly affects the performance model of the group of remote tasks, then the *grpc_local()* function should not be used. This would be the case if a local computation affects whether certain remote tasks get executed or affects the size of computation of tasks. If this were the case, then the local computation should be executed during discovery and any structures, variables etc. that have changed values should be reset back to their original values before the beginning of execution.

On the second iteration, during the execution phase all the code in *grpc_map()* function is executed normally (i.e. the local computation is also executed). The mapping in the code of table 8.3 is performed at every iteration of the main while

```
t_sim=0;
while(t_sim<t_univ) {
  grpc_map("greedy_map"){
    grpc_call(grav_hndl,phiold,...);
    if(t_sim==0){
      grpc_call(initvel_hndl,phi,...);
    }
    grpc_call_async(dark_hndl,&sid_dark,x1,.);
    grpc_call_async(bary_hndl,&sid_bary,n,..);

    /* wait for non blocking call */
    grpc_wait(sid_dark);
    grpc_wait(sid_bary);

    grpc_local(){
      timestep(t_step,...);
      t_sim+=t_step;
    }
  }
}
```

Table 8.3. Hydropad implementation in SmartGridRPC

loop. Generating frequent mapping solutions like this can generate good mapping solutions if the Grid environment is not a stable one. This would be the case if there are other application tasks running on the Grid servers or if other users are using machines on the Grid. If the Grid environment is dedicated, where only one application executes at a time, a better mapping solution may be generated if the area to map contains more tasks, i.e. two or more loop cycles. A simple solution could be including an inner loop within the *grpc_map()* code block (table 8.4). The application programmer could increase the number of tasks mapped together by increasing the number of iterations of the inner loop.

Figure 8.3 is a task graph generated for only two cycles of the evolution step. It is also possible to map a significantly larger number of evolution steps, by increasing the value of the *nb_steps* variable in table 8.4.

This type of coarse mapping would be more favourable on a distributed environment, which is highly stable, for example, a distributed environment that consisted of dedicated servers or servers that are idle. However, if the environment is highly changeable, which would be the case if the distributed environment consisted of workstations currently being used, then it might be more advantageous

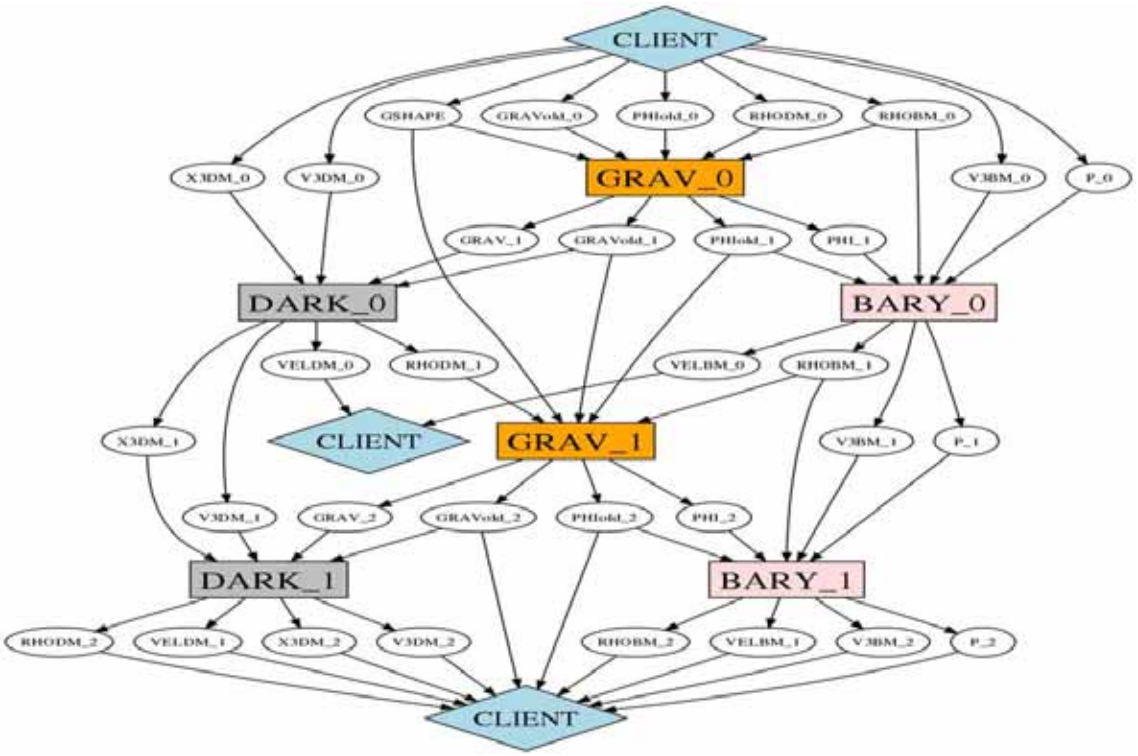

Figure 8.3. Task graph for two evolution steps

to have a higher frequency of mappings. It may also be necessary to increase the frequency of mappings, if the task graph is altered as a result of the execution of one of the remote tasks in the task graph. For example, this may be the case if there is a conditional statement in the group of tasks that is based on an output of a remote task in the group (*task A*). If this conditional statement determines whether another remote task (*task B*) gets executed then the shape of the task graph depends on the output of *task A*. When the shape of a task graph is determined by the outputs of a remote task in the group then it is important to increase the frequency of mappings and perform mappings whenever the task graph is altered. To ensure the shape of the task graph is accurate in the aforementioned case, the task graph should be generated and mapped every time *task A* is executed.

It is also possible to make this mapping frequency more dynamically adaptive. In table 8.4, the value assigned to the variable *nb_steps* indicates how many evolution steps should be mapped collectively at the next point of execution of the application. This value can be fine-tuned during the execution of the application to determine the optimal number of evolutions to map as a group. In this example, the value for *nb_steps* is updated and fine-tuned using an evaluation function *func()*. This may be a function that changes the value of the variable *nb_steps* based on an evaluation of the performances of previous executions of collective mappings.

This approach can be used to find the optimal mapping for an application on any given distributed environment. Once determined, this optimal number can then be assigned statically for each subsequent execution of the application on this environment without the need for an evaluation function.

However, in the case where the environment is highly changeable, this optimal number of evolutions may vary throughout the execution of the application and therefore it may be more beneficial to maintain this dynamic update of *nb_steps* variable at run-time.

```
t_sim=0;
while(t_sim<t_univ) {
  nb steps=func(..); //assign dynamically
  grpc_map("greedy_map"){
     for(i=0;(i<nb_steps)&(t_sim<t_univ);i++){
      grpc call(grav_hndl,phiold,...);
      if(t_sim==0){
        grpc_call(initvel_hndl,phi,...); }
        grpc_call_async(dark_hndl,&sid_dark,..);
        grpc_call_async(bary_hndl,&sid_bary,..);

        /* wait for non blocking call */
        grpc_wait( sid_dark);
        grpc_wait( sid_bary);

        grpc_local(){
          timestep(t_step,...);
          t_sim+=t_step;
        }
      }
    }
  }
}
```

Table 8.4. Dynamically determining the optimal group size to map

Chapter 9

Experimental results

In the experiments performed in this section, we use three different implementations of Hydropad: the original sequential version, a GridSolve version and a SmartGridSolve version. For each version, we present the average computation time of one evolution step and the memory footprints of the application on the client machine. In section 9.1, we compare the GridSolve version versus the local sequential version. Then, in section 9.2 we compare the SmartGridSolve version of Hydropad versus both GridSolve and the local one. Furthermore, in section 9.2 we focus on the performance improvement of each of the key benefits of the SmartGridRPC model over the GridRPC model, which were introduced in section 2.2.

The hardware configuration used in the experiments consists of three machines: a client and two remote servers, S1 and S2. The two servers are heterogeneous. However, they have similar performance, respectively 498 and 531 MFlops, and they have equal amount of main memory, 1GB each. The bandwidth of the communication link between the two servers is 1GB/s. The client machine, C, is a computer with low hardware specifications, 248MFlops of performance. The client to server connection varies depending on the experimental setup. We use two setups, C1 with a 1GB/s connection and C100 with a 100MB/s communication link. For each conducted experiment, table 9.1 shows the initial problem parameters and the corresponding data sizes (the total memory used during the execution of Hydropad on a single machine). The quantity of memory available in the client machine varies as well depending on the experimental setup. We

use two configurations: C1-1 with 1GB of memory, which is large enough to avoid paging, and C-256 with 256MB of memory, that undergo paging for larger problems.

