

The Design of Dynamic Exploration Environments for Computational Steering Simulations

R.G. Belleman and P.M.A. Sloot

Section Computational Science, Universiteit van Amsterdam,
Kruislaan 403, 1098 SJ Amsterdam, the Netherlands
email: [robbel,sloot][@science.uva.nl](mailto:[robbel,sloot]@science.uva.nl)
phone: (+31) 20 525 7462, *fax:* (+31) 20 525 7490

Abstract

Interactive simulation environments are systems that combine simulation, data presentation and interaction capabilities that together allow users to explore the results of computer simulation processes and influence the course of these simulations at run-time. The goal of these interactive environments is to shorten experimental cycles, decrease the cost of system resources and enhance the researcher's abilities for the exploration of data sets or problem spaces.

The conceptual idea of allowing users to interactively steer a computation while it is running is not new and examples of applications that benefit from it are abundant. However, very often these applications provide *ad hoc* solutions that are very specific to the problem at hand. This paper investigates the issues that are involved with building an interactive computational steering environment in an attempt to improve the synergy between computational simulation and interactive exploration through a generalised architecture. The architecture is validated by analysis of a prototypical case study of *simulated abdominal vascular reconstruction*.

1 Introduction: exploration environments

In many scientific computing problems, the complexity of both the simulation and the generated data is too vast to analyse analytically or numerically. For these situations, exploration environments provide essential methods to present and explore the data in a way that allows a researcher to comprehend the information it contains. Exploration environments combine presentation and interaction functions into one system to allow exploration of large data spaces. These data spaces may originate from data acquisition devices or represent results from computer simulations. In our research we discriminate between static and dynamic exploration environments.

In static exploration environments (SEE), the data presented to the user is time invariant; once the data is loaded into the environment, the user is presented with a visual representation of this data. Interaction methods are provided to

change the visualization parameters interactively in order to get the best view to gain understanding. The data itself, however, does not change (see Fig. 1).




Fig. 1: Schematic representation of a static exploration environment (SEE).

An important step towards a successful exploration environment is to involve the researcher into the presentation as much as possible, thereby increasing the researcher’s level of awareness [7]. To achieve this, an exploration system needs the following, often conflicting capabilities:

- High quality presentation – The most common method to provide insight in large multi-dimensional data sets is to represent data as visual constructs that present quantitative and relational aspects to the observer in an apprehensive manner. Many scientific visualization environments are now available that provide means of efficiently achieving this [17, 27, 33, 37].
- High frame rate – While the capabilities of modern graphical hardware allow increasingly complex images to be rendered with relative ease, the level of detail in the presentation should be minimized to avoid information clutter and achieve high frame rates (a compromise is often necessary). For an interactive exploration environment the visual frame rate should be at least 10 frames per second [8].
- Intuitive interaction – A prerequisite of a successful SEE is that a sufficiently rich set of interaction methods is provided that allows a user to extract both qualitative and quantitative knowledge from the data sets. An unfortunate side effect of increasingly richer sets of interactive methods is that user-friendliness is compromised, so careful consideration is required during user-interface design.
- Real-time response – Some delay will always occur between the moment a user interacts with a presentation and the moment that the results are visible. This is caused by low tracking rates of input devices, (re-)computations, communication delays or temporary reduced availability of computational or network resources. To attain accurate control over the environment and to avoid confusion with the user, the amount of lag in an exploration system should be minimized [36].

Provided these capabilities are carefully considered, such environments are well suited for the exploration of *static* multi-dimensional data sets [5].

Static exploration environments can be customized to observe iteratively updated data sets produced by “living” simulations. When interaction with the simulation is also allowed, however, we speak of *dynamic* exploration environ-

ments and radically different considerations come into play, as we describe in the next section.

2 Design considerations for the construction of dynamic exploration environments

Dynamic exploration environments (DEE) extend the previously described static model in that the information provided to the user is regenerated periodically by an external process, in our case a computer simulation. Here, the environment is expected to provide (1) a reliable and consistent representation of the results of the simulation *at that moment* and (2) mechanisms enabling the user to change parameters of the external process (i.e. simulation) (see Fig. 2).




Fig. 2: Schematic representation of a dynamic exploration environment (DEE).

Dynamic environments have additional requirements over static environments. For example, in static interactive systems, the interaction functions can be implemented inside the visualization environment, since the only interaction that takes place is with the visualization. In dynamic environments, interaction influences both the visualization and the simulation environment. Changing a static environment into a dynamic environment therefore requires that at least one module performs additional processing to service the interaction. Such a change makes these environment less suitable for use in other applications without significant modifications. In the sequel we address some of the functionalities required to develop generic dynamic exploration environments. We start with a brief description of the specific time management aspects in DEE and present a top-down description of the associated additional requirements in such systems.

2.1 Time management

An important issue in a DEE is time management. Time management deals with the exchange of time stamped information between components. For a DEE, the four most time demanding components are; the simulation environment, the visualization modules, the rendering layer and the explorer (i.e. the user).

