

The GRID: Computational and Data Resource Sharing in Engineering Optimisation and Design Search

Cox SJ, Fairman, MJ, Xue, G, Wason, JL, ⁺Keane, AJ
Department of Electronics and Computer Science
⁺*School of Engineering Sciences*
University of Southampton

{sc, mjf00r, gx00r, jlw98r}@ecs.soton.ac.uk, ajk@soton.ac.uk

Abstract

We present our GRID architecture for engineering optimisation and design search. Our system will allow seamless access to an intelligent knowledge repository, a state-of-the-art collection of optimisation and search tools, industrial strength analysis codes, and distributed computing and data resources. We focus on the underlying open standards technologies required to implement the system, and give exemplars of how they are being exploited at present. Of particular importance is the interchange of data throughout the system, for which we have adopted W3C standards (e.g. XML), and the ability to link together each of the components in the form of web services. In particular we demonstrate how Condor may be offered as a web service.

1. Introduction

Engineering optimisation and design search is the *process* whereby existing engineering modelling and analysis capabilities are exploited to yield improved designs. In the next 2-5 years intelligent search tools will become a vital component of all engineering design systems. Such facilities will steer the user through the process of setting up, executing and post-processing design search and optimisation activities in a variety of disciplines.

A major driving force in these developments is the need to allow distributed design teams from multiple enterprises to access design and analysis capabilities via the GRID. The GRID has been characterised by a 3 layer model which consists of (i) a computation / data grid, (ii) an information grid, and (iii) a knowledge grid.

In this paper we focus on the technologies which we are using to implement our grid-based system. In section 2 we introduce the Grid architecture for our optimisation system. Section 3 discusses the open standards technologies which we are using. We discuss the computation, data, information and knowledge layers of the grid in sections 4-6, with particular reference to offering the Condor System as a web service and draw our conclusions in section 7.

2. Optimisation and Design Search Grid

2.1. Architecture

Figure 1 shows our grid-based system. In a typical scenario of, for example, wing design, an engineer may couple together Computer Aided Design (CAD) tools, analysis codes for Computational Fluid Dynamics (CFD), or Finite Element Analysis (FEA), and tools for optimisation. The resulting sequence of computations may be performed on local or national machines, or on pay-per-use internet cluster resources e.g. [1].

One important new component in the higher levels of the GRID is the application of database technology, which has conventionally only been used to store information about the products designed using the system. Key questions in the process of, say a typical wing design optimisation, are: What previous designs have been explored and how can I develop them further? Which optimisation strategies are likely to prove effective? Which computational resources have the analysis codes I require, and when will my results be ready? Here there are new requirements for information and knowledge to be extracted from automatically generated databases. Encapsulating and exploiting knowledge at all levels of the grid will enable new designs to be developed more rapidly, or at lower cost. Our system therefore includes a knowledge repository which may be queried.

The architecture of the system consists of four main components which we now discuss in turn.

2.2. Portal

The portal is the web-based point of access, which allows engineers to locate and combine the services they require, giving advice as necessary. It provides grid-based seamless access to an intelligent knowledge repository, a state-of-the-art collection of optimisation and search tools, industrial strength analysis codes, and distributed computing and data resources. It provides the information capabilities of a system like Tambis [2] and has the following roles:

- The engineer interacts with the system through the service portal.