Table 9.1. Input values and problem sizes for the Hydropad experiments

Problem ID	N_p	N_g	Data Size
P1	120^3	60^3	73MB
P2	140^3	80^3	142MB
P3	160^3	80^3	176MB
P4	140^3	100^3	242MB
P5	160^3	100^3	270MB
P6	180^3	100^3	313MB
P7	200^3	100^3	340MB
P8	220^3	120^3	552MB
P9	240^3	120^3	624MB

9.1 Experiments with the GridSolve version of Hydropad

Table 9.2 shows the results obtained by the local sequential version and by the GridSolve version of Hydropad using C1-1 as the client machine which has a fast network connection and a large quantity of memory.

Table 9.2. Experimental results using client C1-1 that has 1GB/s network link to the servers.

	Local	GridSolve	Local vs GS
P. ID	Time Step	Time Step	S_p
P1	14.12s	9.40s	1.50
P2	29.90s	18.38s	1.63
P3	34.84s	20.82s	1.67
P4	52.04s	30.81s	1.69
P5	54.06s	32.00s	1.69
P6	58.56s	36.81s	1.59
P7	66.29s	37.22s	1.78
P8	102.03s	67.04s	1.52
P9	114.83s	112.05s	1.02

One can see that the GridSolve version is faster than the local sequential computation. The speedup obtained is over 1.50 for all problem sizes except for P9. This speedup is due to the parallel execution of the two tasks and the use of servers with greater performance than the client machine. The fluctuation in speedup obtained by GridSolve depends on the varying ratio of data size used by the two parallel tasks for different problem sizes. Furthermore, it should also be noted, that the speedup achieved on P9 is significantly lower due to paging on the server. This is caused by the fact that the GridRPC model maps both tasks to the same server and therefore causes paging on it.

Table 9.3 shows the results obtained by the GridSolve version when the client machine used, C100-256, has a slow client-to-servers connection of 100MB/s and only 256MB of memory available. This hardware configuration simulates a common situation that can happen in real life. A user has access only to a slow client machine with low hardware specification, which is not suitable to perform large simulations, and wants to use a powerful Grid environment through a relatively slow network link. Table 9.3 also presents the scale of paging that occurs on the client machine during the executions. One can see that for the local computation the paging is taking place when the problem size is equal to or greater than the machine main memory, 256MB.

In these experiments, "light" paging means that paging is occurring only in some task calls and the amount of paging is approximately 10% of the main memory (approx. 25MB). "Normal" paging means that paging is occurring on almost every task call and the amount of paging is approximately 40% of the main memory (approx. 100MB). "Heavy" paging means that all task calls cause a memory page and almost 100% of the main memory is paged (approx. 256MB).

The GridSolve version is slower than the local computation when the client machine is not paging. This is happening because there is a large amount of data communication between tasks. So for this configuration, the time spent communicating the data compensates for the time gained by computing tasks remotely. However, as the problem size gets larger and the client machine starts paging, the GridSolve version becomes faster than the local computation, even in the case of slow communication between the client and server machines. This trend is also seen in figure 9.1. For the GridSolve version, the paging is occurring

Table 9.3. Experimental results using client C100-256 that has 100MB/s network link to the servers and 256MB of memory.

P.ID	(a) Local		(b) GridSolve		
	Time Step	Paging	Time Step	Paging	S_p v Local
P1	14.32s	No	20.26s	No	0.71
P2	30.05s	No	38.75s	No	0.78
P3	35.78s	No	48.65s	No	0.74
P4	55.57s	Light	60.48s	No	0.92
P5	62.13s	Light	66.43s	No	0.94
P6	84.33s	Yes	76.76s	Light	1.10
P7	128.22s	Yes	93.74s	Yes	1.37
P8	231.56s	Heavy	150.03s	Heavy	1.54
P9	279.52s	Heavy	183.45s	Heavy	1.52

later than for the local version, when the problem size is around 310MB, as shown in table 9.3. The GridRPC implementation can save memory due to the temporary data allocated remotely in the tasks and consequently increasing the problem size will not cause paging. Furthermore, in the sequential local execution, the paging is taking place during a local task computation, while for the GridSolve version the paging occurs during a remote task data communication. Hence, for the GridSolve version of Hydropad, the paging on the client machine does not negatively affect the execution time of the application.

9.2 Experiments with the SmartGridSolve version of Hydropad

In the first experiment of this section, we use the same hardware configuration of table 9.3. The client machine used, C100-256, has a slow client-to-servers connection of 100MB/s and only 256MB of memory available. As previously mentioned, this is a common situation. Table 9.4 shows the results obtained by the SmartGridSolve version for this configuration. This table shows that the SmartGridSolve version is much faster than the GridSolve and the sequential versions. The speedup is around three times that of GridSolve (figure 9.2) and the speedup versus the local sequential version is over 4 in the case of larger problems.

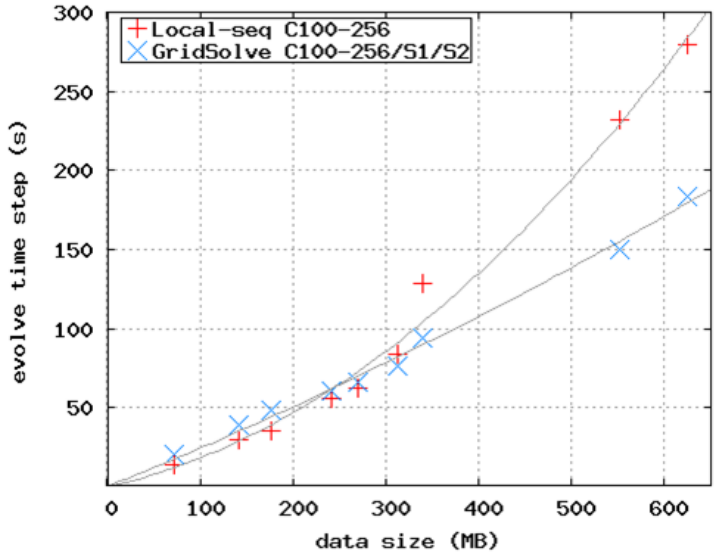

Figure 9.1. Evolution time step of the local and GridSolve computation on client C100-256

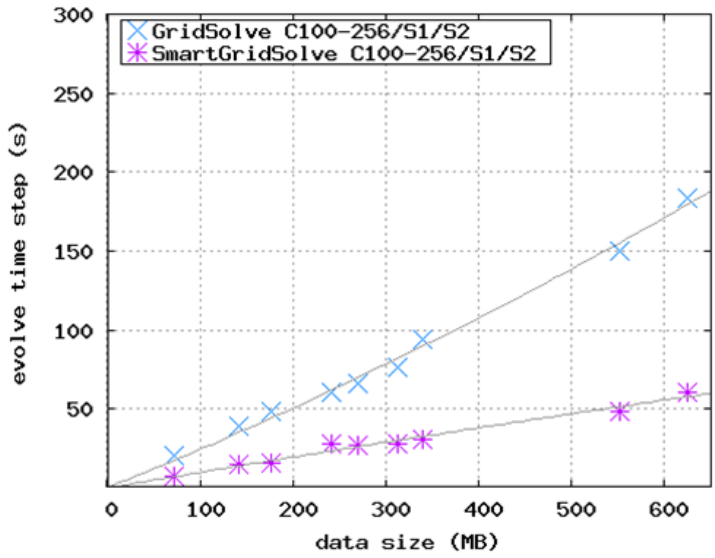

Figure 9.2. Execution times of the GridSolve and SmartGridSolve version of Hydropad on client C100-256

Table 9.4. Experimental results using client C100-256 that has 100MB/s network link to the servers and 256MB of memory

(a) Local

P.ID	Time Step	Paging
P1	14.32s	No
P2	30.05s	No
P3	35.78s	No
P4	55.57s	Light
P5	62.13s	Light
P6	84.33s	Yes
P7	128.22s	Yes
P8	231.56s	Heavy
P9	279.52s	Heavy

(b) GridSolve

Time Step	Paging	S_p v Local
20.26s	No	0.71
38.75s	No	0.78
48.65s	No	0.74
60.48s	No	0.92
66.43s	No	0.94
76.76s	Light	1.10
93.74s	Yes	1.37
150.03s	Heavy	1.54
183.45s	Heavy	1.52

(c) SmartGridSolve

P.ID	Time Step	Paging	S_p v Local	S_p v GS
P1	7.31s	No	1.96	2.77
P2	15.06s	No	2.00	2.57
P3	16.36s	No	2.19	2.97
P4	28.06s	No	1.98	2.16
P5	27.54s	No	2.26	2.41
P6	27.78s	No	3.04	2.76
P7	30.81s	Light	4.16	3.04
P8	48.04s	Light	4.82	3.12
P9	60.74s	Light	4.06	3.02

These performance improvements are due to the key features of the SmartGridRPC model: improved mapping, improved data movement and reduced memory usage. In the next experiments of this section, we show the benefits introduced by each feature by using specific hardware configurations and setup.