Figures 3 and 4 show time frame diagrams illustrating the advancement of time, under two different time management strategies; *lock-step* and *asynchronous*. Time frames are illustrated by rounded boxes. The gaps between time frames on a same level represent the *idle* time of the component on that level. The gaps between time frames on neighbouring levels represent the *delays* that occur between the time one component is done with a time frame and the next component starts working on it. These delays are delineated at the bottom of the figure.

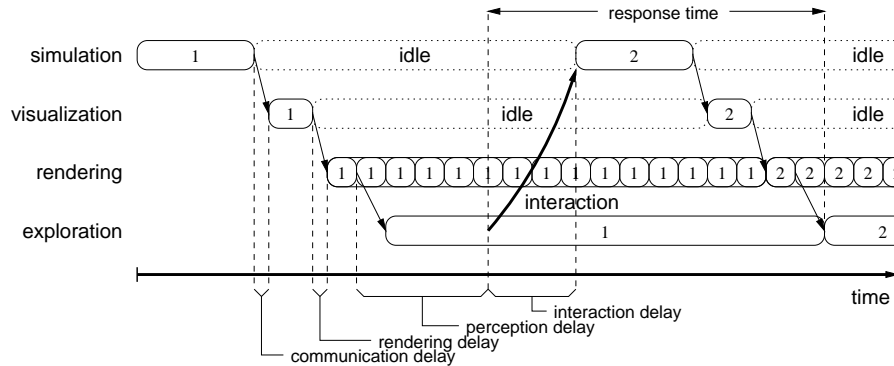


Fig. 3: Time frames and delays in a lock-step interactive dynamic exploration environment.

Fig. 3 shows a time frame diagram in a lock-step interactive dynamic exploration environment. In this strategy, the simulation is allowed to advance only if the user explicitly tells the environment it is alright to do so. While the user is exploring the results rendered by the graphics system, the simulation and visualization modules sit idle. In situations where a single simulation, visualization and rendering time frame takes a negligible amount of time, this strategy may be perfectly adequate since the user will see the result of his interaction in short notice. However, if these time frames are long, it may take a long time before the result of an interaction is shown. This often leads to an unusable environment and confusion with the user.

The time required before simulation updates are presented to the user (i.e. the length of a time frame on the exploration level) can be shortened by allowing the simulation to run asynchronously from the rest of the environment. Fig. 4 shows a time frame diagram in an asynchronous environment. In an asynchronous system, different components are allowed to advance when they have finished processing and communicating the current time frame. Different components may therefore execute at different time scales. In addition, these time scales are mostly nondeterministic because of hardware, software and human imposed delays. As a consequence, time delays occur as the output generated by one component can not be accepted immediately by another component for

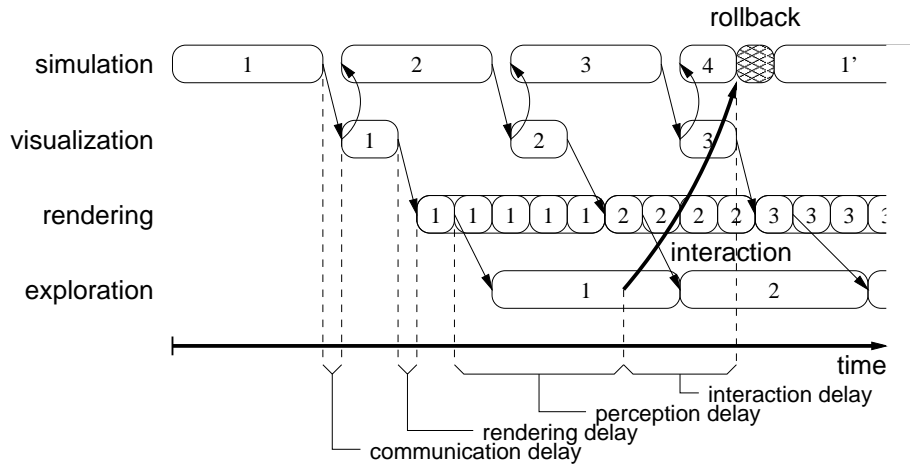


Fig. 4: Time frames and delays in an asynchronous interactive dynamic exploration environment.

processing. When components depend on the output of an increasing number of other components, the time frame that is processed by “later” components fall further apart. This has a causality consequence for the user who explores the final component in the environment and therefore interacts with components at a much “earlier” time compared to what is being processed by a simulation at the same wall-clock time. Time management is responsible for detecting and resolving this causality violation. Methods for resolving time causality problems have been investigated within our group [28].

2.2 Interaction

A DEE provides the opportunity to interact with a living simulation. This interaction can take any form; from typed input for simple types of interaction via graphical user interfaces to fully immersive virtual environments. The main feature of immersive environments over other graphical user interfaces is that user-centered stereoscopic images are presented to a user rather than visualization centered three-dimensional (3D) projections. User-centered stereoscopic images differ from projections on a flat screen in that slightly different picture are generated for the left and right eye, dependent on the position of the viewer. This makes images “pop out of the screen” and react to the user’s movements¹, an important depth-cue to gain understanding in complex multi-dimensional structures.

