DISTAL Take-Up D2.5
Transfer to other markets

Report DISTAL Take-Up D2.5 Issue2.0

Matthew Addis

IT Innovation Centre

28 October 2003

Contents

1 Introduction

This document is Deliverable D2.5 of the DISTAL Take-Up Project, which is part of the IST-2000-28221 EUTIST-AMI Cluster.

This document describes how DISTAL can be used to access software and hardware on demand over the Internet, how DISTAL has been applied to the domain of engineering simulations in the aerospace and automotive sector, and suggests other possible uses of the DISTAL technology.

A series of case studies are included that describe some quite specific potential application areas of DISTAL. These are presented to help the reader picture the scope and applicability of DISTAL. These examples have been specifically chosen based on IT Innovation experience and exposure to these application areas.

We believe that there is (or soon will be) a genuine business need for a DISTAL-like solution in these areas.

2 What is DISTAL?

DISTAL is a software system that provides managed, secure and robust access to software and hardware on demand over the Internet. DISTAL allows computationally intensive software applications (for example simulating how a car chassis will deform during a car crash) to be executed remotely by high-performance computing service providers instead of on in-house systems.

Outsourcing of computationally intensive tasks through DISTAL is performed according to a structured business process and data confidentiality is ensured using PKI security techniques.

Use of remote service providers through DISTAL can be an attractive alternative to provision of in-house facilities since up-front investments are lower, peak loads are more manageable, and work can potentially be done to a higher quality, in a shorter timescale, and at a lower cost.

3 Overview

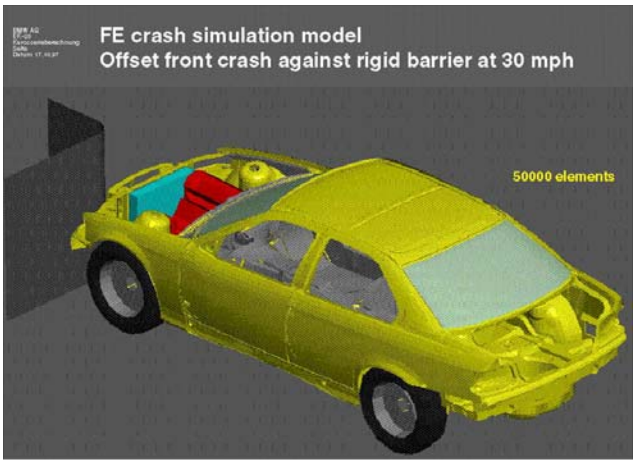

Figure courtesy of BMW AG

Using engineering simulations instead of costly and time-consuming physical experiments can significantly shorten the product development cycle in the aerospace and automotive industry.

For example, it is cheaper, faster and more accurate to simulate car crash behaviour during the design process than to construct and test a large number of prototype vehicles.

Accurate results from numerical simulation methods in engineering require sophisticated analysis of very large-scale and refined models of the real physical problems. For example, reliability of high-tech light weight structures can be investigated by stochastic methods which require execution of hundreds of variants of a single model within 'reasonable engineering time', i.e. within weeks, a few days, or even hours.

Stochastic analysis is typically extremely computationally intensive. For example, the PROMENVIR stochastic case study accumulated 8,000 CPU hours on a 512-CPU Cray T3E for the 128 PAM-CRASH analyses of an offset car crash with a 50,000 finite element model as shown in the car crash picture above. http://62.58.73.21/promenvir/bmw.pdf.

The cost of the in-house hardware and software resources needed to complete simulations in ‘realistic engineering time’ is often unacceptable, even for big companies or organizations like the European Space Research and Technology Centre (ESTEC).

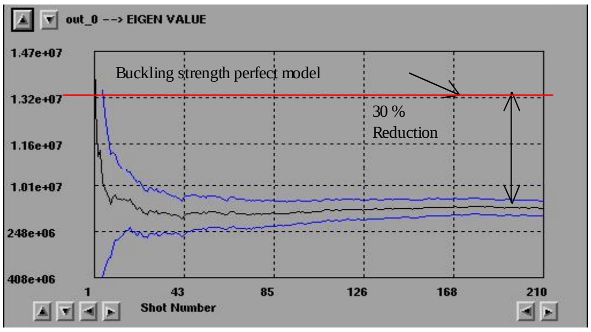

For organisations like ESTEC, outsourcing the execution of large-scale aerospace engineering simulations is considered as an attractive alternative to provision of in-house facilities.

Up-front investment is lower, peak loads are more manageable, and work can be done to a higher quality, in a shorter timescale, and at a lower cost.

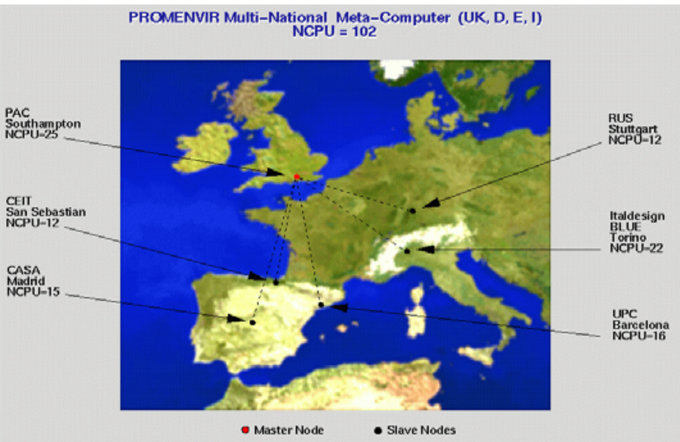

Whilst outsourcing the execution of large-scale engineering simulations at times of peak load according to an ASP model maybe attractive, this is not to say that it is easy to achieve either technically or commercially.

There are many problems to overcome when using the multiple, distributed hardware and software resources to provide the necessary compute power.