Computational load experiments. One important feature of SmartGridRPC is the superior mapping system that permits improved balancing of computational load of tasks compared to standard GridRPC. In the underlying experiments, we compare the average computation time of one evolution step achieved by the GridSolve version versus the SmartGridSolve version of Hydropad, where SmartGridSolve is set up to utilize the same network topology of GridSolve (star-network), i.e. without direct server-to-server communication and server-caching. Consequently, the performance gains obtained by the SmartGridSolve version are due only to the improved mapping method. In these experiments, we use C1-1 as the client machine. This machine has a high speed network connection of 1GB/s

to the servers. Table 9.5 shows that the SmartGridSolve version of Hydropad is faster than the GridSolve version.

Table 9.5. Experimental results using only star-network topology (i.e. no direct server-to-server communication) and client C1-1 that has 1GB/s network link to the servers

	GridSolve	SmartGridSolve	GS vs SmartGS
P. ID	Time Step	Time Step	S_p
P1	9.40s	7.09s	1.33
P2	18.38s	15.27s	1.20
P3	20.82s	16.17s	1.29
P4	30.81s	29.02s	1.06
P5	32.00s	28.99s	1.10
P6	36.81s	29.88s	1.23
P7	37.22s	30.88s	1.21
P8	67.04s	50.05s	1.29
P9	112.05s	53.35s	2.10

Despite Hydropad having only two parallel tasks, the collective mapping of SmartGridRPC can produce a faster execution time than the individual task mapping of GridRPC. The baryonic task is computationally far larger than the dark matter one, $C_{bm} >> C_{dm}$. When a GridRPC system maps these two tasks, it does so without the knowledge that they are part of a group to be executed in parallel. Its only goal is to minimize the execution time of an individual task as it is called by the application. If the smaller dark matter task is called first, it will be mapped to the fastest available server. With the fastest server occupied, the larger baryonic task will then be mapped to a slower server and the overall execution time of the group of tasks will be sub-optimal. As previously mentioned, in the case of problem P9, both tasks were mapped to the same server, which increased the total execution time and caused paging on the server.

Communication load experiments. As mentioned before, another primary improvement of SmartGridSolve is its communication model, use of which minimizes the amount of data movement between the client and servers. This advantage is most prominent when the client connection to the Grid environment is slow. Table 9.5 shows the results obtained by the SmartGridSolve version of

Hydropad using C100-1 as the client machine, which has a slow network connection of 100MB/s. One can see that the SmartGridSolve version is much faster than the GridSolve versions. The increase of speed is over twice that of GridSolve, which is primarily due to the improved communication model of SmartGridSolve.

Furthermore, one can see that the timing results obtained by SmartGridSolve in table 9.5 are similar to those obtained in table 9.6. This shows that when the client-server links are slow and there is direct communication (table 9.5), it is similar to when the client links are fast and there is no direct communication (table 9.6). This shows that the SmartGridRPC model allows the mapping heuristic to generate solutions, which effectively minimize the communication load on the network link.

Table 9.6. Experimental results using client C100-1 that has 100MB/s network link to the servers

	GridSolve	SmartGridSolve	GS vs SmartGS
P. ID	Time Step	Time Step	S_p
P1	19.97s	7.24s	2.76
P2	38.73s	15.17s	2.55
P3	48.20s	16.24s	2.97
P4	61.59s	29.42s	2.09
P5	66.26s	28.91s	2.29
P6	78.16s	29.73s	2.63
P7	93.20s	31.25s	2.99
P8	140.53s	50.20s	2.80
P9	174.14s	53.02s	3.28

Memory usages experiments. In the following experiments, we utilize the client machine C1-256, that has a high speed network connection of 1GB/s to the servers and has 256MB of main memory. Table 9.7 shows the average computation time of one evolution step achieved by the local sequential version, by the GridSolve version and by the SmartGridSolve version of Hydropad. Table 9.7 also presents the scale of paging that occurs on the client machine during the various executions.

One can see that for the SmartGridSolve experiments the paging on the client machine is less penalizing than in the GridSolve experiments and local experi-

Table 9.7. Experimental results using client C1-256 that has 1GB/s network link to the servers and 256MB of memory

(a) Local

P.ID	Time Step	Paging
P1	14.32s	No
P2	30.05s	No
P3	35.78s	No
P4	55.57s	Light
P5	62.13s	Light
P6	84.33s	Yes
P7	128.22s	Yes
P8	231.56s	Heavy
P9	279.52s	Heavy

(b) GridSolve

Time Step	Paging	S_p v Local
8.6s	No	1.67
18.4s	No	1.63
20.1s	No	1.77
31.3s	No	1.77
33.7s	No	1.84
42.3s	Light	1.99
63.1s	Yes	2.03
109.3s	Heavy	2.12
144.3s	Heavy	1.94

(c) SmartGridSolve

P.ID	Time Step	Paging	S_p v Local	S_p v GS
P1	7.0s	No	2.02	1.21
P2	14.4s	No	2.08	1.27
P3	15.8s	No	2.26	1.27
P4	27.5s	No	2.02	1.14
P5	28.1s	No	2.21	1.20
P6	28.8s	No	2.92	1.47
P7	30.0s	Light	4.27	2.10
P8	46.6s	Light	4.96	2.34
P9	55.1s	Light	5.07	2.62

ments. A secondary advantage of the direct server to server communication implemented in SmartGridSolve is that the quantity of memory used on the client machine is lower than that of the GridSolve version. Furthermore, in SmartGridSolve, the memory paging is happening only when data has to be sent to the server. Hence, it happens only at the beginning and at the end of a group of tasks execution. This minimizes the impact of paging on the overall execution of the group of tasks.

Therefore, the SmartGridSolve version of Hydropad can execute larger problems without paging having a serious impact on the execution time. One can see that the computation time of the evolution steps in table 9.7 is similar to that of table 9.5 and 9.6. The speedup of SmartGridSolve over GridSolve, is increasing as the problem gets larger due to paging on the client. This trend is also seen in figure 9.3.

The new features of SmartGridRPC have also a secondary benefit. As previ-

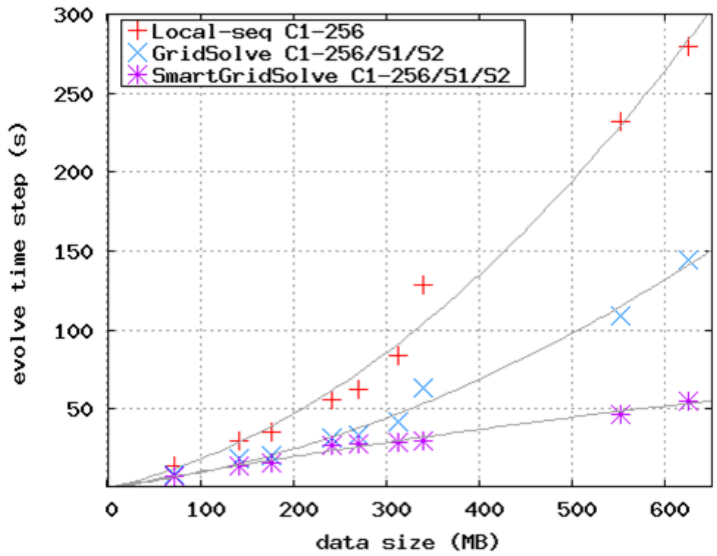

Figure 9.3. Execution times of the GridSolve and SmartGridSolve version of Hydropad when the client machine C1-256 has 256MB of memory

ously mentioned, SmartGridSolve obtains similar results when the client memory and the client-to-server link are largely different. Consequently, the hardware configuration of the client has less impact on the application performance than in the case of GridRPC. Figure 9.4 shows this trend. We compare the results obtained by GridSolve and SmartGridSolve version of Hydropad when the two configurations of the client used are the optimal one, C1-1, and the worst one, C100-256. It is possible to see that in the case of GridSolve the performance changes dramatically depending on the hardware used while for SmartGridSolve the performance is similar.