A minimal requirement for interaction in an immersive virtual environment (VE, see also [19]) is the availability of input devices that can be used by a user to convey intention to the environment. The most common are sensor devices

¹This is commonly referred to as “motion parallax”.

that measure the 6 degrees of freedom (DOF) one has to move around in a 3D space (position and orientation). These sensors can be used to detect the proximity of a physical object (such as the user's hand) to virtual objects, so that the user can interact with them. Interaction with a virtual environment is a key issue, especially in an interactive simulation environment. The following types of interaction are deemed mandatory.

- Object interaction – An “object” is defined here as a visual entity that is in the center of interest to an end-user. These objects are representations of data sets or simulation results but can also be “widgets” (menus, buttons, sliders, etc.). An object has attributes associated with it such as position, scale, level, state, etc. Object interaction is concerned with changing these attributes.
- Navigation and wayfinding – Navigation provides users with methods to move beyond the confinements of the VE's physical dimensions. Objects beyond the VE come into reach by moving the user towards them. Note how this concept places the user of the VE in the center of this type of interaction; the user is transported from one place to another while the objects remain where they are. Wayfinding is a relatively new concept in VE applications and provides the user with a reference on where he is in a virtual environment [14].
- Probing – Although visual presentations allow researchers to qualitatively analyse their data, an instrument for obtaining quantitative information from the visualization is a valuable asset. Previously, we have developed an architecture that allows researchers to probe visual presentations in order to obtain quantitative information [4].

2.3 Coordination: Intelligent Agents

Since the various components in a DEE are independent processes, some means of coordination between them is required to allow a complex environment of this kind to be used. Especially in the case of interactive simulations, interaction involves not only the visualization part, but also the simulation part. For software engineering and efficiency reasons, it is reasonable to move the processing for those general interactions into independent modules, and let them be reusable to all components that need these interactions. Our approach to this is through Intelligent Agents (IAs, see also Fig. 5).

IAs are software modules with the capability of performing three functions: (1) perceiving state changes in the environment through the use of monitors, (2) taking actions that affect conditions in the environment and (3) reasoning to interpret perceptions, solve problems, draw inference, and determine actions [16]. Agents execute autonomously, interfering minimally with the rest of the environment, apart from communicating with other agents or the user.

Feedback is generated when the agent has solved a problem, or has prepared suggestions based on the current running context. This feedback could be information to the user, or the information is sent to environment components. Depending on the permission settled beforehand, an agent could just provide




Fig. 5: A DEE instrumented with monitors and an Intelligent Agent.

feedback without affecting the working status of the environment, or make some changes in the environment based on its reasoning.

We currently use agents to provide feedback to the user concerning the state of a simulation (e.g. accuracy of the simulation, time to completion, convergence rate), user interaction (including speech recognition and synthesis). In the near future we will develop agents for feature extraction (e.g. determining the geometric skeleton of objects, detecting eddies in flow domains) and providing assistance (context sensitive help).

2.4 Distributed execution environment

Although the capabilities of modern computer systems are nearing the requirements for performing both simulation and visualization tasks on the same machine, some performance increase may be attained by running these tasks on dedicated computing platforms. For example, many simulation applications perform better on dedicated hardware such as vector processors, massively parallel platforms or other high performance computing (HPC) machinery, while state-of-the-art graphical systems are now available that are well suited for the visualization tasks. Moreover, a decomposition of the environment into separate communicating tasks facilitates implementation and allows more control over the performance of the system as a whole (in Fig. 5, each block can be considered a separate process or a combination of processes, possibly running on different systems).

Especially in the case of distributed environments, some means of job control is required that starts/stops the execution of the different components of the environment on the various computing platforms. In many organizations this system also needs to allocate the required resources prior to execution (for example in the case of batch execution systems). GLOBUS is one such software infrastructure for computations that integrates geographically distributed computational and information resources [15].

In distributed systems, components execute on different, possibly heterogeneous computing platforms. To be able to communicate data with each other,

components provide access to their attributes which can then be made available to other components. In heterogeneous computing environments, the attributes often have to be converted into different representation format. Furthermore, in many circumstances not all components in an environment will participate in a communication. For these situations a publication and subscription mechanism needs to be provided that limits communication to members of a restricted group. Message passing systems such as PVM and MPI support this data distribution facility for the most part. In general, however, these systems do not support the construction of asynchronous systems as they restrict the application programmer to the Single/Same Program Multiple Data (SPMD) paradigm, which in DEEs is troublesome.

2.5 Attribute ownership

The behaviour of individual components in the environment is defined by one or more attributes (or parameters) which together define the state of that component. In a distributed system, attribute changes (which can be considered events, for example as a result of user interaction) should only be performed by a component that *owns* the attribute to avoid race conditions. In some cases it may be necessary to transfer ownership so that attributes can be changed by other components (for example in a collaborative environment where multiple users manipulate the same components).