The problems associated with on-demand distributed computing via the Internet include the need for:

- Dynamic discovery of remote service providers offering the required applications, hardware and business models;
- Negotiating and agreeing access according to the required level of service;
- Remote execution of the application including secure transfer of large volumes of data;
- Auditing to check for failures which could include networks problems, application errors, incorrect job specification, or insufficient allocation of resources; and
- Attribution of any failures, recovery (if applicable), and then settlement for services rendered.

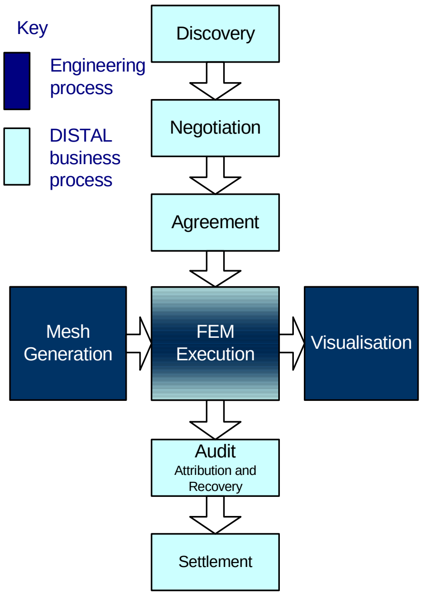

All these issues need to be addressed within an appropriate legal and contractual framework for on-demand business-to-business commerce involving the users, hardware providers, application software vendors and intermediaries such as brokers or trusted third parties.

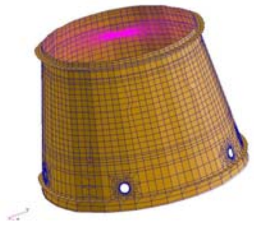

Analysis of effect of geometrical imperfection of a composite thrust cone to its buckling strength

DISTAL Take-Up has proven that it is possible to distribute large-scale industrial engineering problems over the Internet to remote hardware and software providers in a robust, managed and secure way. The DISTAL software covers the registration and selection of appropriate resource providers, negotiation over the terms and conditions of the work, transfer of input and result files, execution of the work, and reporting of work done and resources used.

DISTAL provides a business-to-business solution for software and hardware on demand over the Internet. DISTAL facilitates the dynamic and secure connection of user and resource provider organisations and supports the business process of using one or more of these service providers according to an ASP or outsourcing model.

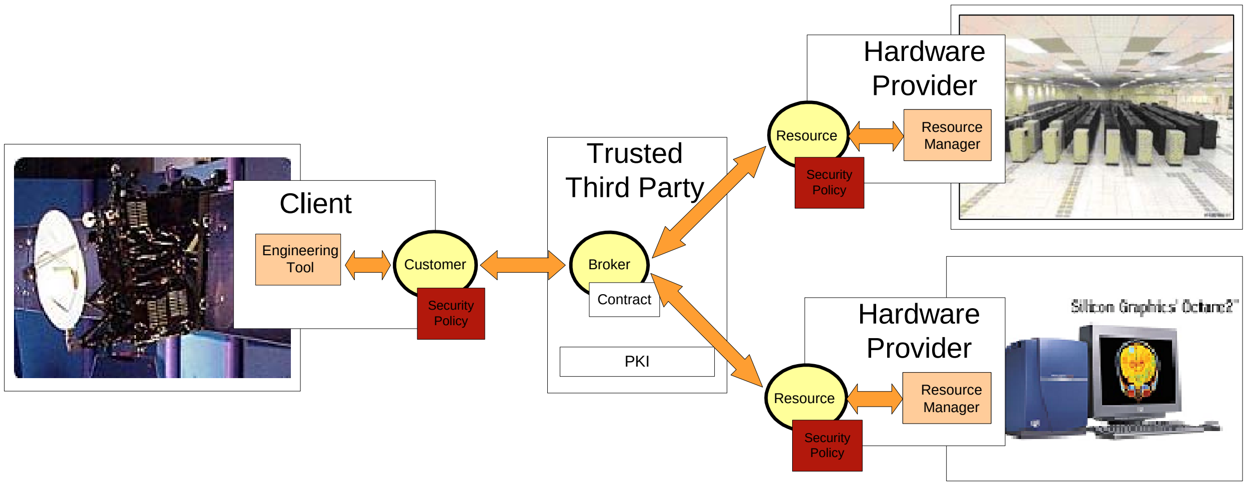

A high level of security is supported, backed by a PKI, to ensure data confidentiality and authentication of all parties involved with no 'holes' required through firewalls for inbound connections. Internal and external resources can be combined, and the system will automatically distribute work according to user requirements with automatic recovery from network problems.

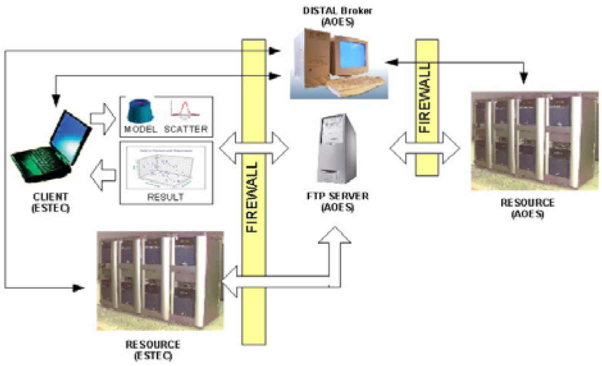

DISTAL is a simple to install and use, and is implemented in Java. A distributed set of software components communicates using standard web protocols to enact the business processes involved in software and hardware on demand over the Internet.

Atos Origin Engineering Services (AOES) and ESTEC have successfully tested DISTAL for large-scale distributed stochastic simulations using finite element analysis.

Web Services	DISTAL
BPEL etc.	DISTAL protocols
UDDI	DISTAL Yellow Pages
WSDL	DISTAL agent interface
SOAP	DISTAL messaging
XML	XML
TCP/IP, HTTP,etc	TCP/IP, HTTP,etc

New applications are easy to support in DISTAL, for example AOES have integrated computational fluid dynamics applications for remote use by ESTEC.