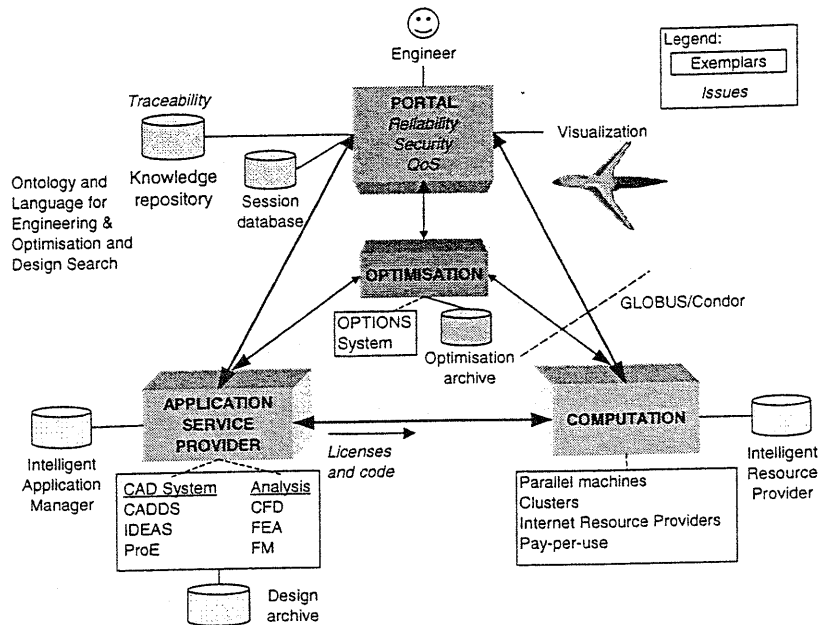


Figure 1: Engineering Optimisation and Design Search Architecture

- To advise the engineer which optimisations work well with which applications on which resources.
- It houses an intelligent knowledge repository to determine which web service databases contain the information required to answer questions it is posed. The knowledge base is built up over time in a manner similar to that used to cache queries on web search engines.
- To provide *traceability* and *security* for results produced on the system.
- To negotiate *quality of service* (QoS) and *reliability* guarantees with other service providers.

2.3. Application Service Provider

The engineer requires access to both design and analysis tools, supported by a database which will provide information about previous designs. Whilst some tools may run adequately as a user application across a server, commercial compute-intensive analysis codes should be also available on a pay-per-use basis to run on external computational facilities provided by the resource provider. An exemplar of this was developed at Southampton [3].

2.4. Optimisation service provider

The optimisation technology is provided by the OPTIONS [4] system, which has been continually developed over the last 15 years. OPTIONS is used for design optimisation and includes a variety (over 40 at present) of different optimisation algorithms.

At the start of the optimisation process, the engineer provides initialisation information— in particular the analysis codes to be coupled together, the permitted methods by which a design may be modified, and an objective function by which each design may be evaluated. The optimisation

service then coordinates the maximization (or minimization) of the objective function to improve the design.

In our current implementation, the optimisation service has a knowledge repository of recommended settings to use for each optimisation control variable. Depending on the level of automation, the optimisation service may automatically enter or suggest to the engineer suitable values to use. The optimisation service then retrieves the required program from the application provider and executes it on computation resources provided by the resource provider.

The optimisation process is divided into several 'stages', which require single or multiple evaluations of the objective function. Each stage is a transactional unit and after successful completion the optimisation's related state and current data is automatically archived. Our framework fits into the model of web orientated application architectures such as Sun Microsystem's Java Enterprise Server where each optimisation code would be an Enterprise Java Bean. Since each step is a transactional unit we can provide consistency and durability and allow for the possibility of multiple independent objective function evaluations to be made in a naturally parallel way.

Furthermore since all data is stored along with session identity information it is flexible enough to meet the requirements of different requesters and can deal with multiple requests at the same time.

2.5. Resource Provider

The resource provider's main role is to execute code to return the value of the objective function. It has a database of computational resources which it provides and may need to negotiate licences for commercial code from the application service provider.

3. Open Standards Technology

Engineering design optimisation requires various computational and data resources to be integrated. To achieve this, technical obstacles brought about by the differences between system environments, software and programming languages must be overcome. One solution to this problem is to wrap various computer systems in a common interface, which 'talks' with the others in a standard format.

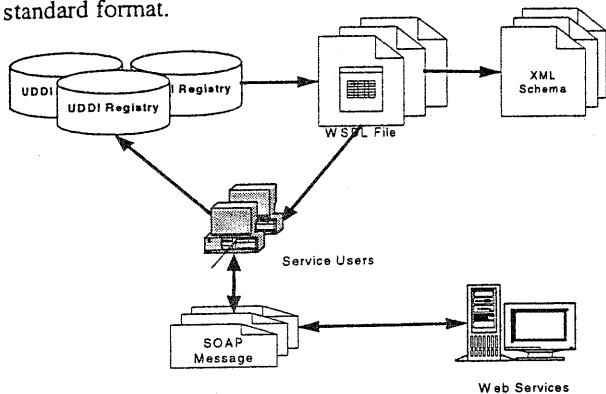


Figure 2: Use of XML technologies in web services

The recent progress in several XML-based open standard technologies [5] has brought major advances in this approach. These contrast with traditional Java/ Jini based approaches by adopting very simple lightweight protocols. Figure 2 shows how the technologies discussed in the next sections are related to each other.