2.6 Runtime support system

From the considerations described in the previous sections, it becomes clear that a generic framework to support the different modalities is required. In our research we have chosen for the “High Level Architecture” as a suitable architecture for constructing a DEE. The High Level Architecture (HLA) aims to establish a common architecture for simulation to facilitate interoperability among simulations and promote the reuse of simulations and their components [12, 25]. As a successor to the DIS (Distributed Interactive Simulation) protocol, HLA provides a robust architecture with which distributed discrete event and other types of simulations can be designed.

HLA simulations are composed of “federates”, these are gathered to form “federations”, interaction between federations is controlled by a “Runtime Infrastructure” (RTI). HLA systems comprise *rules* which govern how federates and federations are constructed; an *interface specification* which governs how federates interact with the RTI; and an Object Model Template which provides a method for documenting key information about simulations and federations.

HLA provides solutions to many of the problems and issues described in the previous sections. Specifically, HLA allows data distribution across heterogeneous computing platforms (including message groups), supports a flexible attribute publish/subscribe and ownership mechanism and offers several methods to do time management.

3 Vascular reconstruction: a full-blown case study

The design considerations described in the previous section cover the issues that are involved with building a DEE. The architecture is validated by analysis of a prototypical case study of *simulated abdominal vascular reconstruction*.

The application we have chosen as a test case combines visualization, simulation, interaction and real-time constraints in an exemplary fashion. By a detailed analysis of the spatial and temporal characteristics of the test case we attempt to recognize generic elements for the design of a computational steering architecture. We begin with a description of the test case.

3.1 Simulated abdominal vascular reconstruction

Vascular disorders in general fall into two categories: *stenosis*, a constriction or narrowing of the artery by the buildup over time of fat, cholesterol and other substances in the vascular wall, and *aneurysms*, a ballooning-out of the wall of an artery, vein or the heart due to weakening of the wall. Aneurysms are often caused or aggravated by high blood pressure. They are not always life-threatening, but serious consequences can result if one bursts.

A vascular disorder can be detected by several imaging techniques such as X-ray angiography, MRI (magnetic resonance imaging) or computed tomography (CT). Magnetic resonance angiography (MRA) has excited the interest of many physicians working in cardiovascular disease because of its ability to non-invasively visualize vascular disease. Its potential to replace conventional X-ray angiography that uses iodinated contrast has been recognized for many years, and this interest has been stimulated by the current emphasis on cost containment, outpatient evaluation, and minimally invasive diagnosis and therapy [40].

A surgeon may decide on different treatments in different circumstances and on different occasions but all these treatments aim to improve the blood flow of the affected area. Common options include *thrombolysis* where a blood clot dissolving drug is injected into, or adjacent to, the affected area using a catheter; *balloon angioplasty* and *stent placement* which is used to widen a narrowed vessel by means of an inflatable balloon or supporting framework; or *vascular surgery*. A surgeon resorts to vascular surgery when less invasive treatments are unavailable. In *endarterectomy* the surgeon opens the artery to remove plaque buildup in the affected areas. In vascular bypass operations, the diseased artery is shunted using a graft or a healthy vein harvested from the arm or leg.

The purpose of vascular reconstruction is to redirect and augment blood flow or perhaps repair a weakened or aneurysmal vessel through a surgical procedure. The optimal procedure is often obvious but this is not always the case, for example, in a patient with complicated or multi-level disease. Pre-operative surgical planning will allow evaluation of different procedures *a priori*, under various physiologic states such as rest and exercise, thereby increasing the chances of a positive outcome for the patient [35].

3.2 What is needed?

The test case described in section 3.1 contains all aspects of an interactive dynamic exploration environment that are of consequence in the construction of a generic dynamical computational steering architecture. Our aim is to provide a surgeon with an environment in which he/she can try out a number of different bypass operations and see the influence of these bypasses. The environment needs the following:

- An environment that shows the patient under investigation with his inflection. A patient's medical scan is 3D, so to obtain best understanding on the nature of the problem, the surgeon should be able to look at his specific patient data in 3D, using unambiguous visualization methods.
- An environment that allows the surgeon to *plan* a surgical procedure. Again, this environment should allow interaction in a 3D world, with 6 DOF. The CAVE environment allows us to interact with 3D computer generated images using 6 DOF interaction devices [1, 11]. Note that *visual* realism is not the primary goal here; what is more important here is *physical* realism, and then only of particular issues in fluid flow, as discussed later².
- An environment that shows the surgeon the effect of his planned surgical procedure. As the aim of the procedure is to improve the blood flow to the affected area, the surgeon must have some means to compare the flow of blood before and after the planned procedure. This requires the following:
 - a *simulation* environment that calculates *pressure*, *velocity* and *shear stress* of blood flowing through the artery,
 - a *visualization* environment that presents the results of the simulation in a clear unambiguous manner,
 - an *exploration* environment that allows the researcher to inspect and probe (qualitatively and quantitatively) the results of the simulation (e.g. means for performing measurements, annotate observations, releasing tracer particles in the blood stream, etc.).

All this should be *interactive*, or in other words; it should be fast enough such that a surgeon does not have to wait for the simulation results.

