PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A

Phil. Trans. R. Soc. A (2009) **367**, 2667–2677
doi:10.1098/rsta.2009.0034

Using high-resolution displays for high-resolution cardiac data

By Christopher Goodyer[1,*], John Hodrien[1], Jason Wood[1], Peter Kohl[2] and Ken Brodlie[1]

[1]*School of Computing, University of Leeds, Leeds LS2 9JT, UK*
[2]*Department of Physiology, Anatomy and Genetics, University of Oxford, Oxford OX1 3PT, UK*

The ability to perform fast, accurate, high-resolution visualization is fundamental to improving our understanding of anatomical data. As the volumes of data increase from improvements in scanning technology, the methods applied to visualization must evolve. In this paper, we address the interactive display of data from high-resolution magnetic resonance imaging scanning of a rabbit heart and subsequent histological imaging. We describe a visualization environment involving a tiled liquid crystal display panel display wall and associated software, which provides an interactive and intuitive user interface. The oView software is an OpenGL application that is written for the VR Juggler environment. This environment abstracts displays and devices away from the application itself, aiding portability between different systems, from desktop PCs to multi-tiled display walls. Portability between display walls has been demonstrated through its use on walls at the universities of both Leeds and Oxford. We discuss important factors to be considered for interactive two-dimensional display of large three-dimensional datasets, including the use of intuitive input devices and level of detail aspects.

Keywords: high-resolution displays; cardiac imaging; visualization

1. Introduction

Many of the problems investigated in the life sciences involve datasets with increasingly large volumes of data that span multiple spatio-temporal scales and modalities. This trend is driven by the increasing resolution of modern imaging technologies. Resulting two-dimensional slices and three-dimensional volumes are significantly larger than those that have typically been handled on desktop computers for viewing and analysis. This means that the methods of data extraction, rendering and display must adapt, ideally in a way that is scalable to address the certain future increases in resolution.

The long-term context is the drive towards personalized medicine, which builds on the appreciation that medical interventions need to be tailored for the patient, not 'just' the disease. Increasingly detailed medical imaging data require

* Author for correspondence (ceg@comp.leeds.ac.uk).

One contribution of 16 to a Theme Issue 'Crossing boundaries: computational science, e-Science and global e-Infrastructure II. Selected papers from the UK e-Science All Hands Meeting 2008'.

new tools for visualization, exploration, annotation, evaluation and training. These tools need to support medical decision-making in real time, or at least within a time frame that is comparable with modern laboratory parameter assessment (hours, not weeks).

In this paper, we illustrate our contribution to data visualization, exploration and annotation, focusing on two datasets from a project addressing the individual histo-anatomy of a rabbit heart. The heart was first scanned non-invasively, using high-resolution magnetic resonance imaging (MRI), providing a uniquely detailed three-dimensional dataset of cardiac anatomy. Subsequently, the whole organ was serially sectioned for histological staining and light microscopy, providing a stack of two-dimensional extended histological sections. The scale of the data (1.5 GB for the MR data, 1.4 TB for the histology stack) prohibits their viewing in full resolution on conventional display equipment (Plank *et al.* in press).

Our aim was, therefore, to develop a visualization environment, both hardware and software, that would allow scientists to examine as much of the image data in high resolution as possible, while maintaining context for the entire scene. This has led us to exploit high-resolution display technology capable of rendering over 50 million pixels. Equally important is the development of software that allows the use of this technology in a way that is both interactive and intuitive.

2. High-resolution displays

Standard computer display devices are designed for use 'at arm's length', for which a density of approximately 100 pixels per inch is generally regarded as optimal for viewing. This corresponds, for example, to a 17 in. monitor, with resolution of 1280×1024 (SXGA), just over 1 million pixels, or a megapixel. The problem for cardiac images with a resolution of $32\,000\times32\,000$, i.e. just over 1 Gpixel (10^9 pixels), is obvious: they contain two orders of magnitude greater detail than can be displayed on standard equipment.

The only option available today is to increase the physical size of the display. One possibility is to use projection; a single projector does not offer sufficient resolution (currently limited to 8 Mpixels), but arrays of projectors can be used to provide a tiled display. For example, the GigaPixel laboratory at Virginia Tech uses an array of VisBlocks (http://www.visbox.com) to provide a large screen, high-resolution facility: each of the 18 VisBlocks has resolution of 1280×720 pixels, totalling over 16 Mpixels. However, there remain limitations: despite successful research into blending, the alignment of projectors still poses a challenge; projectors are expensive; and considerable space is required between the projector and the viewing surface. Moreover, a person standing in front of the screen will cast a shadow unless back projection is used, which further increases demands on space and expense.