During tests, the overhead of using DISTAL for large engineering simulations was not significantly higher than using local resources.

The industrial users in DISTAL Take-Up found the software to be robust, secure and easy to install and use.

DISTAL is ideal for stochastic analysis of mechanical systems; large structural, thermal or fluid simulations; automated design optimisation; and multi-physics simulation.

Benefits can include:

- reduced product design and development time which lowers personnel effort & cost
- improved design through more simulations in a given timeframe
- reduced up-front investment in local hardware resources and lower software costs through on-demand pay-per-use licensing.

DISTAL is not restricted to engineering since similar large simulations are also required in the biological, chemical, medical, multi-media and many other research and industrial product development sectors.

4 Benefits of the DISTAL approach

DISTAL allows computationally intensive software applications to be executed on remote service providers and internal resources. Therefore, the usual arguments of for ASP/outsourcing apply:

- No need for a large upfront investment in hardware, software and systems personnel within an organisation since external resources can be used on a pay-as-you go basis instead.

- Good solution during times of peak load, i.e. when in-house resources are sufficient for day-to-day use, but can't cope with peaks generated by unexpected workloads or deadlines.

- Reduced personnel effort and cost, or better product design and quality by execution of more analyses within a given design period and budget, can be achieved by a lowering the end-to-end time for running applications. This is of course the usual better/faster/cheaper argument for using an ASP.

DISTAL is suitable for general ASP provision of non-interactive software applications. However, DISTAL is most suited for applications that benefit from the distribution and execution of large-scale or massively parallel computations.

- DISTAL provides an easy and uniform way to simultaneously use multiple heterogeneous service providers. The user doesn't need to learn how to use new tools and processes to access each service provider and to cost, execute, and monitor their jobs.

- DISTAL provides a managed and secure process for using external resources on a commercial basis.

- DISTAL includes support for firewalls, which makes it easy to access external resources in a secure way. Standard Web protocols (http and ftp) are used without the need to open holes in the firewall for inbound connections. AOES and ESTEC used this feature to deploy and operate the software without changing their firewalls in any way.

- DISTAL is not just for engineering applications; any application that supports a 'batch mode' execution model can be easily integrated into the system.

5 Experiences with DISTAL in the engineering sector

ESTEC located in Noordwijk in the Netherlands have extensively tested DISTAL. ESTEC used DISTAL to access remote resources over the Internet, which were hosted at AOES located in Leiden, which is also in the Netherlands. In addition, ESTEC have tested DISTAL using a second resource within ESTEC in order to make combined use of internal and external resources.

DISTAL was successfully tested for:

- Stochastic analysis using the ST-ORM stochastic analysis tool from EASi and the MSC.Nastran finite element solver. Up to 500 shots per analysis were distributed across the AOES and ESTEC resources. The DISTAL software managed the execution of this large-scale stochastic analysis with no failures.

Table 1 shows the estimated elapsed CPU time for stochastic analyses run at ESTEC. Each analysis consists of 100 shots of MSC.Nastran where three MSC.Nastran jobs are executed simultaneously on three separate processors of a local resource.

Elapsed time using local resources	Stochastic job size		
	Small	Medium	Large
Pre-processing (local)	15 min	75 min	750 min
100 MSC.Nastran shots (remote)	3.000 min	15.000 min	150.000 min
Post-processing (local)	15 min	75 min	750 min
Total elapsed time in minutes	**3.030 min**	**15.150 min**	**151.500 min**
Elapsed time if max 3 shots run in parallel	»35 hrs	»175 hrs	»1750 hrs
Total working days	» 4 days	» 2 weeks	» 20 weeks

Table 1 Computation time of stochastic analysis using current in-house resources. The figures in the table are based on the average elapsed CPU time that an MSC.Nastran job needs on a single processor, i.e. a small job needs around 30 minutes, a medium-sized job needs around 2.5 hours, and a large job needs around 25 hrs.

The elapsed CPU time estimates when using DISTAL ASP and distributed resources are presented in Table 2. The estimates are based on the assumptions that at least 20 MSC.Nastran and 20 ST-ORM licenses are available from multiple resource providers and that hardware performance of the available remote resources is at least the same as the SGI Origin 2000 on which the estimated run times above are based.

Elapsed time using DISTAL ASP	Stochastic job size		
	Small	Medium	Large
Pre-processing (local)	15 min	75 min	750 min
100 MSC.Nastran shots (remote)	150 min	750 min	7500 min
Post-processing (local)	15 min	75 min	750 min
Total elapsed time in minutes	180 min	900 min	9000 min
Total elapsed time in hours	3 hrs	15 hrs	150 hrs
Total working days	» 0.5 days	» 2 days	» 18 days

Table 2 Computation time of stochastic analysis using DISTAL ASP. Comparison of the total elapsed time figures in Table 1 and Table 2 shows that DISTAL ASP can enable a time reduction by a factor of up to approximately 12.

DISTAL was also successfully tested for:

- Large-scale Engineering Applications using MSC.Nastran with up to 25 MB of input data and 195 MB of results data and a computational fluid dynamics test using CFD-ACE from CFDRC, which included a user written routine.

Outsourcing of large-scale engineering applications is helpful in situations where in-house resources are nearly exhausted. For example, at the time of writing, the average load on one of the hardware resources used for computational fluid dynamics at ESTEC was around 80%. DISTAL ASP can provide substantial timesaving through outsourcing of large-scale engineering applications during times of such high load.

Table 3 shows how the estimated job runtime can increase by up to a factor of six for large CFDRC jobs when the hardware resource is loaded by 90 percent. Timesaving by outsourcing these jobs through DISTAL also means savings of personnel efforts and cost. Such savings can only be measured case by case, but are considered to at least balance the extra cost for using the ASP.

