Interactive Volume Reconstruction and Measurement on the Grid

C. Germain-Renaud[1], A. Osorio[2] , R. Texier[1]

[1]LAL and LRI CNRS-Université Paris-Sud
[2]LIMSI CNRS

Correspondence to :

Cécile Germain-Renaud
LAL
bat. 200 Université Paris-Sud XI
BP 34
F - 91898 Orsay cedex
FRANCE
tél +33 (0)1 64 46 84 02 or 33 (0)1 64 46 83 95
fax +33 (0)1 69 07 94 04
email: germain@lal.in2p3.fr

Summary

1. Objectives

To prove the advantages of integrating grid computing within medical image analysis software, and to discuss the technological, sociological and health care related issues.

2. Methods

Presentation of an instant volume reconstruction and measurement tool (PTM3D) used in clinical practice, including percutaneous nephrolithotomy examples; description of a parallel implementation of volume reconstruction, evaluation of this implementation on lung and body reconstruction, presentation of the technical limitations for clinical use; description and discussion of a prototype grid implementation.

3. Results

Volume reconstruction can broaden its medical scope and use by accessing high-performance computing systems; interactive exploration of medical images can co-exist with the usual batch workload of grid systems; the EGEE grid middleware offers some of the required core services; a fully adequate computing environment needs further evolution to integrate real-time constraints.

4. Conclusions

Clinical experiments of a grid-enabled PTM3D become possible. Widespread adoption of grid technology in the medical images analysis field will benefit from this “early user” project. Convergences do appear between two broadly different fields, High Energy Physics and Medical image, towards the need of a smooth integration of the new resources offered by grid systems into the everyday tools of their respective end-users. It can be expected that the convergence will mature towards truly interactive grids, able to serve the needs of the medical community.

Keywords
Radiology, computing grids, visualization

1. Introduction

Image analysis algorithms can be used in radiological and pathological screening, to identify patients that require additional follow-up, or provide a characterization of disease that minimizes or eliminates inter-observer variability, or at various steps in surgical planning and computer-aided surgery. The health care challenge is to transfer experimental research, first to clinical practice, then to routine clinical practice. Today, classical problems such as segmentation or registration require in some cases high-end parallel computers, do not have a fully reliable automatic solution, and need some degree of user interaction. The task of analyzing large images at a sufficient speed to support interactive use requires substantial computing power. Interactive image analysis is thus potentially a major beneficiary of the grid proposal of transparent access to high-performance computing.

This paper exemplifies the opportunities offered by the advent of grid infrastructures, as well as the new requirements for the grid middleware, on the PTM3D medical images analysis software We show that PTM3D can take advantage of the grid computing power for real-time volume reconstruction, and can be integrated with real-world grids.

Section 2 presents a few medical applications of instant volume reconstruction and measurement. Section 3 gives a brief presentation of PTM3D. Section 4 deals with the dynamic parallelization of the most time-consuming algorithm of PTM3D. Section 5 explores the integration of PTM3D with grid systems. Section 6 discusses the convergences, both at the medical level and at the grid middleware level, towards a better integration of the requirements presented in the paper.

2. Objectives and issues

2.1. Medical objectives

The main objective is 3D reconstruction and instant volume measurement. Volume measurement targets diagnostic, therapy planning and surgery planning.

- Diagnostic and therapy planning: current practice for evaluating treatment efficiency uses the WMO standard, which is based on bi-dimensional measurements; the volume variations are crudely estimated, leading to a ±25% error; a precise volume measurement tool allows a much improved assessment of the efficiency of chemotherapy products and of lesion evolution.
- Surgery planning. A precise measurement of organs (liver, kidney, lung) is required when targeting transplantations or resection for partition surgery.

3D-reconstruction alone targets surgery planning and guided surgery through augmented reality: supplementing real data from some viewpoint, source, date or modality with real data from another such item.

Two use cases related to percutaneous nephrolithotomy (PCNL) are useful to exemplify the clinical situations and the related computing requirements. They are detailed in the next section.

- 3D reconstruction and instant volume measurement of complex and staghorn renal stones, to confirm and facilitate PCNL: 3D reconstruction allows the surgeon to study the stone's shape, which aids in choosing the surgical procedure; also, comparison between in vivo and ex vivo measurement, informs whether the stone extraction was complete, avoiding repeated CT scans.