4 Implementation of a dynamic exploration environment

Parts of the components mentioned in the previous section have already been implemented in the course of previous projects. Others require minor adaptations to fit into our dynamic exploration architecture. In the following subsections we will briefly discuss the current status of the visualization and exploration environment, the interaction environment, the simulation environment and the middleware that combines these together.

²This in contrast to research projects towards virtual operating theatres that include the simulation of tissue deformation and realistic blood spills [3, 6].

4.1 VRE: immersive static exploration

We have previously built a SEE called the *Virtual Radiology Explorer* (VRE [13, 39]) which is capable of visualizing medical CT and/or MRI data in 3D (see also Fig. 6). 3D data sets acquired with computed tomography (CT) or magnetic resonance imaging (MRI) are often displayed and evaluated from various perspectives or at different levels, including sets of single slices, stack mode (cine loop) interactive representation of sets of slices, or multi-planar reformation (MPR) represented as single slices or interactive cine loops. Despite the increased possibilities of acquiring data, clinical use of 3D rendering has been hampered by insufficient computing capacity in the clinical environment.




Fig. 6: The *Virtual Radiology Explorer* (VRE) environment allows medical data from hospitals to be preprocessed on HPC systems for 3D visualization. High speed networking initiatives such as the *GigaPort* project [34] allow hospitals to make interactive use of HPC visualization techniques for patient diagnostics.

An example of the clinical use of 3D rendering is simulated endoscopy³. Simulated endoscopy has several advantages over mechanical endoscopy (shorter acquisition times, increased patient comfort, higher cost-effectiveness, no complications of endoluminal instrumentation, field-of-view extending beyond the surface). In addition, simulated endoscopy can be used in virtual spaces that cannot at all, or only after violation of normal anatomical structures, be reached by mechanical (endo)-scopy.

From a clinical perspective, there is a demand in community hospitals to make an environment available suitable for the interactive rendering and interactive matching of, and switching between the above described data sets with an emphasis on simulated endoscopy. The VRE environment provides various such methods for the visualization of medical scans, including volume rendering (using SGI's *Volumizer* [31]), surface rendering (using Vtk [33] and OpenGL [30]), interactive clipping and surface mapping techniques. Mechanisms have been added that allow the VRE environment to be run in a CAVE or on an ImmersaDesk. The ImmersaDesk allows the VRE environment to be used in the radiology department. Shown in figure 7 is an isosurface representation of

³*Endoscopy* is a diagnostic procedure where an instrument is used to visualize the interior of a hollow organ.

the abdominal aorta from an MRA scan. A geometric probing system (GEO-PROVE, see [4]) is used to perform measurements on this representation.

4.1.1 VRE+

VRE+ extends VRE with methods that allow dynamic exploration. Various methods are added to visualize the results of a simulation while it is running. An intelligent agent system is integrated that constantly monitors a user's actions and provides feedback to the user. Currently, we have implemented a speech recognition agent which enables users to control the environment using hands and voice simultaneously. A second agent monitors the position of the user when using the GEOPROVE probing system and provides feedback on the accuracy of the measurements. Another agent monitors the state of the simulation environment and provides feedback on the state of the simulation.

For the planning part, the VRE+ environment is extended with means to “draw” a surgical procedure using a “grid editor”, as described in section 4.3.

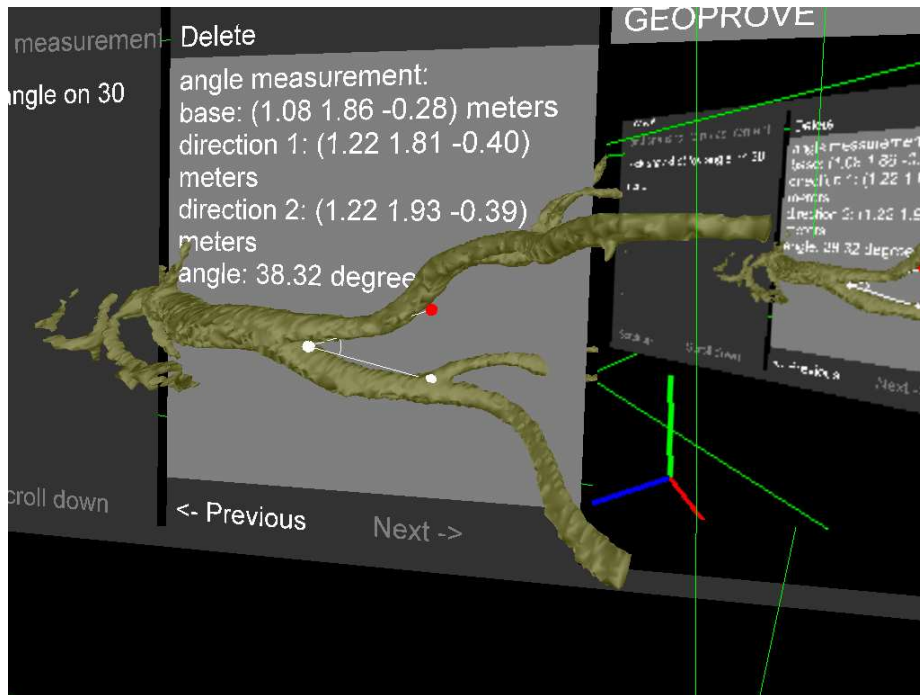


Fig. 7: A snapshot of the VRE environment running in a CAVE. An isosurface representation of an abdominal aorta is shown obtained from an MRA scan. The panel shows the GEOPROVE environment which allows measurements and annotations (such as the virtual “snapshot” on the right) to be made from within the environment.