CFDRC Application	Average job runtime increase rate during high work load periods		
Normal runtime	small job (one day)	medium sized job (2-7 days)	large job (8 and more days)
Work load 50%	1	1	1,2
Work load 70%	1	1,3	2
Work load 90%	1,5	3	6

Table 3 Average job runtime increases of computational fluid dynamics jobs at ESTEC during peak load periods

Overall, the conclusions of AOES and ESTEC were:

- DISTAL works perfectly through the different firewall settings as they exist at ESTEC and AOES.
- During tests, DISTAL has shown the possibility to reduce computation time by distributing ST-ORM shots to multiple service providers.
- DISTAL is suitable for outsourcing of large-scale engineering applications, e.g. fluid, structure or even complex multi-physics simulations.
- General applications can be integrated and executed using DISTAL, which means almost anything as long as it can be executed in 'batch-mode'.
- Full data security is guaranteed via encrypted data transfer over the Internet.
- The DISTAL client is easy to use due to its Graphical User Interface (GUI).
- The DISTAL resource can utilize third-party resource management software for job management and monitoring (Condor was used in the tests).

However, whilst DISTAL provides a proven technical solution, there are several factors that are currently hampering take-up in the aerospace and automotive engineering sector:

- Current pay-per-use license prices of some Independent Software Vendors (ISVs) are defensively high.
- Previous ASP services such as EASi Engineering Stochastics Excellence Centre (SEC) did not survive because:
    - o New product data is usually confidential, in particular in aerospace and automotive design. Therefore there are concerns about allowing this data outside of the organisation.

- o Stochastic simulations have still not gone mainstream and most of them are not time critical. This means big companies tend to schedule stochastics when sufficient local resources are available, e.g. on weekends or during holidays.

- Customer-owned IT services offer lower charges than ASPs can afford.

From a service provider perspective with regard to outsourcing large-scale engineering applications, previous experiences shows that these services cannot bring big business because:

- They are only required from time to time, e.g. to manage peak loads.

- Customers expect to receive the full service, software and hardware at a price around 1 Euro per CPU hour.

- The administration effort is rather high as compared to the potential income.

Current ISV licensing costs and models are the single most significant factor in preventing take-up of DISTAL (or similar approaches) in the aerospace and automotive sector today. In general, application services are becoming more accepted by the user community, but nowhere near as much by the Independent Software Vendors (ISVs). This is because ISV business is focused on yearly software lease contracts and most of them are not prepared, or even not willing, to provide short term on-demand licenses and support at reasonable conditions.

However, the issues above are economic, social or political and are not problems with the DISTAL software itself, or the way that it operates. Therefore, DISTAL could be an ideal solution for other markets that have the same need for DISTAL-like solutions, but don't share the same barriers as the automotive and aerospace sector.

6 Comparison of DISTAL with similar technologies

The DISTAL technology is similar in many respects to Grid and Web Services. DISTAL essentially is a Grid computing approach that has a technical infrastructure very similar to Web Services.

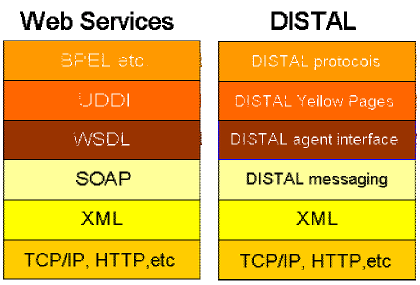

The primary focus of Grid technology is still academic resource sharing, although the Grid 'vision' does extend to dynamic creation and operation of virtual communities [1].

This vision potentially includes commercial collaboration of the form targeted by DISTAL Take-Up. However, there is little work on return on investment, business models, and the end-to-end economics of the Grid and its viability for all the necessary participants.

There is some work on economic models for the Grid, for example GridBus (Grid computing and business) [2] and in the Global Grid Forum working and research groups, for example the Grid Economic Services Architecture working group [3] and the Accounting Models research group. [4]. However, all these activities are research-oriented and often use simplistic models such as auctions with limited regard for more complex discriminants such as business relationships and confidence, and other intangibles.

From a user accessibility perspective, neither web services or Grid technologies come with much support for building vertically focused, easy to use tools since they are primarily middleware technologies. There are relevant PSE efforts, for example the open-source Cactus PSE [5] designed for scientists and engineers, and the SciRun scientific PSE [6] that allows the interactive construction, debugging and steering of large-scale, typically parallel, scientific computations. However, these have limited integration with Grid or Web technologies, and do not provide support for the business transactions involved in outsourcing to third parties when fulfilling a value-chain.

From a services hosting perspective, there are also some efforts towards enabling portal construction and service provision, for example the Grid Portal Development Kit [7]. These typically focus on providing remote access to resource management systems and statistics on resource usage. However, they too do not address the need to support business transactions, business models and processes, and prioritisation between different classes of users.

Web Service offerings have similar limitations. However, this is understandable since the typical application of Web Services is for systems integration either within an organisation or between parties in a well-defined collaboration such as an extranet or supply chain. In these cases, the collaboration between the entities is far less dynamic than required by DISTAL and many of the commercial issues are resolved both in advance and 'out-of-band' of the software systems that are being integrated.

DISTAL addresses these shortfalls. In particular, DISTAL provides:

- An end-to-end business process for outsourcing of software and hardware. The process includes service discovery, provider selection, contract formulation, data transfers, service execution, return of results and settlement.

- Support for license management. The DISTAL architecture is designed to support license procurement, transfer and use. The DISTAL model allows application software to be licensed from the software vendor in a dynamic way so that an application can be executed by a service provider as part of a service that is used by a third-party.

- Business-to-business integration that allows the business models and applications to be easily integrated on a case-by-case basis. Vertically focused applications (such as stochastic analysis using ST-ORM) can be integrated quickly and used according to appropriate business models (such a pay-per-shot) without need to modify the DISTAL infrastructure (security, messaging, business process).

- Integration of PKI security. This provides strong authentication between all parties, confidentiality of all data, and the technical infrastructure for creation of commercially binding contracts between users and suppliers (although this is subject to legal issues).