- An augmented reality system for PCNL: renal stones, kidneys and external skin of the body are reconstructed from CT images; the real-time projection of organs and lesions on the patient in the surgical position leads the surgeon towards the target region and choosing the best route for percutaneous surgery.

2.2. *Computing requirements*

Some of the requirements for a computing system offering 3D reconstruction and instant volume measurement are common to the clinical situations sketched above.

1. Transfer of medical files in DICOM format coming from various sources (TDM, MR, echography) and various proprietary implementations of the DICOM standard.
2. Interactive real-time navigation inside the voxel volume.
3. Interactive real-time visualization, manipulation and 3D-segmentation of radiological images.
4. Interactive real-time generation of a 3D representation of regions of interest.

In this framework, real-time has differentiated meaning. For visualization related actions (navigation, manipulation), true real-time is required, that is the image must smoothly follow the user interaction at the screen frame rate. For generation of 3D representation, waiting a modest amount of time is acceptable; a delay of 2 minutes is generally considered practical by radiologists.

The next section shows that requirements 1, 2 and 3 be achieved on a standard PC. Following the target clinical application, the computing power requirements for the fourth step are quite different. While real-time reconstruction of tumors or calculi can be achieved on a PC, reconstruction of lungs or external skin will very often be too time-consuming to pretend to interactivity on a PC. The problem is not transitory: the increase in processor computing power will be more than counterbalanced by the increase in data size (e.g. multi-slice spiral CT-scan provide more than 1000 slices). We show in section 4 that volume reconstruction is amenable to efficient parallelization, achieving the above-mentioned real-time requirement.

2.3. *Motivation and requirements for a grid-enabled reconstruction*

The motivation is not different from many other medical or non-medical fields: to propose a high-end computing resource that can be accessed freely by any authorized organization, and which provides the appropriate Quality of Service without the burden of maintaining the resource at the individual organization level.

Parallel execution of the volume reconstruction algorithm does not require a grid as absolutely as e.g. analysis of distributed databases of medical images, because data are not initially distributed, and the computation is a classical parallel one. However, the key concept of grids, service, virtual organizations and cross-domain access are present: a volume reconstruction algorithm is a service, that should be managed and provided transparently, reliably accessed from everywhere from authorized persons. A very concrete example is the current situation of our work with PTM3D: we have developed and integrated a completely functional parallel version, but it cannot be used routinely (which is the goal) from a hospital or radiology center if this institution does not include a cluster in its internal network (which is not at all the standard case), because of access rights, firewalls and other mundane but inescapable impediments.

The requirements, however, are quite new with respect to the grid most frequent use cases. The challenge is to smoothly integrate local and remote computing. The radiologist wants a unique tool, to browse the images, to measure a tumor, to reconstruct a complex organ or to create a virtual reality situation. Technically, this implies to move grid systems from batch/throughput-oriented to interactive/response time-oriented processing. In a broader view, the need is to switch from the exploitation model of large, but isolated, virtual high performance computing centers, to a commodity resource available on-demand.

3. The PTM3D software

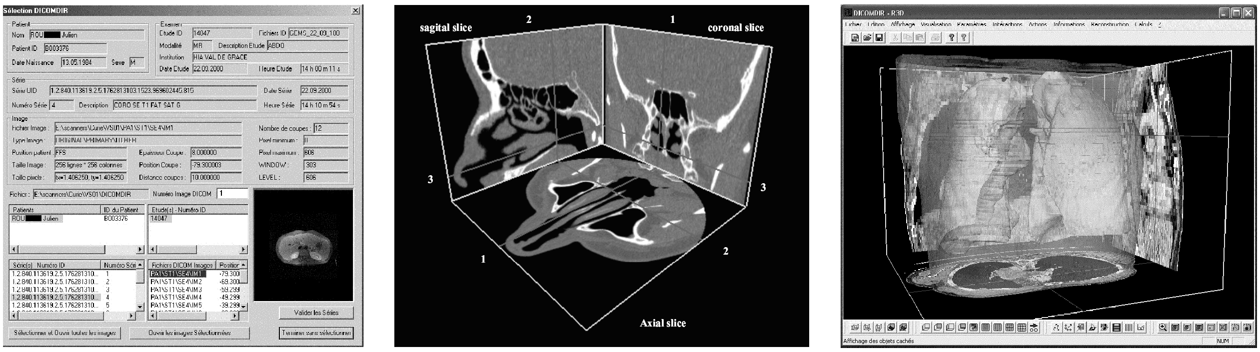

Sélection DICOMDIR

Patient
Nom: ROU[redacted] Julien
Patient ID: B003376
Date Naissance: 13.05.1984
Sexe: M