3.1. XML and XML Schema

XML provides a method to contain data of structured format in plain text. Data in the form of XML are transferable between almost any computer system and software application that have the ability to read plain text. XML, as plain text, also works well with existing Internet Protocols like HTTP or SMTP. Thus XML is an ideal medium for communication between different computer systems. Furthermore security can be provided by tunnelling the XML through a secure layer such as SSL [6].

XML Schema is a W3C alternative specification to the Document Type Definition (DTD) for describing the structure of an XML instance document. Compared to the DTD, XML Schema is both more powerful, incorporating a set of rich data types, and more flexible, supporting sophisticated user-defined data structures. Using XML Schema: (i) it is possible to use XML to reflect the complexity of the existing computer resources, and (ii) validation of data is transferred from the application role to the parser, thus making the separation between data and application logic possible.

3.2. SOAP and XML Protocol

The Simple Object Access Protocol (SOAP) [7], provides a simple and extensible framework to define how an XML message is structured. It is a lightweight protocol

for the exchange of information in a decentralized, distributed environment and is easily carried via various Internet protocols, such as HTTP and SMTP. SOAP allows security systems like a firewall to identify an XML message without needing to understand its contents, thus it is feasible to prevent the blocking of unknown HTTP requests. SOAP also provides rich semantics for indicating encoding style, array structure, and data types. SOAP is currently under inspection by the W3C consortium and is a prototype of the future XML Protocol.

3.3. Web Services and WSDL

With XML Schema, SOAP and their supporting software, it is feasible to wrap various computer resources into similar XML-interfaced web components, which are called web services. These communicate with each other using XML messages. To access a web service properly, description about the service interface is required.

Web Service Description Language (WSDL) is one of the emerging technologies for this purpose. WSDL, written in XML, describes a web service as a set of endpoints. Each endpoint consists of several operations and is bound to a specific network protocol (e.g. HTTP or SMTP). The input and output messages are defined for an operation and XML Schema is used in the WSDL file to provide the data types and structures to define the messages. Together with the XML Schema, it provides a complete definition for the interface of a web service and allows programmatical access to the service, in the manner of an API. Tasks like data requests or execution of a piece of code are then performed by sending and receiving XML messages using the SOAP mechanism. For the publishing, discovery and integration of a web service, additional technologies are needed.

3.4. UDDI

The Universal Description, Discovery and Integration (UDDI) specification defines a way to publish and discover information about web services [8]. It is a collaboration between Ariba, IBM and Microsoft who each provide UDDI services. The UDDI project includes a UDDI business registry and a set of operations on it. The UDDI registry, an XML file, identifies a web service and provides information about the service. Other programs use the registry to get the information about the web service and check compatibility with it. Categorisation of web services in the registry enables location and discovery. UDDI together with WSDL, provides the ability to locate and programmatically interface to the web service. This allows the web service to be used in a program in a similar way to a software component, and hence simplifies the service collaboration.

We now discuss the levels of the grid in more detail.

4. GRID: Computation

This section considers the concept of resource management and computational resources. Condor [9] is used as an example of an existing resource sharing system. It

is chosen because its concepts and architecture are well suited to the GRID model and are applicable to other resource sharing systems. We demonstrate how it can be wrapped and offered as a generic web service.

4.1. Condor Introduction

Condor is a software system that creates a High-Throughput Computing (HTC) environment by harnessing the power of UNIX and NT clusters and workstations. It can manage dedicated clusters (in a similar way to Mosix) however its main appeal is that it can make use of pre-existing resources which may be computers sitting on people's desks.

When jobs are submitted to Condor it finds an available machine in its organisation's pool to run the job. Machines become available once they have been idle for a specified period for example when the owner goes to lunch. Jobs are migrated over the network to the machine. If the machine becomes unavailable and the job has not finished (such as if a user returns to their desk after lunch), Condor checkpoints it and either migrates the job to another machine or queues it to disk until a machine becomes free.