4.2 Fluid flow simulation: the lattice-Boltzmann method

The lattice-Boltzmann method (LBM) is a mesoscopic approach for simulating fluid flow based on the kinetic Boltzmann equation [10]. In this method fluid is modeled by particles moving on a regular lattice. At each time step, particles propagate to neighboring lattice points and re-distribute their velocities in a local collision phase. This inherent spatial and temporal locality of the update rules makes this method ideal for parallel computing [21]. During recent years, LBM has been successfully used for simulating many complex fluid-dynamical problems, such as suspension flows, multi-phase flows, and fluid flow in porous media [23]. All these problems are quite difficult to simulate by conventional methods [20, 22].

The data structures required by LBM (cartesian grids) bare a great resemblance to the grids that come out of CT and MRI scanners. As a result, the amount of preprocessing can be kept to a minimum which reduces the risk of introducing errors due to data structure conversions. In addition, LBM has the benefit over other fluid flow simulation methods that flow around (or through) irregular geometries (like a vascular structure) can be simulated relatively easy. Yet another advantage of LBM is the possibility to calculate the shear stress on the arteries directly from the densities of the particle distributions [2]. This may be beneficial in cases where we want to interfere with the simulation while the velocity and the stress field are still developing, thus supporting fast data-updating given a proposed change in simulation parameters from the interaction modules.

4.3 Lattice-Boltzmann grid generation and editing

As mentioned earlier, the basic structure of the grids used in LBM bare great resemblance to the medical scans obtained from a patient. To convert the medical scans into LBM grids, the raw data from the medical scanner is first segmented so that only the arterial structures of interest remain in the data set (see also Fig. 8). The contrast fluid injected into the patient in a MRA scan provides sufficient contrast in the vascular structures to do this quite efficiently. The segmented data set is then converted into a grid that can be used in LBM; boundary nodes, inlet nodes and outlet nodes are added to the grid using a variety of image processing techniques.

A surgical procedure is simulated through the use of a 3D grid editor. This system allows a user to interactively add and/or delete areas in the LBM grid corresponding to the procedure that is simulated. Similar grid generation techniques as described above are used to ensure the grids comply to the demands imposed by LBM.

4.4 Middleware

The different components involved in our interactive simulation system are shown in Fig. 9. As can be seen from this figure, the visualization & exploration




Fig. 8: LBM grids are generated from raw medical scans through a combination of segmentation and image processing techniques.

system runs on a different system (a CAVE) than the simulation system (a massively parallel Origin 2000). HLA is used as a middleware layer to connect these components together. By using HLA, the different components can run asynchronously while spatial and temporal effects as described in section 2.1 can be controlled. In addition, HLA provides attribute ownership management and does efficient data distribution between heterogeneous (if needed) systems.

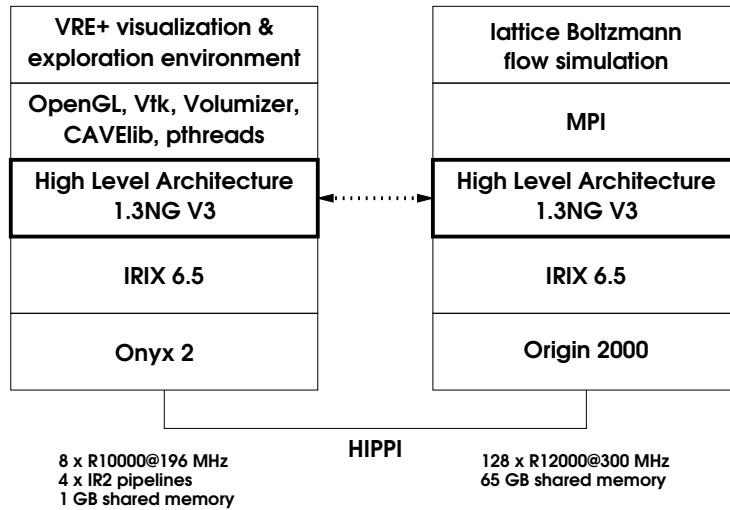


Fig. 9: The visualization & exploration environment (on the left, running in a CAVE) and the simulation system communicate via the High Level Architecture (HLA).

5 Discussion and future work

We have presented our views on dynamic exploration environments that support distributed interactive simulation. We have provided an overview on the requirements of such an environment and the issues involved in its construction. We have described how the High Level Architecture offers all requirements that are needed in its basic architecture. The case described in the final section is presented as a prototypical case study to validate these assumptions.

Preliminary measurements on the test case environment show that HLA is a suitable architecture to build a relatively efficient interactive and distributed simulation environment. Compared to raw network performance, communication overhead and delays imposed by HLA are acceptable. Implementing a HLA federation however, requires a substantial effort but is mostly due to the lack of proper software development tools.