Examen
Etude ID: 14047
Fichiers ID: GEMS_22_09_100
Modalité: MR
Description Etude: ABDO
Institution: HIA VAL DE GRACE
Date Etude: 22.09.2000
Heure Etude: 14 h 00 m 11 s

Série
Série UID: 1.2.840.113619.2.5.1762813103.1523.969602445.815
Date Série: 22.09.2000
Numéro Série: 4
Description: CORO SE T1 FAT SAT G
Heure Série: 14 h 10 m 54 s

Image
Fichier Image: E:\scanners\Curie\VS01\PA1\ST1\SE4\IM1
Nombre de coupes: 12
Type Image: ORIGINAL\PRIMARY\OTHER
Pixel minimum: 0
Position patient: FFS
Epaisseur Coupe: 8.000000
Pixel maximum: 606
Taille Image: 256 lignes * 256 colonnes
Position Coupe: -79.300003
WINDOW: 303
Taille pixels: tx=1.406250, ty=1.406250
Distance coupes: 10.000000
LEVEL: 606

Fichier: E:\scanners\Curie\VS01\DICOMDIR
Numéro Image DICOM: 1

Patients	ID du Patient
ROU[redacted] Julien	B003376

Etude(s) - Numéro ID
14047

Série(s) Numéro ID	Numéro Séri
1.2.840.113619.2.5.176281310...	1
1.2.840.113619.2.5.176281310...	2
1.2.840.113619.2.5.176281310...	3
1.2.840.113619.2.5.176281310...	4
1.2.840.113619.2.5.176281310...	5

Fichiers DICOM Images	Position
PA1\ST1\SE4\IM1	-79.300
PA1\ST1\SE4\IM2	-69.300
PA1\ST1\SE4\IM3	-59.299
PA1\ST1\SE4\IM4	-49.299
PA1\ST1\SE4\IM5	-39.299

(a) DICOM interface (b) Navigation (c) 3D reconstruction

Figure1: PTM3D snapshots

3.1. *Overview*

PTM3D is a fully featured DICOM images analyzer developed at LIMSI. PTM3D transfers, archives and visualizes DICOM encoded data (fig. 1.a); besides moving independently along the usual three axes(fig1.b), the user is able to view the cross-section of the DICOM image along an arbitrary plane and to move it. PTM3D provides computer-aided generation of three-dimensional representations from TDM, MRI, PET-scan, or echography 3D data (fig. 1.c). A reconstructed volume (organ, tumor) is displayed inside the 3D view. The reconstruction also provides the volume measurement, required for therapeutic decisions. The system currently runs on standard PC computers and it is used on line in radiology centers (1).

3.2. *3D segmentation and volume reconstruction*

The reconstruction involves three main steps: segmentation, triangle mesh generation, and rendering. The segmentation process is semi-automatic, based on a few graphical hints from the user, and an active contour method. From the hint, the contour is automatically adjusted to the region of interest, and vertically propagated and deformed upward and downward the slices. The rationale of a user-guided, thus interactive, segmentation is the following. Because of the complexity of the human body, and the high anisotropy of medical images, atlases are not currently available; besides, specific models of known lesions cannot be assumed, as each day carries its share of new aspects of new malformations. The result is the need of on-demand reconstruction of any zone of interest in a very noisy context: from the computing point of view, medical images appear as a conglomerate of very noisy textures. Thus, an initial hint must be given by the radiologist, and he/she must be able to continually supervise the contour deformation, e.g. to require more iterations of the deformation algorithm at points of steep change in the geometry of the reconstructed region.

In the mesh generation step, all surfaces joining two adjacent contours are generated as a set of triangles, which will be used for visualization and measurements. The resulting data are stored as OpenGL data structures. This, and the overall development under C++, ensures portability across various graphical and windowing user interfaces: the system is currently implemented on MS-Windows, X11R6 and Motif 1.2.