4.2. Resource Management

Condor provides a distributed resource management system, which matches the resource consumer's needs with the resource owners. Instead of attaching properties to a job queue directly, Condor implements a publishing system where each machine advertises their resource properties called ClassAds. Each submitted job specifies a resource request ad, which defines the required and preferred run properties of the job. Condor performs brokering by matching and ranking the ClassAds with a job's resource requests.

4.3. Condor and the GRID

Whilst we treat Condor as fundamentally a resource provider at the lowest level of the GRID, its overall architecture fits into the component layers of the GRID model (Figure 3).




Figure 3: Condor GRID Model

If we view the model from the bottom up, at the computation resource level are the machines in Condor's pool with data and control between machines and the Condor control daemons being performed through Remote Procedure Calls. Each machine's ClassAds publishing provides resource description and discovery at the

information layer. Condor provides the resource integration of the information layer with its resource management and migration. At the knowledge layer is the resource manager, which performs sophisticated resource brokering.

There are however some fundamental differences between Condor and the GRID. As a project it began development before the concept of the GRID came into existence. Therefore some parts of the system (e.g the communication system), use older technologies (RPC) and Condor uses its own proprietary systems for resource description, discovery and integration. Furthermore, unlike the components in our GRID architecture, Condor does not expose a programmatic interface as a web service: it is more UNIX-like in that interaction between users and Condor is through the command line.

4.4. Condor as a Web Service

Condor can be integrated into the GRID because of its clear separation of its computational and information layers. We now show how it is possible to expose Condor's resource management ability programmatically using Microsoft's .NET framework [10] to wrap Condor as a C# web service (Figure 4).

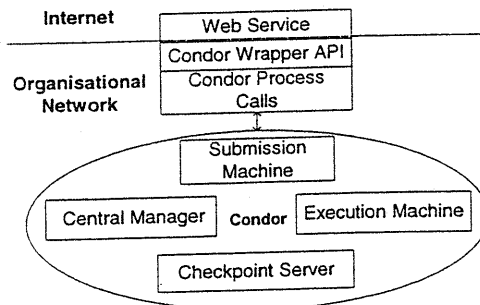


Figure 4: Condor Web Service

4.5. Legacy System Integration

The two most important Condor programs are `condor_submit` and `condor_status`, which respectively perform submission of jobs and provide the user with the ability to monitor and query the machine pool. Other commands provide control, monitoring and querying of submitted jobs.

It is possible to launch any command line executable as a process from another program. Input to the command line executable is through the standard input, read files, pipes, and process start properties. Output is retrieved from the standard output, standard error, written files, and pipes. Construction of an API is then achieved using standard IO and process function calls.

Code wrapped around Condor by modelling each of its programs as a process facilitates the creation of an API. It then becomes possible to make use of these APIs in the web service code to expose a programmatic interface to Condor on the Internet.

ClassAds and Submission Description files (Job Requirement descriptors) are flat files, which are translatable to and from XML documents and have structures that are directly describable as XML Schema documents, as shown in Figure 5.

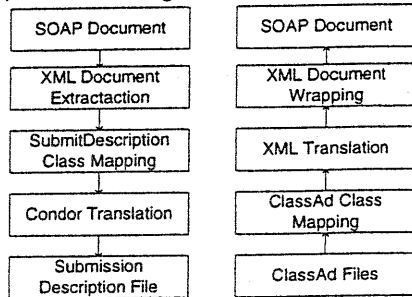


Figure 5: XML Wrapping Process

WSDL provides a standard way for the Condor web service to describe its abilities: containing information about its functionality, method of interaction (SOAP), job sessions, and the structure of the data types it exchanges.

4.6. Computation Issues

At the knowledge level sits the resource manager, which provides a brokering service to match resource requests with resource offers. However, now that each computational resource is separate from any particular resource manager, a user or other GRID service may wish to access the computational resource directly or in combination with the resource manager. This creates an environment where resource managers and computational resources compete to provide superior services and lower costs.