This is not to say that support for some of these areas won't become available in Grid/Web Service standards and implementations in the future. Indeed, WS-Security [8] provides a promising security model for Web Services. Furthermore, Web Service business process languages [9] can be used for implementing end-to-end business process support. However, these are not yet mature so are not a viable alternative to DISTAL.

Of course, considering the rate at which the technology and standards are progressing, Web Services and Grid technologies may well be viable options to consider when further developing and productizing the DISTAL software.

7 Some potential application areas for DISTAL

The business benefits of the DISTAL approach (presented in Section 4) include:

- Reduced upfront investment in hardware, software and systems personnel within an organisation since external resources can be used on a pay-as-you go basis instead.

- Alternative to in-house systems during times of peak load, i.e. when in-house resources can't cope with peaks generated by unexpected workloads or deadlines.

- Reduced personnel effort and cost, or better product design and quality through lowering of the end-to-end time for running applications

DISTAL has the potential to be applied wherever there is a requirement for large-scale computation to be distributed across a set of independent, third-party service providers in a managed and secure way.

At a technical level, there are many parallels between DISTAL functionality and current/emerging Grid and Web Services solutions. Therefore, it would be easy to list all the applications of Grid technology and then state that DISTAL could be used in these areas. Whilst it may be true that DISTAL could be applied in these areas, such a statement would be ignoring the likelihood of transfer due to the effects of competition, rapidly evolving standards and significant market flux.

In a similar vein, the use of DISTAL purely in an intranet setting is limited since other products already deal with this area, and the use of DISTAL for one-to-one outsourcing to a single supplier may also be limited since HPC portals already provide existing solutions in the marketplace.

Therefore, this section of the document concentrates on some specific (and sometimes very niche) areas where (a) there is potential business benefit from using a DISTAL approach and (b) DISTAL is a good technical solution. The example application areas below have been chosen based on direct IT Innovation exposure to these areas (i.e. we are already working with these people) and where we believe that the barriers to DISTAL take-up experienced in the engineering sector (see Section 5) may not be such an issue.

The overall intention is to help the reader to better understand and imagine the scope and potential for DISTAL, based on some concrete and quantified examples where DISTAL is a good technical fit. It is then up to the reader to assess the business case for using DISTAL.

7.1 Digital film restoration

The UNESCO estimate of the world audiovisual holdings, in archives, is 200 million hours. 50 million hours is European.

Tape-based audiovisual materials (70% of the total) have a shelf-life from 5 to 100 years, depending upon storage conditions. Twenty years is a reasonable life expectancy. Audio tape recordings have been made on a large scale for over fifty years, and videotape for over thirty years. ALL these original recordings are at risk. In addition, players for audiotape and older formats of videotape are obsolescent. Film has a life expectancy of up to 400 years, but older films have problems of colour fading, and acetate-based film can produce acetic acid (vinegar syndrome) after as little as 20 years even in reasonable storage conditions.

Figure 13: Original still image from the travelogue "Een Autotocht in de Pyreneeën"

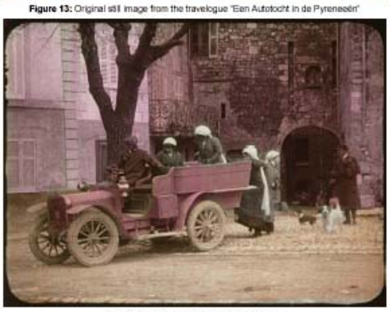

Figure 14: Corresponding restored image

Additionally, new blank film stock, for reprinting or re-mastering existing films, may cease to be produced in as little as 10 years. Indeed, UNESCO estimates that some 2.2 billion meters of nitrate-based film are currently stored at national and international film archives and these films are subject to continual chemical decomposition and have often suffered from improper treatment during their screening lives.

It is the experience of the major audiovisual content-holders that conservation is, ultimately, a dead end. At best it reduces the rate of chemical deterioration. At worst it imprisons content inside obsolete carriers, making access more and more expensive, difficult – and unlikely. Audio and video archives, and their professional bodies, have reached near-universal agreement favouring migration. Whilst an on-going cycle of migration to new analogue media is possible, preservation of legacy audiovisual materials is best achieved by migration to the digital environment. Therefore, massive efforts are underway to digitise the audiovisual assets of small and large content holders alike.

The digitisation and preservation process offers an ideal opportunity to apply restoration techniques to the audiovisual content involved. These techniques are currently evolving fast from analogue to digital ones. Digital restoration already offers techniques without an analogue equivalent, for example motion stabilisation and dust removal. There is still a need of new algorithms to detect and remove image and sound defects, like severe scratches, hair, grain, dye fading, geometrical distortions, echoes, streaking, severe resolution loss, saturation, wow and flutter. However, these are already in development.

The problem of applying computationally intensive digital restoration algorithms to large volumes of digitised audiovisual content is a significant one. One hour of film digitised and uncompressed at HDTV resolution requires 1 terabyte of storage. Digital restoration of this material at only 1 frame per second (1 hour of film will take one day to restore) using existing algorithms requires a high-performance cluster computing solution with upwards of 100 Intel Pentium IV class processors.

All but the largest content-holders will not be able to afford the hardware and software necessary for large-scale digital restoration, especially when restoration needs to be performed in short timescales or on an ad-hoc basis. However, digital restoration can be easily parallelised and executed across multiple, independent compute resources. Furthermore, the choice of restoration techniques and their parameters can be readily investigated locally on a small subset of the data. The resulting 'restoration decision list' can then be applied in batch mode to the full dataset on one or more remote machines.

One potential problem is the movement of digitised content to and from the remote service providers due to the large data volumes and hence bandwidth requirements. In the case of digital restoration of material for long-term archival and preservation, this is indeed a problem since high-resolution is required and only lossless compression will be tolerated. However, there are plenty of restoration scenarios where both lossy compression techniques can be used and where the resolution requirements are lower, for example restoration for screening on television networks. In this case, it is quite feasible to transfer MPEG encoded content to and from the resource providers over current infrastructure. For example, 1 hour of MPEG2 encoded SDTV resolution video can be transferred over a 2Mbit link in only a few hours. Indeed, this method of digital film distribution is already used in cinemas for new film releases where resolution is much higher (HDTV or 4k) which shows that digitised film can be distributed over current network infrastructure.