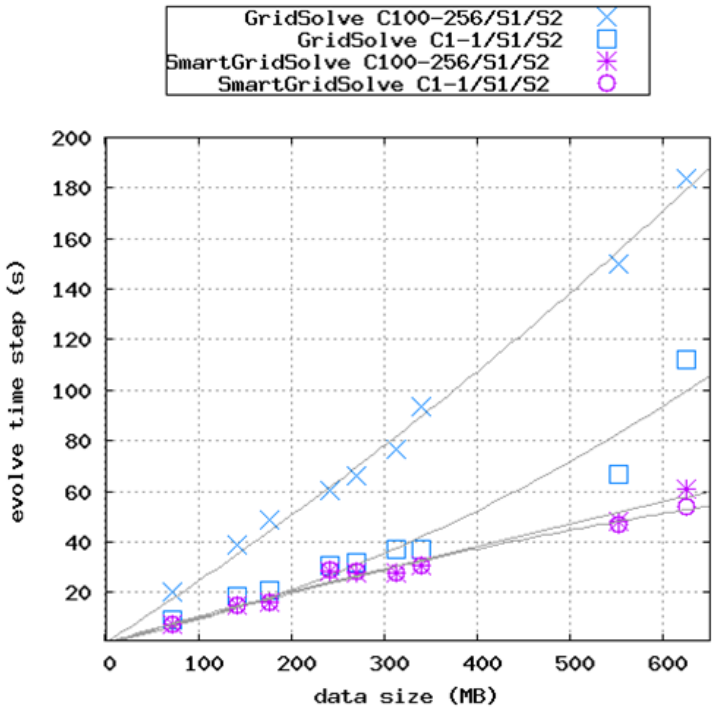

Figure 9.4. Execution times of the GridSolve and SmartGridSolve version of Hydropad when the client machines are C1-1 and C100-256

Chapter 10

Conclusion

In this thesis, we have presented the SmartGridRPC model, which is an extension to the GridRPC model that aims to achieve high performance.

The GridRPC model maps each task in an application separately and independently of other tasks of the application. Given this restriction, the model can only support the minimization of the execution time of each individual task of the application rather than the minimization of the execution time of the whole application.

Another restriction of GridRPC is its communication model. The communication model of GridRPC is based on the client-server model or star network topology. This means that tasks can be executed on any of the servers and inputs/outputs can only traverse the client-server links.

Mapping tasks individually on a star network results in several limitations. Firstly, the computation load of a group of parallel tasks cannot be analysed collectively and therefore the computation load will not be optimally balanced over the available servers. Also, if dependencies exist in a group of tasks bridge communication is forced. This is because dependencies between tasks cannot be analysed and servers cannot communicate directly. This increases the communication load on the network. Bridge communication results in intermediate results being sent back to and stored on the client. Consequently, this increases the amount of memory used on the client and may result in paging, which could significantly increase the execution time of the application. Furthermore, since dependencies are not known and the network is based on the client-server model,

it is impossible to employ any parallelism of communication between the tasks in the group. For example, this could be implemented if there is a dependency between two tasks and the destination task is not executed in parallel or immediately after the source task. In theory, this dependent data could be sent to the destination task in parallel with any computation or communication on any other machine (client or other servers) which happens in the intervening time. But since tasks are mapped individually on to a star network, this parallelism of communication cannot be realized using the GridRPC model.

In related research (chapter 3) we outlined several extensions to the GridRPC model which fall into the following categories:

- extensions to the GridRPC model which extend the client-server model to implement direct communication between servers or data persistence.
- extensions to the GridRPC model which extend the system so that a group of tasks can be collectively mapped.

The papers introduced in this section also have a number of limitations which were outlined in the related research chapter.

The SmartGridRPC model enhances the GridRPC model to allow a group of tasks to be mapped collectively on to a fully connected network. This removes each of the aforementioned limitations of the GridRPC model and the limitations of the extensions of the GridRPC model described in chapter 3.

The SmartGridRPC model can increase the performance of an application by:

- Improving the load balancing of computation.
- Improving the load balancing of communication.
- Reducing the overall volume of communication.
- Reducing memory usage on the client (reduce paging).
- Parallelising communication.

The SmartGridRPC model provides an API, which allows the application programmer to specify a block of code, in which a group of GridRPC task calls should be mapped collectively. Then, when the application is run, the specified group of tasks in this block of code is processed collectively and each operation in the GridRPC call is separated and done collectively for all tasks in the group. Namely, all tasks in the group are discovered collectively, mapped collectively and executed collectively on the fully connected network. In the discovery phase, performance models are generated for estimating the execution time of the group of tasks on the fully connected network. In the mapping phase, the performance models are used by the mapping heuristic to generate a mapping solution for the group of tasks. In the execution phase, the group of tasks is executed on the fully connected network according to the mapping solution generated.

In chapter 7 we described the implementation of the SmartGridRPC model in SmartGridSolve, which is an extension to the GridSolve middleware. It described a possible implementation of the performance models which are used to estimate the execution times of executing a group of tasks on a fully connected network.

In chapter 8, we presented Hydropad, a real-life astrophysics application that simulates the evolution of clusters of galaxies in the Universe. This application simulates the evolution of clusters of galaxies in our Universe from the beginning of time till present. The reason this application was chosen as a benchmark application was that it is not well suited to be implemented in Grid environments and consequently it can show the eventual limits and benefits of the two models tested.

In chapter 9 we gave an experimental evaluation of the SmartGridRPC model by comparing the execution of Hydropad application using a SmartGridSolve version (SmartGridRPC model), a GridSolve version (GridRPC model) and a local sequencial version of the application.

In these experiments we evaluated these models using a hardware configuration that consisted of three machines, a client and two servers. The client is approximately half the speed of the two servers which simulates a common situation that happens in real life where the user has access to a low specification client machine and wants access to more powerful machines. We used two network configurations, either the client-server link was slow (100MB/s) or fast (1GB/s)

and the client either had a small amount of memory available (256MB) or a large amount (1GB).

For the 1GB/s link speed and 1GB client memory configuration it was shown that the GridSolve version improves the performance over the local sequential version for all the given problem sizes. The reason for this speedup is the due to the fact that server machines are more powerful than the client machine and there is parallelism of the baryonic and darkmatter tasks of Hydropad. Also for larger problem sizes, the memory footprint on the local sequential version causes a greater level of paging on the client. This is reduced using GridSolve as the memory allocation for computation is done on the servers as opposed to on the client so therefore the level of paging is either decreased or avoided using GridSolve. In GridSolve, memory is allocated on the client for communication but not for computation.

For the 100MB/s link speed and 256MB client memory configuration, it was shown that GridSolve improves the performance over the local sequential version only for larger problem sizes. The reason for this is due to the heavy penalty of communicating the data over the slow client-server network link. However, when the problem sizes is large enough as to cause paging on the client, the GridSolve version begins to outperform the local sequential version due to the temporary memory allocation for computation being done on the servers as opposed to on the client.

The experiments also showed how the SmartGridSolve could further improve the performance of the Hydropad application. This is due to improving the load balancing of computation, reducing the overall volume of communication, increasing the parallelism of communication and decreasing the level of paging on the client.

Firstly, SmartGridSolve improved the load balancing of computation by improving the distribution of the load of the two parallel tasks. It assigned the task with larger computational load (baryonic matter) to the faster server and the task with the smaller computational load (dark matter) to the slower server. In addition, it reduced the volume of communication on the network by eliminating bridge communication through direct server-server communication. Also, the performance was increased due to asynchronous remote communication as

server-server communication was done asynchronously and therefore it is done in parallel with other communication and computation. Moreover, it further decreased the paging on the client. This is because SmartGridSolve further reduces the memory allocation on the client because it eliminates bridge communication and therefore intermediate results are not sent back to and stored on the client. As a result of these performance enhancements, SmartGridSolve shows significant speedup over both the local sequential execution of Hydropad and the GridSolve execution for each of tested problem sizes.