We have not addressed support for check pointing and migration. If computational resources are to be shared among resource managers, then it is necessary to define a standard mechanism for this. Who does the linking, the resource manager or the computation resource, or neither? What about code security and how to prevent malicious or bad code from effecting a computational resource or interfering with another's code? A solution to these issues may be to use a platform independent programming solution such as Java [11] or the Microsoft Common Language Runtime System [12], which both run code in a sand box environment with a definable access policy.

5. GRID: Data

5.1. Metadata

Data is an essential part of the optimisation and design search process. Integration of the Grid relies heavily on the data exchanged between various computer systems. The web services use XML for data exchange and XML Schema, which uses XML itself, to define the types and structures of XML data. It provides rich support for basic data types like integer and string as well as common data structures available in computer programming languages. It is also

possible to construct user defined data formats, such as postcodes.

The flexibility of XML enables the XML Schema to support the sophisticated data structures necessary to define complex user types for web services. As an example, a reply to the Condor_Status operation is to bring a list of "ClassAd"s consisting of two attributes: "Arch" and "OpSys", corresponding to the system architecture and operating system, which are assigned specific values. A fragment of the XML Schema describing the message is shown in Figure 6.

Another advantage of using XML Schema as the data type and structure information provider is that XML Schema works with the XML parser, rather than the application. This provides a separation between the application logic and the data logic, which simplifies (and reduces) any maintenance required as a result of changes on either the data or application side.

XML Schema is designed to be reusable. A new piece of schema can be built based on several existing ones. This brings more efficiency and basic elements for a knowledge system. There have already been several services for publishing XML Schemas on the Internet, (such as BizTalk.org).

```

<?xml version="1.0" encoding="UTF-8"?>

<xsd:schema xmlns:xsd="http://www.w3.org/2000/10/XMLSchema"
  elementFormDefault="qualified">
  <!-- The other parts... -->
  <xsd:element name="Condor_State">
    <xsd:complexType>
      <xsd:sequence>
        <xsd:element name="ClassAd" type="ClassAdType" minOccurs="1"/>
      </xsd:sequence>
    </xsd:complexType>
  </xsd:element>

  <xsd:complexType name="ClassAdType">
    <xsd:sequence>
      <xsd:element name="Arch" type="Architecture" minOccurs="1"
        maxOccurs="1"/>
      <xsd:element name="OpSys" type="OperatingSystem" minOccurs="1"
        maxOccurs="1"/>
    </xsd:sequence>
  </xsd:complexType>

  <xsd:simpleType name="Architecture">
    <xsd:restriction base="xsd:string">
      <xsd:enumeration value="Intel"/>
      <xsd:enumeration value="Sun"/>
      <!-- and so on ... -->
    </xsd:restriction>
  </xsd:simpleType>

  <xsd:simpleType name="OperatingSystem">
    <xsd:restriction base="xsd:string">
      <xsd:enumeration value="Windows"/>
      <xsd:enumeration value="Linux"/>
      <!-- and so on ... -->
    </xsd:restriction>
  </xsd:simpleType>

  <!-- The other parts... -->
</xsd:schema>
  
```

Figure 6: Example XMLSchema Listing

5.2. Role of Databases

It is desirable to database information at each web service for the following reasons:

1. Each web service (resources, application service provision and optimisation) can use the archive itself to make intelligent decisions about its own function.
2. The entry portal may query these databases to build a knowledge archive for the system. Queries may be offered as part of the WSDL specification for the service.
3. Information about users, times, machines, optimisation strategies and results can be transparently logged by the web service as XML conforming to an XML Schema and automatically archived using JDBC.

We have implemented a set of tools for database generation and data storage using an XML Schema with a combination of JDBC, XML parsing, and some conversion rules [13, 14]. If a database does not yet exist for this data it can be generated automatically from the XML Schema. When the structure of the XML Schema changes, for example if a new optimisation method is added; the database must also evolve—we are currently working on this.

As a simple relational database example, the XML Schema fragment in Figure 6 can be translated into a table called `ClassAd` with two columns of type `varchar` called `Arch` and `OpSys`. Each of these columns has a check constraint to ensure the validity of its data values, for example column `Arch` can only contain one of the values 'Intel', 'Sun', etc. However, it must be noted that these check constraints are unlikely to fail, as the parser will already have validated data against the XML Schema before it is automatically archived. The table will also include a column to uniquely identify a record, such as a session ID, from further up in the XML Schema hierarchy.