3.3. Use cases related to percutaneous nephrolithotomy

a. 3D reconstruction and instant volume measurement of complex and staghorn renal stones (2)

As said before, this example does not require the grid computing power, but is useful to give evidence of the validity of the volume reconstruction and measurement tool. Validation compares the stones volumes after PCNL with those obtained by physical measurements (laser and immersion). CT scans were performed for 43 patients (5 mm collimation every 3 mm). All stones were also reconstructed and their volumes measured after PCNL according to the same CT protocol. For patients with residual fragments, a new CT was done, and the residual fragments were calculated. Comparison between in vivo and ex vivo volume measurements showed a 5 % discrepancy. In the ex vivo case, the discrepancy with physical measurements was about 2%. This difference is probably explained by the loss of stone powder during the extraction.

b. An augmented reality system for percutaneous nephrolithotomy (3).

Surgical planning for this intervention is a complex process. The volume reconstruction part is thus quite modest, but in the order of half an hour for sequential (PC-based) execution. Interactive real-time response time for will allow for more accurate measurements, and has thus the potential to improve health care.

The surgery guiding strategy is (fig. 2):: to reconstruct the 3D shape of renal stones, kidneys and the external skin of the body; to project the reached forms on the patient in the surgery position; to align interactively 3D forms with the patient body using control points until the forms match precisely their exact positions in the body. The device was validated on phantoms and clinical evaluations are in progress.

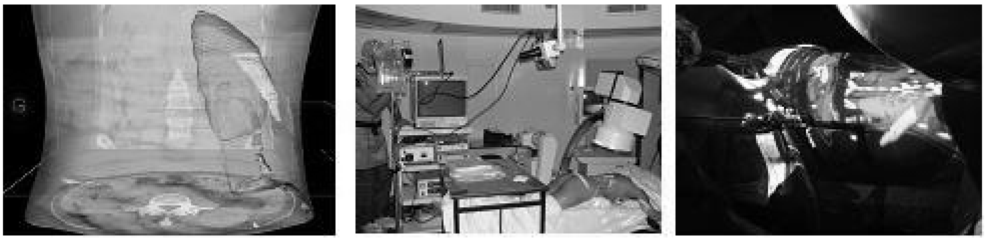

Figure 2: Augmented reality in surgical planning for percutaneous nephrolitotomy

4. A Cluster-enabled PTM3D

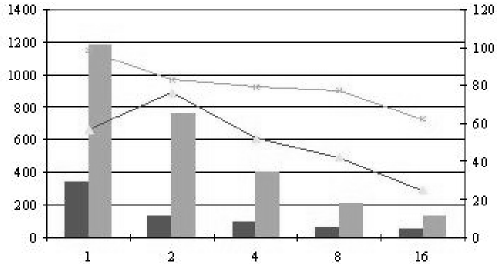

Figure3. Parallel scheduling

Clusters have become the standard parallel commodity; they are much more widely available than high-end parallel machines, and are the elementary node of deployed grid infrastructures. We consider first the easiest (but unrealistic) case, where the cluster is fully available for exclusive use. The main purpose of this section is to show that volume reconstruction requires only medium to low-end configurations.

From the parallelization point of view, the mesh generation component of PTM3D is master-slave: it can be divided in a fairly large number of independent tasks, with a small input and a large output. Each task is the tessellation of one pair of adjacent slices. While static (pre-planned) parallel scheduling has been extensively studied (for a review of recent advances, see (4)), the PTM3D mesh generation can take little profit of this research, because nothing can be known in advance about the number or duration of tasks.

To accommodate such an extremely dynamic profile, we use centralized self-scheduling (for a comparison with other scheduling schemes, see (5)). Self-scheduling requires two types of agents: one common scheduling agent, and a worker agent for each processor. The scheduling agent dispatches work to the workers following an on-demand scheme: when a worker finishes its task, it asks the scheduler for a new one. Besides performance, the scheduler/worker scheme naturally fits with interactivity. The flow of input data may be altered on the fly by some user interaction on the front-end, typically when stopping or reshaping the segmentation.