A cost-effective solution is to build an array of liquid crystal display (LCD) screens, arranged so as to provide a single large display surface. These tiled LCD panel displays are becoming increasingly popular: they allow a pixel density equivalent to a desktop monitor. This density is much higher than can be achieved for equivalent expense using projectors. Moreover, the space required is considerably less. This makes them attractive for applications such as biomedical image inspection. Another advantage of the tiled LCD panel is with respect to the brightness of the display, which is greater than that with a projection solution, allowing operation under normal room lighting.

A broad overview of high-resolution display technologies is given by Ni *et al.* (2006). In this paper, we focus on our own experience of using a tiled LCD panel display, the *LeedsWall*. This tiled display, shown in figure 1, comprises 28 flat panels, each of resolution 1600×1200, arranged in four rows of seven. This constitutes a 53.7 Mpixel display. The LCD panels are connected to seven computers, all equipped with two nVidia 7800 GTX graphics cards running two panels each. The computers are connected via gigabit (Gbit) Ethernet to each other and to a central filestore. The total hardware cost, including the custom-made stand, was under £30 000, which compares favourably with multi-projector solutions (Johnson *et al.* 2006). For full details of the construction of the wall, see Hodrien *et al.* (2007).

One disadvantage of the tiled LCD panel displays is the potential for distraction by the monitor frames, or bezels. We have experimented with two approaches to rendering. The first is to simply ignore the gaps, and render every pixel in the image leaving the user to neglect the borders. In our experience, users can do this easily, especially when immersed in a scene. The second option is to set up the display software to automatically adjust for the gaps, as though the user was looking through a window with bars across it. The disadvantage of this method is that additional processing power is required to set up the multiple viewports. Also, parts of the data are obscured from view. The latter may be an important consideration if the display wall is used to search for detailed targets that may be obscured. Mackinlay & Heer (2004) considered this problem in depth, arguing for a solution where the window metaphor is used so that geometry looks natural (diagonal lines will cross boundaries as straight lines with a gap, which the user finds easy to 'fill in'), but where any labels are always displayed in full.

The successful application of a large display wall to biomedical applications does depend not only on the screen hardware, but also on the input devices used to control applications. A standard keyboard and mouse would not be appropriate for users standing in front of the display, possibly walking along the length of the display. We have therefore experimented with a GyroMouse, a FrogPad keyboard, Flock of Birds controllers and a wireless games controller. For our cardiac application, we have used primarily the games controller joypad. This provides both analogue and digital control options, including a variety of configurable buttons. This has proved to be a suitably intuitive device for new users, regardless of whether they have experience of using similar input tools on computer games consoles.

3. Software for tiled displays

Several different packages support synchronized visualization across multiple displays, based on a range of computational resources. We have evaluated Chromium, Scalable Adaptive Graphics Environment (SAGE), Distributed Multihead X (DMX), Virtual Network Computer (VNC) and VR Juggler. The last package proved to be most suitable for our cardiac application.

Figure 1. View of a histological slice of a rabbit heart on the LeedsWall. Histology data courtesy of Rebecca Burton, Fleur Mason and Fahd Mahmood, University of Oxford.

CHROMIUM (Humphreys *et al.* 2002) is a tool for cluster-based rendering. It wraps the existing programs, intercepting the OpenGL calls and distributing them to the machines of the cluster for rendering. This not only minimizes the effort required to adapt the existing software for a tiled display, but it also imposes fairly important limitations. Significantly, CHROMIUM can be used only with OpenGL applications. In addition, we found that CHROMIUM is very demanding on the network, especially the head node, unless one uses display lists. When the network is heavily loaded, synchronization between the screens can be compromised.

SAGE (Jeong *et al.* 2006) was designed to provide a flexible environment for running multiple applications on high-resolution tiled displays. It allows for separation between the back-end systems that create visualizations and the front-end systems that render them. This separation, as with CHROMIUM, means that high-bandwidth interconnects are required for good performance. Even with 1 Gbit networking, we found that SAGE provided low frame rates and delayed interaction, and the performance appeared highly dependent on the resolution of the displayed windows, unlike non raster-based systems.

DMX (http://dmx.sourceforge.net) allows a logical X-server to be created across multiple X-servers. The entire wall can thus be treated as a desktop allowing windows to be dragged around freely. Our attempts to use DMX were not particularly successful: testing on four screens suggested potential, but with 28 screens the latency of interactions outweighed the benefits of this solution.

VNC (http://www.realvnc.com) provides a very simple method for rendering onto high-resolution displays. A single VNC server has an off-screen X-server, set to the size of the wall. Each cluster machine then runs a VNC viewer and