10.1 Contributions

Below, we present more precisely the contributions of this work

The design of the SmartGridRPC model and API: The main goal of the design of the API was to allow the application programmer to easily convert an existing GridRPC application into a SmartGridRPC enabled application. This facilitates a simple transformation from a standard GridRPC application into a SmartGridRPC application which can fully exploit the potential of collective mapping of tasks on an execution model that is based on a fully connected network. It has been shown that the collective mapping of tasks on a fully connected network improves the chances of finding a closer to optimal solution over the individual mapping of tasks on to the star network of the GridRPC model. Therefore an application programmer with a few simple changes to their client application can easily benefit from all the performance enhancements of the SmartGridRPC model.

The SmartGridRPC model was designed to be both incremental and interoperable with the underlying GridRPC middleware. It is incremental to the GridRPC model which means a SmartGridRPC client application can be executed on a standard GridRPC network where all servers can only communicate with the client. Furthermore, a SmartGridRPC application can be executed on a network consisting of both GridRPC enabled servers and SmartGridRPC enabled servers (*SmartServers*) that can communicate with each other. It is interoperable with GridRPC model which means that if the SmartGridRPC extension is

installed, the application programmer has the option to implement their application using the GridRPC model or the SmartGridRPC model. As a result of these interoperability and incremental features SmartGridRPC is both upward and downward compatible with the GridRPC model.

The design and implementation of SmartGridSolve: The GridSolve middleware was extended to be SmartGridRPC compliant and this extension is called SmartGridSolve. Details of this implementation were described in this thesis to demonstrate how a developer of a GridRPC middleware can extend their middleware so that it is SmartGridRPC enabled.

The implementation of SmartGridSolve required extensions to the following aspects of GridSolve:

- Network Discovery.
- Task Discovery.
- Mapping Heuristics.
- Execution model.

In GridSolve, in order to map an individual task on to a star network, the programming system discovers the performance model of the star network (network discovery) and a representation of an individual task (task discovery). Then the mapping heuristic generates a mapping solution for the individual task based on the performance model of the star network and representation of the individual task . Then this task is executed in the distributed environment which is based on the client-server model (execution model).

In GridSolve, the discovery of the performance model of the star network includes the discovery of the dynamic performance of each server on the network and the dynamic performance of client-server links. The discovery of the representation of an individual task includes the discovery of the computational load of the task and the communication load of the input and output arguments. Using the individual task representation and performance model of the star network, the mapping heuristic times the communication time and computation time of that task on each server and returns a mapping solution with specifies the servers

with the lowest execution time. This mapping solution is then executed on an execution model which is based on the client-server model. The inputs are sent to the mapped server, the mapped server executes the task and sends the output arguments back to the client.

In order to map a group of tasks, the extensions illustrated in figure 10.1 had to be implemented in GridSolve.

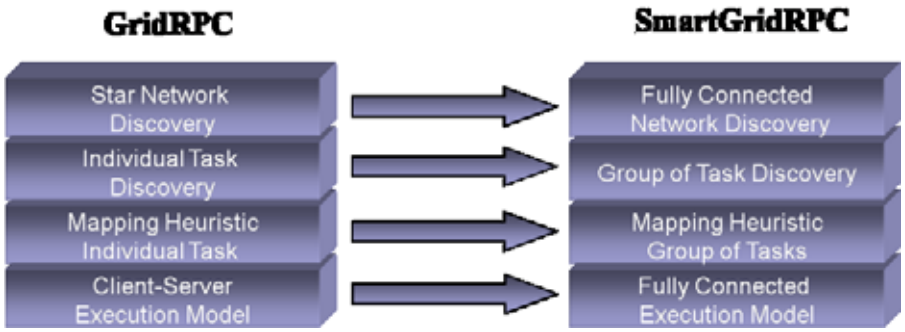

Figure 10.1. SmartGridSolve extensions

The GridSolve programming system was extended to discover the performance model of the fully connected network. In addition to performing the discovery of the dyanamic performance of each server on the network and the dynamic performance of client-server links, the programming system also discovers the dynamic performance of the links connecting *SmartServer* to other *SmartServers*. Furthermore, the programming system was extended to discover a representation of a group of tasks as opposed to an individual task. This was made possible by the addition of the *grpc_map()* function. The implementation of this function and how it generates a task graph was described in this thesis. Furthermore, the mapping heuristics of GridSolve which maps an individual task on to a star network were extended so that they map a group of tasks on to a fully connected network. Finally, the execution model was extended so the mappping solution generated by these heuristics can be executed on a fully connected network. The GridSolve servers were extended so that they can perform asynchronous push communication to other servers on the network. The GridSolve servers were also extended so they can perform asynchronous push broadcast communication where

they communicate outputs to multiple servers. Server caching was also implemented, where servers can store outputs to the local filesystem where they can be retrieved by subsequent tasks executing on the same server. This communication and caching of outputs is asynchronous and therefore can be done in parallel with other computation or communication which further increases the opimization of the executing group of tasks.

Appendix A

Appdx A - SmartGridSolve Manual

This manual demontrates how to configure SmartGridSolve and implement applications using the SmartGridRPC API. SmartGridSolve can be downloaded from the Heterogeneous Computing Laboratory website.
http://hcl.ucd.ie/project/SmartGridSolve
The SmartGridSolve extension is also included in the official release of GridSolve which can be download from the GridSolve website.
http://icl.cs.utk.edu/netsolve/software/index.html

A.1 Using SmartGridSolve

To use SmartGridSolve, one should enable the SmartGridSolve feature both on the GridSolve client and server.
Type
% ./configure -enable-smartgridsolve
% make
% make services
during the initial configuration of GridSolve. Note that services/tasks require compilation when GridSolve system has been configured with SmartGridsolve extension enabled.

A.2 SmartGridSolve API

The SmartGridSolve API allows a user to specify the scope of a group of tasks to be mapped collectively.

A.2.1 The grpc_map() function

This function is used for specifying the scope of the group of tasks and the mapping heuristic to implement.

```
grpc_map(char * mapping_heuristic_name)
```

Parameters:

- mapping_heuristic_name - Name of the mapping heuristic to implement when mapping the group of tasks.

Usage:

```
grpc_map(char * mapping_heuristic_name){
  ...
  // group of tasks to map collectively
  ...
}
```

Description:

The *grpc_map()* "function" is in fact a macro that inserts a while loop around the code block specified by the parenthesis. When the *grpc_map()* function is called the code within its parenthesis will be iterated through twice. On the first iteration each *grpc_call()* and *grpc_call_async()* is discovered but not executed. From these discovered calls a task graph is generated. At the beginning of the second iteration the mapping heuristic specified by the *grpc_map()* parameter will generate a mapping solution based on the task graph and the performance model of the network. The mapping solution outlines a task to server mapping and also the communication scheme between tasks.

During the second iteration through the code, the tasks will be executed according to the generated mapping solution. It should also be noted that handles and sessionids should be created, initialised, destroyed and deleted outside the scope of the parenthesis of the *grpc_map()* function.
The communication scheme may implement:

- Client-server communication
  - Standard GridRPC communication
- Server-server communication
  - Server sends a single argument to another server
- Client broadcasting
  - Client sends a single argument to multiple servers.
- Server broadcasting
  - Server sends a single argument to multiple servers.
- Server caching
  - Server stores an argument locally for future tasks.

Example:
In the example in table A.1, the *grpc_map()* function is the only addition required to make this code SmartGridSolve enabled. As previously explained the handles and sessionids should be created, initialised, destroyed and deleted outside the scope of the parenthesis of the *grpc_map()* function.