Fig. 3 shows the performance of this scheme. The experimental setup is the LRI cluster (AMD 1800+, Linux 2.4.20, Ethernet 100 switch); the PTM3D front end runs on a desktop at LRI. The histogram records the execution time in seconds (left axis), and the curves the percentage of processor occupation averaged over all processors (right axis). The two examples are: *e1* - two lungs (225 slices) and *e2* - a body as in fig. 2, left size (156 slices). The body volume requires much more triangles than the lungs. The main result is that, for both cases, interactive time can be reached with moderate parallelism. As shown by average occupation, the scheduler is not able to keep all processors busy all the time. A detailed analysis of the execution steps shows that the contribution of network contention is minor, the major reason being temporary slackness between the PTM3D front-end and the cluster when the segmentation does not generate slices fast enough.

5. Towards a grid-enabled PTM3D

5.1. Grid protocols penalty

The first experimented grid middleware was Globus 2. The performance criterion is the completion time, with exclusive access. The difference with the previous scenario is: first, the latency of WAN access; second the penalty of grid protocols

The experiments were performed on the TAG grid at ICPS, running Globus 2.4. The processors are AMD 2000+ and Pentium III 800MHz. The *regular dynamic* scheme, where the scheduler agent is kept local and launches one Globus job per slice, or group of slices, is attractive, because it allows for adaptation to a dynamic behavior of the grid resources. However, the overhead of launching and terminating a job in the Globus framework is not negligible, in the order of 40s on the TAG local grid, and much more in a real (nation-wide) grid. Thus, we switched to a *regular static* scheme: on activation of the PTM3D reconstruction tool by the user (clicking a button), a Globus job, including a scheduler and some workers, is submitted through a *globus_job_run* command; static refers to the fact that the number of workers is part of the job specification and cannot be changed.

Experiments for the Regular Static scheme gave results very similar to the parallel scheduling scenario, which was expected: the Globus startup cost is paid only once, and the latency of the WAN access is hidden by the interleaving of communications and computations.

The real difficult problem is elsewhere, at the level of the communication between the Globus resources and the outside world. The basic interface is through files, when we need streams. While simulating the latter with the former is obviously feasible, this option is unnatural and contrived. An adequate solution has been proposed in the EU CROSSGRID project (related to the large grid projects Datagrid-EDG (6) and its successor EGEE), through the interactive option of the *edg-job-submit* facility, which redirects the standard streams of the launched grid processes to named pipes. We are currently migrating to the EDGE/GEE platform.

5.2. Resource sharing

The scenarios considered so far assume the unrealistic hypothesis of exclusive access. In practice, an interactive job will have to share the grid resources with other processes (7). Advance reservation is the basic option for true real-time situations (e.g. intra-operative imaging). For less critical situations (e.g. surgical planning), the appropriate goal is soft real-time schedulers, defined by a high probability of QoS with the possibility of rejecting a request when the resources cannot provide the required performance. Interactive jobs are thus subject to regulation through the grid queuing policy like batch ones. The overall management of queuing should ensure the following contract.

- The real-time requirements of the interactive jobs are met.
- The delays incurred by the batch jobs because of the interactive ones remain bounded
- Interactive jobs do not degrade resource utilization.
- Interactive jobs can be smoothly integrated with queuing policies governing the batch jobs.

The first requirement targets real-time interactivity. The other ones mainly target sociologic acceptance. Grid resources (hardware, software, manpower) are not yet commodity resources as desktop PCs are, meaning that the funding organizations are not willing to be exceedingly disturbed by small jobs, and ask for some way of measuring the efficiency of their investment. The simplest and most frequent one (if not always meaningful) is resource utilization. Site managers design complex queuing policies in order to balance the requirements of different classes of customers amongst themselves and versus resource utilization. Because of the current structure of grid funding, such policies are often designed with long batch jobs in mind. Bootstrapping the interactive use of grids will not be possible if it requires a complete redesign of the sites policies, thus the fourth requirement.

The simplest solution is to allow the interactive jobs to run concurrently with the batch ones. To ensure reproducibility, we have tested this situation on the LRI cluster and the TAG one, running only artificial background jobs (namely 1, 2 and 4 jobs), and our application, on the above examples. When all processes have equal processor priority, and on a 16 processors configuration, it appears that one long job has limited impact; four long jobs are hardly acceptable (within the 2 minutes limit). If there are not enough resources for the interactive job to meet its deadline with no special processor priority, the running long job should get less processor priority (be niced). The reason lies in the fact that the volume reconstruction underutilizes the processor resource; one background job (or many niced) exerts an appropriate pressure towards more efficient resource utilization, dynamically achieving task molding (4).