Therefore, DISTAL should provide a simple solution to distributing and executing digital restoration on a set of remote resources.

7.2 Digital film post-production

A large number of films and commercials use computer graphics animations to implement the special effects that the artists want to depict on the screen. There are many small national and multinational companies that include in their working scope the design of 3D environments. Thus, a variety of software solutions (including environment design and rendering) to produce the final 2D result, are available in the worldwide market. Well-known commercial products are: 3D Studio Max (by Discreet), Maya (by Alias Wavefront) and RenderMan (by Pixar). There are additional software packages that are focused on rendering purposes, for example EXLUNA Entropy (EXLUNA) and MentalRay (Discreet).

The common practice for a 3D based movie company, is to purchase and use an integrated solution for the design of 3D environments, and in some cases where additional rendering features are required, to deploy a standalone rendering software package. The current trends in the worldwide market require the purchase of high-end computer workstations to provide the needed huge amount of computational resources.

Digital post-production methods are both computationally and data intensive. However, many companies who provide digital post-production services find they don't need large-scale computing resources every day, and so cannot justify owning the facilities needed to use digital methods routinely. Furthermore, a complete movie production usually involves numerous artistic and production companies, so it doesn't make sense to hold all the data in-house. A close collaboration between the client and the designer is required in order to have a result close to the client’s needs.

The requirements of an iterative digital post-production process, the need for ad-hoc and sometimes very significant computing power, and the need to complete post-production projects to strict deadlines could be addressed using DISTAL. Rendering could be outsourced to one or more remote resource providers on a pay-as-you go basis at times of peak load. This would allow in-house production of small jobs with urgent and significant jobs being outsourced to external providers. Digital rendering is also well suited to remote execution since the input data is relatively small and for most applications the results can be returned in a compressed form, e.g. MPEG2.

7.3 Multimedia content-based analysis and indexing

Multimedia online resources are becoming increasingly commonplace, especially for 2D digital images. For example, the Hulton Getty archive alone contains over 40 million images, and is an unparalleled resource of unique illustrative material, covering every facet of people's experiences and environment, recording history to the present day. 230 thousand of these images are already online. The National Geographic group has amassed 10.5 million images, stored in its archives. Over 10 thousand of these have now been digitised and many have been made available through Getty Images. Indeed, the total number of on-line images available through Getty Images is now almost half a million. Even specialised archives at the national level can be significant in size, for example the Bildarchiv Foto Marburg contains over 2 million digital images on art and architecture in Germany.

Overall, the volume of digitised images is increasing at a tremendous rate, especially when the digitisation programmes of the national museums and galleries are taken into account.

The search and retrieval of images from such vast repositories presents many problems and is currently only dealt with through the searching of simple textual metadata that describes each image. However, several content-based analysis techniques exist that allow images to be searched and retrieved by colour, texture, brightness, spatial distribution of features and others. Searches are enabled where the user submits an image to the system and requests images to be returned that are similar to the query in some way such as colour or texture and sub-image can be presented with a request to search for the parent image. This provides a particularly powerful way of searching very large collections.

Content-based analysis algorithms can be applied to the images in a collection to generate a set of image-content descriptors called feature vectors.

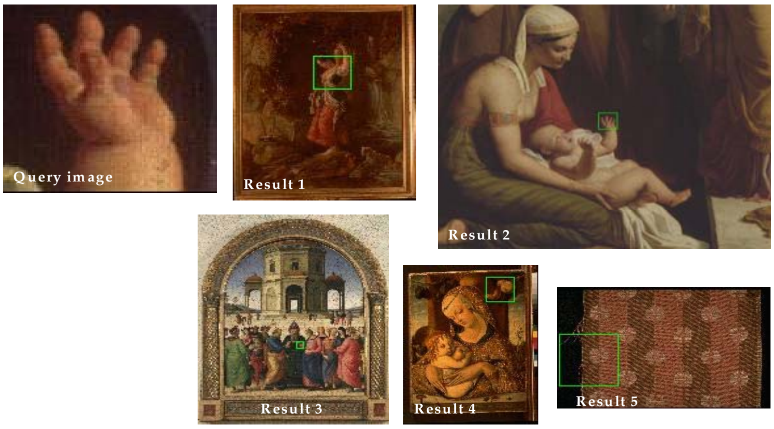

A feature vector can be considered as a way of indexing an image to describe an aspect such as colour distribution or texture. The feature vectors are then integrated and stored with the textual metadata for each image in the collection database.

When a search needs to be made, the required algorithm is run on the query image to create a query feature vector. This query feature vector is then compared with all the corresponding feature vectors for the images in the collection. The comparison of feature vectors results in a measure of distance between the query image and each image in the collection.

The comparison process can be performed very quickly, which is essential for fast search and retrieval. However, the initial indexing process can be very computationally intensive. For example, indexing only a few thousand images can take days on desktop scale hardware. Clearly, the processing requirement for indexing hundreds of thousands of images is huge. The hardware necessary is often beyond the means of the content-holders, especially when it is only used sporadically for batches of images rather than as part of a continuous production process.

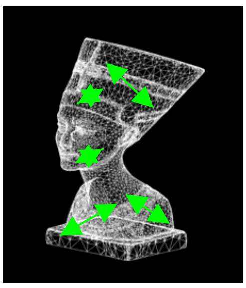

The situation is made much worse for 3D or video content. Digital capture of 3D objects, for example cultural objects held in museum collections, is increasingly common-place and is driven by the need to improve access to cultural heritage through electronic channels such as web sites or gallery information systems. There is also a huge volume of digital audiovisual material (content already produced directly in a digital format + digitised and restored analogue content). Indexing and searching this material presents similar challenges to the still images archives described above.