A.2.2 The grpc_map_ft() function

This function is a fault tolerant version of the *grpc_map()* function.

```
grpc_map_ft(char * mapping_heuristic_name)
```

Table A.1. Example implemention of the grpc_map() function

```
grpc_function_handle_t*handles=
   (grpc_function_handle_t*)calloc(iters, sizeof(grpc_function_handle_t ));

grpc_sessionid_t * sessionIDs=( grpc_sessionid_t *)calloc(iters, sizeof(grpc_sessionid_t));
int * status=(int *)calloc(iters, sizeof(int));

 for(i=0; i<iters; i++){
     if(grpc_function_handle_default(&handle[i], "dgesv") != GRPC_NO_ERROR){
       fprintf(stderr,"Error creating function handle");
       die(EXIT_FAILURE);
   }
}

...
...

 grpc_map("ex_map"){
    for(i=0;i<iters;i++){
      status[i]=grpc_call_async(&handle[i], &sessionID[i], .., .., ..);
   }
}

.....
.....

 for(i=0;i<iters;i++){
   if(grpc_function_handle_destruct(handle[i]) = GRPC_NO_ERROR) {
      fprintf(stderr,"Error destroying function handle1");
      die(EXIT_FAILURE);
   }
}
```

Parameters:

- mapping_heuristic_name - Name of the mapping heuristic to implement when mapping the group of tasks.

Usage:

```
grpc_map_ft(char * mapping_heuristic_name){
  ...
  // group of tasks to map collectively
  ...
}
```

Description:

This is the same as *grpc_map()* function, except that the mapping solution generated does not implement server-server communication. The mapping solution outlines a task to server mapping and a communication scheme which only implements communication between client and server. The communication scheme may implement

- client-server communication.
  - standard GridRPC communication.
- client broadcasting.
  - client sends a single argument to multiple servers.

If any server that is part of the mapping solution fails, then the tasks mapped to those servers will be mapped to the next server which is estimated to give the lowest execution time for that task.

Example:

Implementation is the same as *grpc_map()* in table A.1, just change the function from *grpc_map()* to *grpc_map_ft()*.

A.2.3 The grpc_local() function

This function is used to specify the code that should be ignored during the first iteration through the scope of *grpc_map()* (i.e. code that should be ignored when building the task graph).

```
grpc_local()
```

Usage:

```
grpc_map(char * mapping_heuristic_name){
  //reset variables which have been updated
  //during task discovery
  ...
  if(grpc_local()){
    // code to ignore during task discovery
  }
  ...
  // group of tasks to map collectively
  ...
}
```

Description:

Any segment of client code that is not part of the GridRPC API should not be executed during task discovery. To achieve this, such code must be enclosed in the conditional that tests the *grpc_local()* function. This function will return false during task discovery and true during execution.

There is one exception to this rule, when the client code directly affects any aspect of the task graph. For example, if a variable is updated on the client that determines which remote tasks get executed or the size of inputs/outputs of any task, then the operations on this variable should not be encapsulated by *grpc_local()*. If any variables or structures are updated during the task discovery cycle then they should be restored to their original values before the execution cycle begins.

Example:

In the example in table A.2 the variable x determines which tasks get executed and therefore any computation on x should not be encapsulated by the *grpc_local()*. However the variable y does not affect the task graph therefore computations on y should be encapsulated by the *grpc_local()*.

Table A.2. Example implemention of the grpc_local() function

```
grpc_function_handle_t*handles=
   (grpc_function_handle_t*)calloc(iters, sizeof(grpc_function_handle_t ));

grpc_sessionid_t * sessionIDs=( grpc_sessionid_t *)calloc(iters, sizeof(grpc_sessionid_t));
int * status=(int *)calloc(iters, sizeof(int));

 for(i=0; i<iters; i++){
     if(grpc_function_handle_default(&handle[i], "dgesv") != GRPC_NO_ERROR){
       fprintf(stderr,"Error creating function handle");
       die(EXIT_FAILURE);
   }
}

...
...

 x_old=x;
 grpc_map("ex_map"){
    \\Reset updated variable
    x_old=x;
    for(i=0;i<iters;i++){
       x=func1(x);
      if( x==1){
        status[i]=grpc_call_async(&handle[i], &sessionID[i], .., .., ..);
        if( grpc_local()){
          y=func2();
        }
      }
   }
}

.....
.....

 for(i=0;i<iters;i++){
   if(grpc_function_handle_destruct(handle[i]) = GRPC_NO_ERROR) {
      fprintf(stderr,"Error destroying function handle1");
      die(EXIT_FAILURE);
   }
}
```

Appendix B

Appdx B - Hydropad Manual

In this section we show how to install and run the Hydropad application.This application can be downloaded on the Heterogeneous Computing Laboratory website. http://hcl.ucd.ie/project/SmartGridSolve

B.1 Introduction to Hydropad

Hydropad is an astrophysical application that simulates the evolution of clusters of galaxies in the Universe. Hydropad is a cosmological application, originally written by Claudio Gheller, which simulates the evolution of clusters of galaxies in the Universe. The cosmological model that this application is based on has the assumption that the Universe is composed of two different kinds of matter. The first is baryonic matter, which is directly observed and forms all bright objects. The second is dark matter, which is theorised to account for most of the gravitational mass in the Universe. The evolution of this system can be simulated by examining the mutual interaction between these components which is regulated by a gravitational component.

The Hydropad application is composed of two parts: the initialisation part which initialises the initial state of the Universe and the evolution part. The evolution part of the application consists of a number of iterations that simulate the discrete time steps used to represent the evolution of the Universe from the Big Bang to present time. This part consists of three tasks: the gravitational

task the dark matter task and the baryonic matter task. For every time step in the evolution of the Universe, the gravitational task generates the gravitational field using the density of the two matters calculated in the previous time step. Hence the dark and baryonic tasks use the newly produced gravitational forces to calculate the movement of the matter that happens during this time step. Then the new density is generated and the lapse of time in the next time step is calculated from it. The dark matter task and baryonic matter task are independent of each other.

B.2 Installing Hydropad

The installation procedure in Hydropad uses the GNU auto-tools (autoconf, automake and libtools) and Makefile. The auto-tools generates a *configure* shell script that automatically check if the computer contains all the necessary programs and libraries to compile the application. At this point of development Hydropad was tested only in a x86 platform with a Linux environment. Hydropad computational code is written in Fortran 90 while the kernel is written in C language. To be able to compile Hydropad the host machine needs to have these two compilers installed. Hydropad uses the library FFTW, in the gravitational task, to compute the discrete Fourier transform. The *configure* script checks if this library is installed in the host machine. If this is not the case the package contains a x86 version of the library. Hydropad application is composed of three executable files: *hydropad_seq*, *hydropad_gs*, *hydropad_smart*. The first file is the original Hydropad application, it executes the computation sequentially in a local computer. The other two executables use the GridRPC protocol to compute the different tasks of Hydropad in a Grid environment. The difference between the two files is that the *hydropad_gs* uses the standard GridSolve middleware while the *hydropad_smart* utilises SmartGridSolve extension. The last two executables need to be compiled with the GridSolve and SmartGridSolve libraries, if the host machine does not contain these libraries only the sequential executable will be generated. The configure shell script will automatically check the presence of these libraries. The various computational tasks, to be used inside GridSolve or

SmartGridSolve, need to be compiled using the special GridSolve problem compiler. To compile these tasks the libraries that contains the code of the task and a gsIDL file that describe the task are required. The Hydropad package contains these libraries and gsIDL files to compile these tasks. Furthermore it contains a specific makefile command to simplify the compilation. The installation procedure is composed of the following steps:

Step 1
Retrieve the package `hydropad-1.tar.gz` from the hcl website repository (hcl.ucd.ie) and unpack the files in a local directory with the shell command:

```
$cd /path/local/
$tar zxfv hydropad-1.tar.gz
```

Step 2
Execute the configure shell script to generate the makefile.

```
$./configure --prefix=/path
```

With the argument `--prefix=/path` it is possible to choose the directory to install Hydropad. The default location is `/usr/local` so the user needs to have write permission for this directory to install the application. The GridSolve and SmartGridSolve libraries have to be installed in the machines to generate the respective executables. If the configure script locates the libraries it will print the following message:

```
checking for grpc.h... yes
checking for GridSolve library... yes
checking for GridSolve... yes
checking for gs_smart_clib.h... yes
checking for SmartSolve library... yes
checking for SmartSolve... yes
```

Step 3
Compile the application with the command:

```
$make
```

Step 4

Install the Hydropad libraries and executables with the command:

```
$make install
```

To run the application, the directory that contains the executable files has to be included in the environment variable "$PATH" while the directory that contains the libraries has to be included in the variable "$LD LIBRARY PATH".