The performance of the test case simulation environment will be validated through a comparison of fluid flow simulation results and the results of other simulation methods as well as *in vivo* measurements of blood flow through phantom structures and pre- and post-operative MRA scans.

6 Acknowledgments

This research is funded through grant 612-21-103 from the Netherlands Organization for Scientific Research (NWO). We are also greatly indebted to Charles A. Taylor (Dept. of Mechanical Engineering, Stanford University) for his insightful and inspiring discussions and for allowing us to use his data sets, Sean A. Spicer (Dept. of Mechanical Engineering, Stanford University) for providing his *Volumizer Convenience Classes* and enabling them for use in the CAVE, Silicon Graphics, Inc. for their patience in answering all our questions and fixing bugs in the course of this research, Drona Kandhai (Section Computational Science, Universiteit van Amsterdam) for his work on the lattice-Boltzmann fluid simulation environment and Zhiming Zhao (Section Computational Science, Universiteit van Amsterdam) for his work on HLA and intelligent agents. Finally, we would like to thank Eva Rombouts for her medical input and Alfons Hoekstra for his helpful remarks while reading this document.

References

1. Academic Computing Services Amsterdam (SARA), Amsterdam, the Netherlands. *SARA - CAVE Homepage*, 1998. <http://www.sara.nl/hec/vr/cave/>.
2. A.M. Artoli, D. Kandhai, A.G. Hoekstra, and P.M.A. Sloot. Accuracy of shear stress calculations in the lattice Boltzmann method. Accepted for the 9th International Conference on Discrete Simulation of Fluid Dynamics.
3. Cagatay Basdogan, Chih-Hao Ho, and Mandayam A. Srinivasan. Simulation of tissue cutting and bleeding for laparoscopic surgery using auxiliary surfaces. In James D. Westwood, Helene M. Hoffman, Richard A. Robb, and Don Stredney,

- editors, *Medicine Meets Virtual Reality*, pages 38–44, Amsterdam, 1999. IOS Press.
4. R.G. Belleman, J.A. Kaandorp, D. Dijkman, and P.M.A. Sloot. GEOPROVE: Geometric probes for virtual environments. In P.M.A. Sloot, M. Bubak, A. Hoekstra, and L.O. Hertzberger, editors, *High Performance Computing and Networking (HPCN'99)*, pages 817–827, Amsterdam, the Netherlands, April 1999. Springer-Verlag.
 5. R.G. Belleman, J.A. Kaandorp, and P.M.A. Sloot. A virtual environment for the exploration of diffusion and flow phenomena in complex geometries. *Future Generation Computer Systems*, 14(3-4):209–214, 1998.
 6. Uli Bockholt, Ulrich Ecke, Wolfgang Müller, and Gerrit Voss. Realtime simulation of tissue deformation for the nasal endoscopy simulator (NES). In James D. Westwood, Helene M. Hoffman, Richard A. Robb, and Don Stredney, editors, *Medicine Meets Virtual Reality*, pages 74–75, Amsterdam, 1999. IOS Press.
 7. Steve Bryson. Virtual reality in scientific visualization. *Communications of the ACM*, 39(5):62–71, 1996.
 8. Steve Bryson and Sandy Johan. Time management, simultaneity and time-critical computation in interactive unsteady visualization environments. In *Proceedings of Visualization '96*, page 255. IEEE Computer Science Press, Los Alamitos, CA, 1996.
 9. Jim X. Chen, David Rine, and Horst D. Simon. Advancing interactive visualization and computational steering. *IEEE Computational Science & Engineering*, pages 13–17, Winter 1996.
 10. S. Chen and G.D. Doolen. Lattice Boltzmann method for fluid flows. *Annu. Rev. Fluid Mech.*, 30:329, 1998.
 11. C. Cruz-Neira, D.J. Sandin, and T.A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *SIGGRAPH '93 Computer Graphics Conference*, pages 135–142. ACM SIGGRAPH, August 1993.
 12. Defense Modeling and Simulation Office (DMSO), Department of Defense, US. *High Level Architecture Run Time Infrastructure Programmer's Guide (1.3 version 7)*, 1999. <http://hla.dmsomil/>.
 13. Laura Durnford. Virtual reality: More than just a game. Radio Netherlands Wereldomroep, May 14th 1999. On the web: <http://www.rnw.nl/science/html/virtualreality990514.html>.
 14. T. Todd Elvins. Virtually lost in virtual worlds – wayfinding without a cognitive map. *Computer Graphics*, Spring 1997. See also <http://www.sdsc.edu/~todd/>.
 15. I. Foster and C. Kesselman. Globus: A metacomputing infrastructure toolkit. *Intl. J. Supercomputer Applications*, 11(2):115–128, 1997.
 16. B. Hayes-Roth. An architecture for adaptive intelligent systems. *Artificial Intelligence (special issue on agents and interactivity)*, 72:329–365, 1995.
 17. IBM Corporation, Armonk, NY. *Data Explorer Reference Manual*, 1991.
 18. Christopher R. Johnson and Steven G. Parker. Applications in computational medicine using SCIRun: A computational steering programming environment. In H.W. Meuer, editor, *Supercomputer '95*, pages 2–19, 1995.
 19. J.A. Kaandorp, editor. *Future Generation Computer Systems, special double issue on Virtual Reality in Industry and Research*, volume 14, numbers 3-4. Elsevier Science, August 1998.
 20. D. Kandhai. *Large Scale Lattice-Boltzmann Simulations (Computational Methods and Applications)*. PhD thesis, Universiteit van Amsterdam, Amsterdam, the Netherlands, 1999.