DISTAL provides a potential solution to the problem of content-based indexing since the digital content to be indexed can be readily distributed and indexed on multiple remote resources. Since the physical storage requirement of the indices created is an order of magnitude smaller than the content being indexed, it is straightforward to transfer an image or video sequence to a remote resource once and then retrieve the set of indices generated for that content as and when appropriate indexing algorithms are available. Furthermore, since content-based analysis often uses relatively coarse aspects of the images, for example colour distribution, then significant amounts of compression can be applied to the images to reduce the network bandwidth requirements during transfer to remote resources.

7.4 Medical

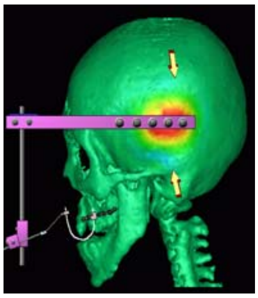

A disproportional growth of mandible and maxilla is amongst the most frequent defects with patients suffering from certain congenital deformations. Often a surgical correction is the method of choice. Medical practitioners in the area of maxillo-facial surgery have to do individual pre-operative planning for each patient to be treated. Numerical tools can allow a doctor to compare different treatment plans and to predict their outcome based on computer simulations. Simulations require the use of high performance computing platforms with sufficient compute power and memory. The scale of resources required is not usually available in a clinical environment.

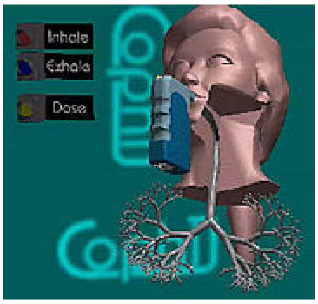

Simulation techniques are often used in inhaled drug delivery and integrates: medical images; mesh generation; computational fluid dynamics (CFD); compartment modelling of the lungs; and simulation of inhalations devices. It is also a comprehensive tool for the study of new lung treatments and drug delivery to the lungs, with a focus on determining the conditions necessary for targeted delivery of medication. The numerical fluid dynamics simulations demand a vast amount of compute resources, creating a possible bottleneck that would limit the use of the simulation process.

Intracranial lesions are routinely treated using stereotactic radiosurgery. This technique uses intense levels of gamma radiation in order to control a tumour or a vascular abnormality within the brain. The use of many sources firing from many angles results in a high radiation dose within the lesion whilst minimising the dose to the surrounding normal brain.

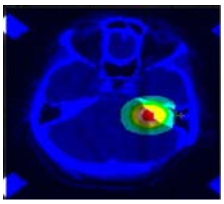

The determination of the radiation dose requires the calculation of the dose from many sources of radiation for a matrix of points positioned within and around the lesion. For each point in the matrix, the dose from each beam must be determined, taking into account the direction of travel and the resultant beam attenuation. Complex treatment plans require the summation of the dose from several treatment positions.

All the examples above involve computationally intensive numerical simulation techniques that typically require significant amounts of hardware to be executed within the required medical or design timescales. This hardware is frequently not available in-house. This type of problem is very similar to large-scale engineering simulations in the automotive and aerospace sector and hence DISTAL should provide a good technical and business solution.

7.5 Combinatorial Chemistry

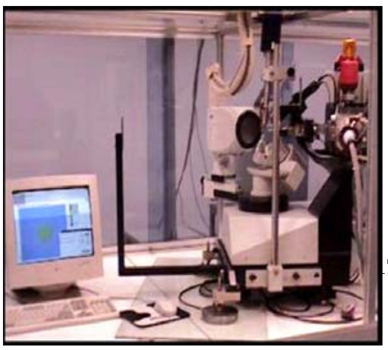

Combinatorial chemistry provides a powerful means of understanding how the molecular structure of materials affects their properties, and of finding new and better compounds ranging from drugs to advanced construction materials.

The combinatorial approach to synthetic chemistry has been made possible by developments in solid phase synthesis. This has led to an enormous explosion in the production of new compounds, usually based upon systematic variation of a given base structure.

The use of this idea in organic synthesis, directed towards the discovery of new pharmaceuticals, has been followed by the application of the "parallel synthetic approach" to the investigation of both inorganic and other organic compounds, for a variety of applications targets, including catalysis and materials properties.

The combinatorial approach involves synthesising and testing families of compounds, finding their molecular structure, and fitting the results into the bigger picture obtained from previously measured compounds. The bottleneck in this process has traditionally been the determination of structure via X-ray crystallography. However, the latest generation of instruments can handle dozens of new samples per month, and with robotic sample handling they will produce hundreds of samples per month. World-wide, hundreds of thousands of new structures will become known each year.

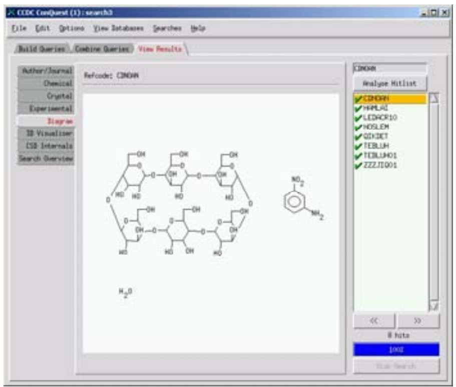

Significant parts of the combinatorial chemistry process are software-based and use computationally intensive techniques, for example chemical and structural property prediction is used both in the design of the combinatorial experiment and during the screening and characterisation stages.

A DISTAL approach is well-suited for outsourcing computationally intensive activities like property prediction within the context of combinatorial chemistry due to the inherently parallel nature of the technique.

The result would be a more complete characterisation database, which would not just provide more results of relevance to the specific combinatorial chemistry experiment being conducted, but would greatly improve the combinatorial chemistry process in general since more data would be available for analysis and comparison as part of future experiments.

