Garibaldi Project Report

Andries van Dam, Donnie Kendall, Ferdi Adeptura, Yudi Fu

Summary, September 2010

This report summarizes the work so far on the Garibaldi project. Our goal was to provide a research and visualization tool for a specific large piece of artwork – the Garibaldi Panorama. The work so far has been positively received – the application will be on show at the British Library as a centerpiece of their "Growing Knowledge – The Evolution of Research" exhibit for approximately 9 months beginning in mid October.

This project also led to the new research project also funded by Microsoft Research, code-named HumBub, which is shorthand for Bubbles for the Humanities. Bubbles are a light-weight, minimal-chrome, window-like container for fragments of code (the Code Bubbles system is also partially sponsored by MSR), as well as other text and multi-media content.

Large format art

The Garibaldi panorama is a massive piece of artwork, measuring nearly 300 feet long and 4.5 feet tall. The high resolution scans take up over 10GB of disk space. For phase 1 of the project we opted for the easy solution of scaling down the scans to make working with them easier. A software framework was developed to manage the 49 images of the scan and associated metadata for each individual scene plus additional documentation that was associated with each scene. The full panorama is described via a series of XML files, thus making the framework general enough to be used for other panoramic pieces of art. Note that the current UI and interaction paradigm assumes that the artwork is much wider than it is tall, so it is not currently suitable for other aspect ratios.

Collaboration

In order to facilitate collaborative exploration of the panorama an external display was connected to the Surface. This display does not merely mirror the Surface display, but instead displays a subset of content that is determined by the user of the Surface via simple flicking gestures. This allows for both collaborative exploration by allowing more people than can easily fit around the relatively small form-factor of the Surface to view the artwork simultaneously, and for teaching by allowing a professor to lecture while selectively displaying parts of the panorama to a class.

Exploration

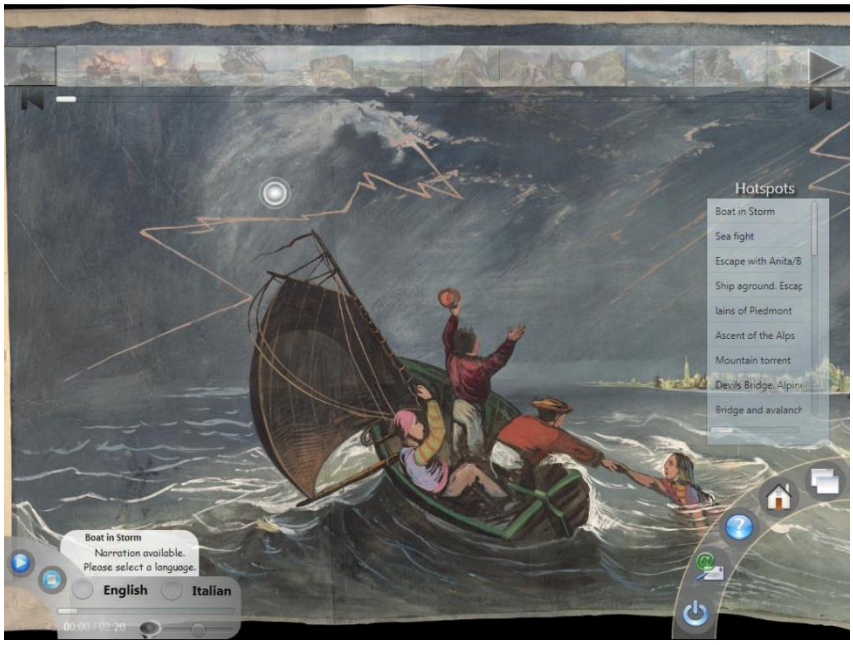

The Panorama can be explored in a variety of ways. The expected touch-based panning gesture is supported to allow browsing along the length of the panorama. A panning bar at the top shows your current location relative to the surrounding portion of the panorama and allows for quicker navigation by tapping or dragging along it. Pinch-zooming is supported to give a closer view of the current scene. Additionally, by using Surface byte tags physical props can be used to create a magnifying glass with adjustable zoom level, allowing a user to get a closer look at details while keeping the surrounding scene visible, something not possible with pinch-zooming.