Step 5

Compiling the tasks :

```
$make services
```

Other than the three version of the Hydropad executable, the package contains also another application, *grafic*. This application is used to generate the initial computational value and it is utilised by the task *usegrafic*. To be able to run Hydropad with GridSolve/SmartGridSolve, the *grafic* file needs to be installed and the directory that contains it has to be inserted in the environment variable "$PATH". After this procedure the Hydropad application is ready to be executed sequentially or with GridSolve/SmartGridSolve.

The Hydropad application also needs a specific input file that contains the initial value of some of the physics variables used in the simulation. The name of this file is passed at command line by using the specific argument *"—input"*. Table B.1 shows an example of the input file. The most important values are the third and forth one. The third value indicates the number of particles utilised in the dark matter N-Body method while the fourth value indicates the number of cells for grid used for the baryonic matter in the simulation. By increasing these values it increases the percision of the simulation. However more memory will be used and the computation will take a longer amount of time. These values have to be even and the number of particles cannot be more than double the number of cell for side grid. Table B.2 shows example values and the resulting quantities of memory used to run the application.

Table B.1. Hydropad Input File

LCDM Michele Guidolin	**Cosmological Model** Owner of simulation
64	Number of particles (Nparmax = Np^3)
64	Ng Grid side sizes, X Y Z
30.0	Box size (Mpc/h)
0.71	Hubble parameter
0.226	Omega dark matter
0.044	Omega baryonic matter
0.73	Cosmological constant
0.01	Present BM average temperature
0.75	Hydrogen mass fraction
0.948	Long-wave spectral index
0.772	Desired normalisation
314159265	Random number seed (9-digit integer)

Hydropad uses other arguments in the command line to change the behaviour of the application. The arguments are:
–**input** specify input file.
–**notsc** use faster CIC for interpolation instead of slower triangular shaped clouds.
–**nmap** number of cycles to map.
–**nmap** total number of cycles in the simulation.

The following example executes a local sequential computation of Hydropad with the input file of table B.1 and 10 evolution cycles of the simulation.

```
$hydropad_seq --input input.in --cycles 10
```

The following example executes a GridSolve computation of Hydropad and 10 evolution cycles of the simulation.

```
$hydropad_gs --input input.in --cycles 10
```

The following example executes a SmartGridSolve computation of Hydropad and 10 evolution cycles of the simulation.

Table B.2. Example of input values and problem sizes for Hydropad

Np	Ng	Data Size
120	60	73MB
140	80	142MB
160	80	176MB
140	100	242MB
160	100	270MB
180	100	313MB
200	100	340MB
220	120	552MB
240	120	624MB

```
$hydropad_smart --input input.in --cycles 10
```

References

[1] M. Yang, Y. Jiang, L. Wang, and Y. Yang, "High performance computing architectures," *Computers and Electrical Engineering*, 2009. 1

[2] C. Lee and D. Talia, "Grid programming models: Current tools, issues and directions," *Grid Computing: Making the Global Infrastructure a Reality*, p. 555, 2003. 1

[3] B. Nelson, "Remote procedure call," 1981. 1

[4] K. Seymour, H. Nakada, S. Matsuoka, J. Dongarra, C. Lee, and H. Casanova, "Overview of GridRPC: A remote procedure call api for grid computing," *Lecture notes in computer science*, pp. 274–278, 2002. 1, 18

[5] T. Downing, *Java RMI: remote method invocation*. IDG Books Worldwide, Inc. Foster City, CA, USA, 1998. 1

[6] E. Freeman, K. Arnold, and S. Hupfer, *JavaSpaces principles, patterns, and practice*. Addison-Wesley Longman Ltd. Essex, UK, UK, 1999. 1

[7] P. Eugster, P. Felber, R. Guerraoui, and A. Kermarrec, "The many faces of publish/subscribe," *ACM computing Surveys*, vol. 35, no. 2, pp. 114–131, 2003. 1

[8] M. Snir, S. Otto, D. Walker, J. Dongarra, and S. Huss-Lederman, *MPI: The complete reference*. MIT Press Cambridge, MA, USA, 1995. 1

[9] L. Dagum, R. Menon, and S. Inc, "OpenMP: an industry standard API for shared-memory programming," *IEEE Computational Science & Engineering*, vol. 5, no. 1, pp. 46–55, 1998. 1

[10] M. Sato, T. Boku, and D. Takahashi, "OmniRPC: a grid RPC system for parallel programming in cluster and grid environment," *Proceedings of CCGrid2003*, pp. 206–213, 2003. 1

[11] L. Gong, "JXTA: A network programming environment," *IEEE Internet Computing*, pp. 88–95, 2001. 1

[12] I. Foster and C. Kesselman, "Globus: A metacomputing infrastructure toolkit," *International Journal of High Performance Computing Applications*, vol. 11, no. 2, p. 115, 1997. 1

[13] G. Allen, K. Davis, T. Goodale, A. Hutanu, H. Kaiser, T. Kielmann, A. Merzky, R. Vanniewpoort, A. Reinefeld, and F. Schintke, "The grid application toolkit: toward generic and easy application programming interfaces for the grid," *Proceedings of the IEEE*, vol. 93, no. 3, pp. 534–550, 2005. 1

[14] T. Goodale, S. Jha, H. Kaiser, T. Kielmann, P. Kleijer, G. Von Laszewski, C. Lee, A. Merzky, H. Rajic, and J. Shalf, "SAGA: A Simple API for Grid Applications. High-level application programming on the Grid," *Computational Methods in Science and Technology*, vol. 12, no. 1, pp. 7–20, 2006. 1

[15] G. Allen, W. Benger, T. Dramlitsch, T. Goodale, H. Hege, G. Lanfermann, A. Merzky, T. Radke, and E. Seidel, "Cactus grid computing: Review of current development," *Lecture Notes In computer Science*, pp. 817–824, 2001. 1

[16] J. Novotny, M. Russell, and O. Wehrens, "GridSphere: a portal framework for building collaborations," *Concurrency and Computation: Practice & Experience*, vol. 16, no. 5, pp. 503–513, 2004. 1

[17] R. Orfali and D. Harkey, *Client/server programming with Java and CORBA*. John Wiley & Sons, Inc. New York, NY, USA, 1998. 1

[18] G. Von Laszewski, I. Foster, and J. Gawor, "CoG kits: a bridge between commodity distributed computing and high-performance grids," in *Proceedings of the ACM 2000 conference on Java Grande*, pp. 97–106, ACM New York, NY, USA, 2000. 1

[19] A. Natrajan, A. Nguyen-Tuong, M. Humphrey, M. Herrick, B. Clarke, and A. Grimshaw, "The legion grid portal," *Concurrency and Computation: Practice and Experience*, vol. 14, no. 13-15, pp. 1365–1394, 2002. 1

[20] M. Antonioletti, M. Atkinson, R. Baxter, A. Borley, N. Hong, B. Collins, N. Hardman, A. Hume, A. Knox, M. Jackson, *et al.*, "The design and implementation of Grid database services in OGSA-DAI," *Concurrency and Computation: Practice & Experience*, vol. 17, no. 2, pp. 357–376, 2005. 1

[21] A. Birrell and B. Nelson, "Implementing remote procedure calls," *ACM Transactions on Computer Systems (TOCS)*, vol. 2, no. 1, pp. 39–59, 1984. 2

[22] A. YarKhan, K. Seymour, K. Sagi, Z. Shi, and J. Dongarra, "Recent developments in GridSolve," *International Journal of High Performance Computing Applications*, vol. 20, no. 1, p. 131, 2006. 2

[23] Y. Tanaka, H. Nakada, S. Sekiguchi, T. Suzumura, and S. Matsuoka, "Ninf-G: A reference implementation of RPC-based programming middleware for Grid computing," *Journal of Grid Computing*, vol. 1, no. 1, pp. 41–51, 2003. 2, 18, 29

[24] E. Caron and F. Desprez, "DIET: A scalable toolbox to build network enabled servers on the grid," *International Journal of High Performance Computing Applications*, vol. 20, no. 3, p. 335, 2006. 2, 18, 23

[25] T. Brady, M. Guidolin, and A. Lastovetsky, "Experiments with SmartGridSolve: Achieving Higher Performance by Improving the GridRPC Model," in *Proceedings of the 9th IEEE/ACM International Conference on Grid Computing (Grid 2008), Tsukuba, Japan*, vol. 29, 2008. 5, 20