6. GRID: Information and Knowledge

We have shown that systems like Condor for managing computer resources can be exposed and offered as a web service accessible to applications on the internet. The same open standard methodologies can be used at the information and knowledge layers of the Grid. The information layer adds meaning to data and computation and the ability to access it programmatically. The knowledge layer provides an intelligent problem solving environment to guide the engineer through the process of combining the different components required to perform optimisation and design search efficiently. The intelligent knowledge repository is built by querying the archives of each web service from which, for example, new heuristics about optimisation can be deduced and used to improve future design processes. This may be achieved using formal knowledge elicitation methods and held in rule bases. Design activity also creates domain specific knowledge that designers may wish to archive and reuse in various ways to enhance their design capabilities in the future. Various knowledge management and re-use tools can be used to underpin this process.

7. Conclusions

We have demonstrated that open standards can be used at all layers of the GRID to provide transparent access to applications, optimisation codes, distributed computational resources and knowledge repositories. We view each of the components in the GRID model as web services which may be integrated together to provide, in our case, the framework for an engineering optimisation and design search system. In particular we have demonstrated how the Condor system for exploiting idle compute cycles can be offered as a web service and how automated databasing techniques fit naturally into the web service architecture.

8. References

- [1] <http://www.entropia.com/>
- [2] P.G. Baker, A. Brass, S. Bechhofer, C. Goble, N. Paton, R. Stevens, TAMBIS: Transparent Access to Multiple Bioinformatics Information Sources. An Overview in Proc. Sixth Int. Conf. on Intelligent Systems for Molecular Biology, ISMB98, Montreal, 1998. <http://img.cs.man.ac.uk/tambis/>
- [3] M J Addis, P J Allen, M Surridge. Negotiating for Software Services (2000) Proc. IEEE 11th International Workshop on Database and Expert System Applications p.1039-43 (DISTAL project)
- [4] Keane, A.J. (1999) The Options Design Exploration System – Reference Manual and User Guide, Computational Engineering and Design Centre, University of Southampton. <http://www.soton.ac.uk/~ajk/options.ps>
<http://www.soton.ac.uk/~ajk/options/welcome.html>
- [5] W3C - World Wide Web Consortium (1998). Extensible Markup Language (XML) 1.0. Editors Bray, T., Paoli, J & Sperberg-McQueen, C.M
<http://www.w3.org/TR/REC-xml>
- [6] A. O. Freier, P. Karlton, and P. C. Kocher. The SSL protocol version 3.0. Netscape, Transport Layer Security Working Group, November 1996
<http://home.netscape.com/eng/ss13/draft302.txt>
<http://www.netscape.com/eng/ss13/>
- [7] Simple Object Access Protocol (SOAP) 1.1, W3C Note 08 May 2000. <http://www.w3.org/TR/SOAP/>
- [8] Ariba, IBM, Microsoft. UDDI Executive White Paper, 6th Sept 2000. From: <http://www.uddi.org>
- [9] J Basney and M Livny, Deploying a High Throughput Computing Cluster, in *High Performance Cluster Computing*, Rajkumar Buyya, Editor, Vol. 1, Chapter 5, Prentice Hall PTR, May 1999.
<http://www.cs.wisc.edu/condor/>
- [10] <http://msdn.microsoft.com/net/framework/default.asp>
- [11] Java TM 2 Platform, Standard Edition White Papers
<http://java.sun.com/j2se/1.3/white.html>
- [12] <http://msdn.microsoft.com/library/dotnet/cpguide/cpconwhatiscommonlanguageruntime.htm>
- [13] JL Wason, SJ Cox, and AJ Keane, "Flexible Knowledge Repositories for Problem Solving Environments," pp. 199-205 in Proc. Int. Workshop on Advanced Data Storage / Management Techniques for HPC, ed. R. Allan and K. Kleese, CLRC, Daresbury, ISSN 1362-0223 (2000).
- [14] Sun (2000) JDBC TM Data Access API, Sun Microsystems, Inc. <http://java.sun.com/products/jdbc/index.html>