21. D. Kandhai, A. Koponen, A.G. Hoekstra, M. Kataja, J. Timonen, and P.M.A. Sloot. Lattice Boltzmann hydrodynamics on parallel systems. *Computer Physics Communications*, 1998.
22. D. Kandhai, D. Vidal, A. Hoekstra, H. Hoefsloot, P. Iedema, and P.M.A. Sloot. Lattice-Boltzmann and finite element simulations of fluid flow in a SMRX mixer. *Int. J. Numer. Meth. Fluids*, 31:1019–1033, 1999.
23. A. Koponen, D. Kandhai, E. Hellen, M. Alava, A. Hoekstra, M. Kataja, K. Niskanen, P. Sloot, and J. Timonen. Permeability of three-dimensional random fiber webs. *Physical Review Letters*, 80(4):716–719, January 26 1998.
24. David N. Ku. Blood flow in arteries. *Annu. Rev. Fluid Mech.*, 29:399–434, 1997.
25. Defense Modeling and Simulation Office (DMSO). High level architecture (HLA) homepage, 1999. On the web: <http://hla.dmsomil/>.
26. Jurriaan D. Mulder and Jarke J. van Wijk. 3D computational steering with parameterized geometric objects. In G.M. Nielson and D. Silver, editors, *IEEE Visualization '95*, pages 304–312. IEEE CS, November 1995.
27. The Numerical Algorithms Group Ltd., Oxford, UK. *Iris Explorer User's Guide*, 1998.
28. Benno J. Overeinder and Peter M. A. Sloot. Extensions to time warp parallel simulation for spatially decomposed applications. In David Al-Dabass and Russell Cheng, editors, *Proceedings of the Fourth United Kingdom Simulation Society Conference (UKSim 99)*, pages 67–73, Cambridge, UK, April 1999.
29. Steven G. Parker and Christopher R. Johnson. SCIRun: A scientific programming environment for computational steering. In *Supercomputing '95*, 1995.
30. Silicon Graphics Inc. Software Products. OpenGL homepage, 2000. On the web: <http://www.sgi.com/software/opengl/>.
31. Silicon Graphics Inc. Software Products. Volumizer homepage, 2000. On the web: <http://www.sgi.com/software/volumizer/>.
32. T.M. Roy, C. Cruz-Neira, T.A. DeFanti, and D.J. Sandin. Steering a high performance computing application from a virtual environment. *Presence: Teleoperators and Virtual Environments*, 4(2):121–129, Spring 1995.
33. Will Schroeder, Ken Martin, and Bill Lorensen. *The Visualization Toolkit, an object-oriented approach to 3D graphics (2nd edition)*. Prentice Hall, Upper Saddle River, NJ, 1997. ISBN 0-13-954694-4.
34. Surfnet. Gigaport homepage, 2000. On the web: http://www.gigaport.nl/en_index.html.
35. Charles A. Taylor, Thomas J.R. Hughes, and Christopher K. Zarins. Finite element modeling of three-dimensional pulsatile flow in the abdominal aorta: Relevance to atherosclerosis. *Annals of Biomedical Engineering*, 26:975–987, 1998.
36. Valerie E. Taylor, Jian Chen, Terrence L. Disz, Michael E. Papka, and Rick Stevens. Interactive virtual reality in simulations: Exploring lag time. *IEEE Computational Science and Engineering*, pages 46–54, Winter 1996.
37. C. Upson, T. Faulhaber Jr., and D. Kamins et al. The Application Visualization System: a computational environment for scientific visualization. *IEEE Computer Graphics and Applications*, 9(4):30–42, July 1989.
38. Robert van Liere, Jurriaan D. Mulder, and Jarke J. van Wijk. Computational steering. In H. Liddell, A. Colbrook, B. Hertzberger, and P. Sloot, editors, *High-Performance Computing and Networking*, pages 696–702. Springer-Verlag, April 1996.
39. Leslie Versweyveld. Exploring the medical applications of virtual reality techniques in the Dutch CAVE. Virtual Medical Worlds, March 2nd 1998. On the

web: <http://www.hoise.com/vmw/articles/LV-VM-04-98-13.html>.

40. E. Kent Yucel, Charles M. Anderson, Robert R. Edelman, Thomas M. Grist, Richard A. Baum, Warren J. Manning, Antonio Culebras, and William Pearce. Magnetic resonance angiography (update on applications for extracranial arteries). *Circulation*, 100:2284-2301, 1999.