[26] T. Brady, E. Konstantinov, and A. Lastovetsky, "SmartNetSolve: High level programming system for high performance Grid computing," in *Proc. of the 20 th International Parallel and Distributed Symposium (IPDPS 2006), IEEE Computer Society*, 2006. 5, 18

[27] H. Casanova and J. Dongarra, "NetSolve: A network-enabled server for solving computational science problems," *International Journal of High Performance Computing Applications*, vol. 11, no. 3, p. 212, 1997. 5

[28] A. Lastovetsky, "mpC: a multi-paradigm programming language for massively parallel computers," *ACM SIGPLAN Notices*, vol. 31, no. 2, pp. 13–20, 1996. 6

[29] A. Lastovetsky and R. Reddy, "HeteroMPI: Towards a message-passing library for heterogeneous networks of computers," *Journal of Parallel and Distributed Computing*, vol. 66, no. 2, pp. 197–220, 2006. 6

[30] H. Casanova and J. Dongarra, "NetSolve: A network-enabled server for solving computational science problems," *International Journal of High Performance Computing Applications*, vol. 11, no. 3, p. 212, 1997. 14

[31] M. Sato, H. Nakada, S. Sekiguchi, S. Matsuoka, U. Nagashima, and H. Takagi, "Ninf: A network based information library for global world-wide computing infrastructure," *Lecture Notes in Computer Science*, vol. 1225, pp. 491–502, 1997. 14

[32] H. Casanova, M. Kim, J. Plank, and J. Dongarra, "Adaptive scheduling for task farming with grid middleware," *International Journal of High Performance Computing Applications*, vol. 13, no. 3, p. 231, 1999. 15

[33] D. Arnold and J. Dongarra, "The netsolve environment: Progressing towards the seamless grid," in *Parallel Processing, 2000. Proceedings. 2000 International Workshops on*, pp. 199–206, 2000. 15

[34] F. Desprez and E. Jeannot, "Add data persistence and redistribution to netsolve," *LIP, ENS Lyon, Tech. Rep*, 2002. 16

[35] F. Desprez, E. Jeannot, I. LIP, and F. Lyon, "Improving the GridRPC model with data persistence and redistribution," in *Parallel and Distributed Computing, 2004. Third International Symposium on/Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks, 2004. Third International Workshop on*, pp. 193–200, 2004. 17

[36] B. Del-Fabbro, D. Laiymani, J. Nicod, and L. Philippe, "Data management in grid applications providers," in *Distributed Frameworks for Multimedia Applications, 2005. DFMA'05. First International Conference on*, pp. 315–322, 2005. 17

[37] E. Caron, "Managing data persistence in network enabled servers," *Scientific Programming*, vol. 13, no. 4, pp. 333–354, 2005. 17

[38] D. Arnold, H. Casanova, and J. Dongarra, "Innovations of the NetSolve grid computing system," *Concurrency and Computation: Practice and Experience*, vol. 14, no. 13-15, pp. 1457–1479, 2002. 17

[39] K. Seymour, A. YarKhan, S. Agrawal, and J. Dongarra, "Netsolve: Grid enabling scientific computing environments," *Grid Computing and New Frontiers of High Performance Processing*, 2005. 18

[40] I. Foster and C. Kesselman, "Globus: A toolkit-based grid architecture," *The Grid: Blueprint for a New Computing Infrastructure, Morgan Kaufmann*, vol. 259, p. 278, 1999. 18

[41] M. Guidolin and A. Lastovetsky, "ADL: An Algorithm Definition Language for SmartGridSolve," in *Grid Computing, 2008 9th IEEE/ACM International Conference on*, pp. 322–327, 2008. 18

[42] A. Lastovetsky, X. Zuo, and P. Zhao, "A Non-intrusive and Incremental Approach to Enabling Direct Communications in RPC-Based Grid Programming Systems," *Lecture Notes In Computer Science*, vol. 3993, p. 1008, 2006. 19

[43] Y. Tanimura, H. Nakada, Y. Tanaka, and S. Sekiguchi, "Design and implementation of distributed task sequencing on gridrpc," in *CIT06: Proceedings of the Sixth IEEE International Conference on Computer and Information Technology (CIT06)*, p. 67, 2006. 19

[44] A. Amar, R. Bolze, A. Bouteiller, A. Chis, Y. Caniou, E. Caron, P. Chouhan, G. Le Mahec, H. Dail, B. Depardon, *et al.*, "DIET: New developments and recent results," *Lecture Notes in Computer Science*, vol. 4375, p. 150, 2007. 20

[45] E. Loureiro, "All in one Graphical Tool for the management of DIET a GridRPC Middleware," 2008. 20

[46] S. Fitzgerald, I. Foster, C. Kesselman, G. Von Laszewski, W. Smith, and S. Tuecke, "A directory service for configuring high-performance distributed computations," in *Proc. 6th IEEE Symp. on High Performance Distributed Computing*, vol. 365, p. 375, 1997. 22

[47] Y. Tanaka, H. Nakada, S. Sekiguchi, T. Suzumura, and S. Matsuoka, "Ninf-G: A reference implementation of RPC-based programming middleware for Grid computing," *Journal of Grid Computing*, vol. 1, no. 1, pp. 41–51, 2003. 22

[48] J. Schopf, M. DArcy, N. Miller, L. Pearlman, I. Foster, and C. Kesselman, "Monitoring and discovery in a web services framework: Functionality and performance of the globus toolkits mds4," *Preprint ANL/MCS-P1248-0405, Argonne National Laboratory, Argonne, IL*, 2005. 23

[49] R. Higgins and A. Lastovetsky, "Managing the construction and use of functional performance models in a grid environment," in *Proceedings of the 23rd International Parallel and Distributed Symposium (IPDPS2009)*, 2008. 29

[50] A. Lastovetsky, R. Reddy, and R. Higgins, "Building the functional performance model of a processor," in *Proceedings of the 2006 ACM symposium on Applied computing*, pp. 746–753, ACM New York, NY, USA, 2006. 29

[51] R. Wolski, N. Spring, and J. Hayes, "The network weather service: A distributed resource performance forecasting service for metacomputing," *Future Generation Computer Systems*, vol. 15, no. 5-6, pp. 757–768, 1999. 29

[52] Y. Caniou and E. Jeannot, "Experimental study of multi-criteria scheduling heuristics for GridRPC systems," *Lecture notes in computer science*, pp. 1048–1055, 2004. 29

[53] T. Braun, H. Siegel, N. Beck, L. Boloni, M. Maheswaran, A. Reuther, J. Robertson, M. Theys, B. Yao, D. Hensgen, *et al.*, "A comparison of eleven static heuristics for mapping a class of independent tasks onto heterogeneous distributed computing systems," *Journal of Parallel and Distributed Computing*, vol. 61, no. 6, pp. 810–837, 2001. 30

[54] J. Dongarra, "The LINPACK benchmark: An explanation," *Evaluating Supercomputers*, pp. 1–21. 39

[55] D. Arnold, H. Casanova, and J. Dongarra, "Innovations of the NetSolve grid computing system," *Concurrency and Computation: Practice and Experience*, vol. 14, no. 13-15, pp. 1457–1479, 2002. 39

[56] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, *et al.*, *LAPACK Users' guide*. Society for Industrial Mathematics, 1999. 43

[57] D. Ryu, J. Ostriker, H. Kang, and R. Cen, "A cosmological hydrodynamic code based on the total variation diminishing scheme," *Astrophysical Journal*, vol. 414, pp. 1–1, 1993. 66

[58] M. Guidolin and A. Lastovetsky, "Hydropad: a Scientific Application for Benchmarking GridRPC-Based Programming Systems," *Proceedings of the 23rd International Parallel and Distributed Symposium (IPDPS2009)*, 2009. 66

[59] R. Hockney and J. Eastwood, *Computer simulation using particles*. Institute of Physics Publishing, 1988. 67

[60] P. Colella and P. Woodward, "The piecewise parabolic method(PPM) for gas-dynamical simulations," *Journal of computational physics*, vol. 54, no. 1, pp. 174–201, 1984. 68

[61] E. Bertschinger, "COSMICS: cosmological initial conditions and microwave anisotropy codes," *Arxiv preprint astro-ph/9506070*, 1995. 69