What remains to be done is to integrate this scheme within a real grid scheduler. We have decided to target the EGEE middleware, as explained in 5.1. Raising requirements for a new and prospective application in real-world grids requires an incremental approach. The first step is to define and implement access control, that is the definition of a class of interactive job at the grid scheduler level with guaranteed QoS: e.g. if *N* clusters are required to have always at least 16 processors running at most one batch jobs (or alternatively be ready to nice their batch jobs), *N* simultaneous accesses can be granted for our application. This is obviously only an initial step: more complex and adaptive scheduling algorithms should be studied, for instance increasing the resources allocated as the deadline approaches.

6. Conclusion

We have exemplified to simultaneous need of grid resources and interactivity on the volume reconstruction case. A large number of medical images analysis applications share the same requirements, as decribed in more details in (9): in high-degree-of-freedom optimization problems such as non-linear image registration (8) or radiation treatment planning, computational steering is useful to help algorithms avoid local minima. Another very important area is analysis of massive distributed datasets. The example of the EU Mammogrid project, and its relation with the High Energy Physics (HEP) experiment ALICE (10) is significant. In both cases, remote operations are launched over the distributed datasets (respectively mammograms and LHC data) to *extract* the relevant information from huge datasets; however, *analysis* of this information remains, at least partially, at the local level: many users (respectively clinical researchers and physicists) design tools that suit their expertise, hypothesis or goals. This scheme might be even more relevant when drug design is concerned, the analysis tools and working hypothesis being highly sensitive information.

What comes out of this example is the striking convergence of two broadly different fields: the computing center model is HEP tradition, while radiologists, and often medical epidemiologists, have no interest in it, and consider only local computing resources. The major European efforts in the grid area (LCG and EGEE) are strongly influenced by the HEP community, but, at least for EGEE, target other end-users. From this convergence, it can be expected that the software will mature towards truly interactive grids, able to serve the needs of the medical community.

Acknowledgments

We thank Stéphane Génaud and the TAG team, who have given us access to the Globus infrastructure at ICPS, and greatly helped us in setting up the experiments; we also thank the referees for their helpful comments.

References

1. Osorio A., Valette P-J, Mihalcea A., Atif J., Ripoche X. A new PC based software to perform semi-automatic hepatic segmentation using CT and MR images. InfoRAD-RSNA 2002
2. Traxer O., Merran S., Osorio A., Atif J., Ripoche X., Thibault P. 3D reconstruction and instant volume measurement of complex renal calculi: a treatment preparation tool for percutaneous nephrolithotomy. InfoRAD-RSNA 2003.
3. Osorio A., Traxer O., Merran S., Atif J., Ripoche X., Tligui M. An augmented reality system for percutaneous nephrolithotomy. InfoRAD-RSNA 2003.
4. Beaumont O., Legrand A., Robert Y. Optimal algorithms for scheduling divisible workloads on heterogeneous systems. Procs 12th IEEE Heterogeneous Computing Workshop. 2003
5. Germain C., Osorio A., Texier R. A case study in medical imaging and the grid. 1st Health grid Conference. pp 110-118. 2003. EC-IST.

6. Breton V., Medina R., Montagnat J. Datagrid, Prototype of a BioMedical grid. MIM 42(2), pp 143-148, 2003

7. R. Raman, M. Livny, and M. Solomon. Resource Management through Multilateral Matchmaking. *Procs 9th IEEE Symp. on High Performance Distributed Computing (HPDC9)*, pp 290-291. 2000

8 R. Stefanescu, X. Pennec and N. Ayache. Grid Enabled Non-Rigid Registration with a Dense Transformation and A Priori Information. In *Proc. of MICCAI'03*, LNCS, 2003. Springer Verlag.

9 Berry D., Hill D., Knopp D., Pieper S., Saltz J., Germain C. Report of the IMAGE03: Images, Medical Analysis and Grid Environments workshop. To appear in UK e-Science Technical Report Series. http://www.nesc.ac.uk/technical_papers/

10 Estrella F., McClatchey R., Rogulina D., Amendolia R., Solomonides T.. A Service-Based Approach for Managing Mammography Data. 11th World Congress on Medical Informatics (MedInfo'04). San Francisco, USA. 2004