7.6 Large-scale engineering simulations

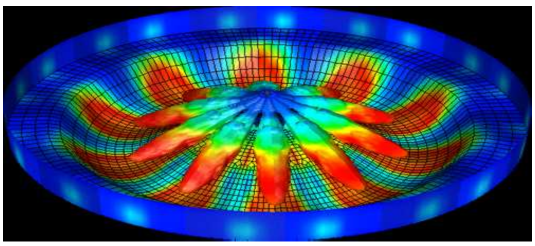

Whilst there are barriers (social, economic, political) to take-up of DISTAL in the automotive and aerospace sectors, the use of computationally intensive numerical simulation techniques such as finite element analysis (FEA) and computational fluid dynamics (CFD) are widely used in other application domains where these barriers may not be present.

This can be seen from the range of applications for CFD-ACE and CFD-FASTRAN from CFRDC. CFD-ACE is used in application areas including biomedical, biotechnology, cavitation, combustion, environment & safety, fuel cells, MEMS & microfluidics, microelectronics, plasma, propulsion, semiconductor processes and turbomachinery.

Full details are available from the CFDRC website http://www.cfdrc.com.

DISTAL is ideal for outsourcing of large-scale simulations in any of these areas and CFD-ACE has already been integrated with the DISTAL software.

7.7 Stochastic analysis

Stochastic analysis is a core application area for DISTAL in the aerospace and automotive sector and has been a driver behind the DISTAL Take-Up project.

DISTAL is ideal for stochastic analysis in engineering since the large number of computationally intensive shots can be readily distributed across independent resource providers. Furthermore, the ST-ORM stochastic analysis tool incorporates techniques to minimise data transfer volumes so that stochastic analysis can be done with little more than the bandwidth needed to execute a single shot at each resource provider. This means that stochastic analysis can be done over relatively low bandwidth connections between user and resource providers.

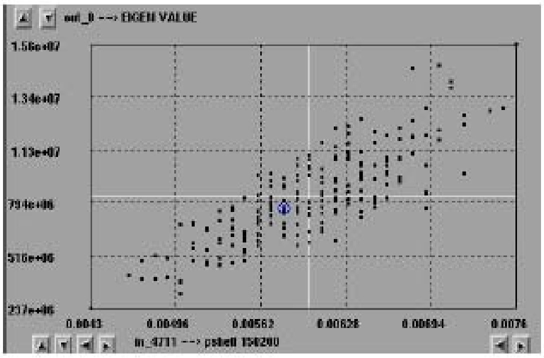

Other application areas of stochastic analysis that are computationally intensive should be a natural fit with DISTAL, especially considering that DISTAL integrates ST-ORM and has already been proven for distributed stochastic analysis in engineering design.

Application areas of stochastic analysis are diverse and include finance, genetics, ecology, chemistry and telecommunications.

8 Conclusions

DISTAL provides a business-to-business solution for software and hardware on demand over the Internet. DISTAL facilitates the dynamic and secure connection of user and resource provider organisations and supports the business process of using one or more of these service providers according to an ASP or outsourcing model.

In engineering design, DISTAL can enable significant cost savings to be achieved by:

- Reducing up-front investment for upgrading of local hardware resources;
- Using medium-sized hard- and software configurations to cover the small and medium-sized jobs; and
- Outsourcing stochastics and large-scale jobs during times of peak load.

Reduction of engineering simulation computation time through use of DISTAL for ASP or outsourcing can:

- Either significantly reduce personnel effort and cost;
- Or allow more analyses within a given design period and budget;
- And, in turn provide a better product design and quality.

In particular, DISTAL offers Small and Medium-Sized Enterprises (SMEs) the possibility to tune their internal resources to their current needs whilst having access to appropriate external resources for covering peak loads and demanding cases.

DISTAL is not restricted to engineering since similar large simulations or computationally intensive applications are also used in the biological, chemical, medical, multi-media and many other research and industrial product development sectors.

9 Contacts and further information

DISTAL software and services
IT Innovation Centre

Dr Matthew Addis
2 Venture Road, Chilworth Science Park
Southampton, SO16 7NP, UK
Phone: +44 23 80 760834
FAX: +44 23 80 760833
Mailto:mja@it-innovation.soton.ac.uk
http://www.it-innovation.soton.ac.uk

DISTAL Take-Up Coordinator
Atos Origin GmbH
Engineering Services

Dr.-Ing. Martin M. Mai
PO Box 80 05 65
D-70505 Stuttgart / Germany
Phone: +49-711-7377-237
FAX: +49-711-7377-108
Mailto:Martin.Mai@atosorigin.com
http://www.distal@62.58.73.21/

DISTAL Take-Up public deliverables, flyers, presentations, and information on the earlier PROMENVIR and DISTAL projects can be accessed from the DISTAL Take-Up Web site: http://www.distal@62.58.73.21/

Examples of IT Innovation's expertise and current involvement in a range of Grid projects is available from http://www.it-innovation.soton.ac.uk/research/grid

10 ACKNOWLEDGEMENTS

The authors, on behalf of the DISTAL Take-Up consortium, would like to thank the responsible project officer, Ms. Annalisa Bogliolo, and the EUTIST-AMI cluster coordinator, Mr. Massimo Busuoli, for their continuous support, and the reviewers, Dr. Bernhard Bauer and Dr. Tony Gore, for all their helpful comments and constructive recommendations which have greatly contributed to the success of the project. The consortium is also grateful to the European Commission for having part-funded the DISTAL Take-Up Activity, to the management of the partnering organisations for having invested in the development of this new technology, and to the application software vendors who have provided free test licenses for the DISTAL robustness tests.

---

1 http://www.globus.org/research/papers/anatomy.pdf
2 http://www.gridbus.org/
3 http://www.doc.ic.ac.uk/~sjn5/GGF/GESA-WG3.htm
4 http://www.gridforum.org/5_ARCH/ACCT.htm
5 http://www.cactuscode.org/
6 http://software.sci.utah.edu/scirun.html
7 http://doesciencegrid.org//projects/GPDK/
8 http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnwssecur/html/securitywhitepaper.asp
9 http://www.ebpml.org