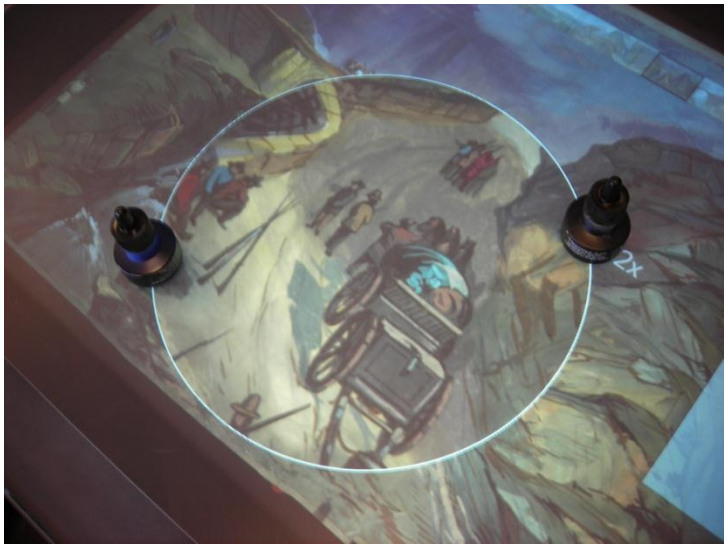

Users may also copy, or “clip” out a portion of the panorama to set aside for later reference by a two finger clipping gesture. These clips may be flicked to the external monitor for sharing and set aside for comparison to later parts of the panorama.

While exploring the panorama most scenes have one or more “hotspots” which contain scans of associated documents (most typically newspapers contemporary with the panorama). Touching one of these allows the user to view a list of the documents, and tapping on a document displays it. Documents may be interacted with via dragging, rotating, and pinch-zooming, and may be flicked to the external display for collaborative sharing.

Hover

By using the Surface’s raw video feed it is possible to see approximately 6 to 9 inches out from the display. A hover API is being developed to enable easier utilization of input which occurs before the user actually touches the Surface.

Development of this API is still in its early stages, with only a rough internal beta being done so far. Work is continuing as we see this as providing a very useful additional interaction technique beyond touch.

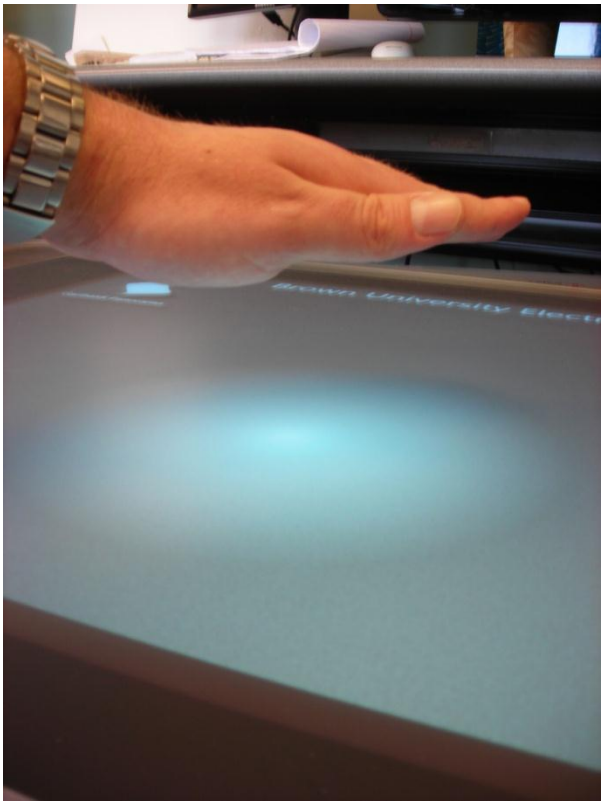

Annotation

Using work also done at Brown (in the Hands-on-Math project) on combining pen with touch input we also allow simple annotations of the panorama. This method of interaction is likely to be of interest to art historians and researchers.

Technical Hurdles

During development we ran into a few technical hurdles with the Surface API and the .NET runtime. First, due to the large amount of RAM that we consume on the Surface (which has a relatively small amount of installed memory) we had occasional problems with the memory manager reporting out of memory errors (even when RAM was still available). This was found to be due to the garbage collector not turning over free objects rapidly enough.

Additionally, some Surface controls have problems dealing with very large graphical objects. The panorama, even in its scaled down state, is on the order of 100,000 pixels wide. In the long term this is being fixed by incorporating Deep Zoom into the application.

Future and on-going work

Work is proceeding on what is being called “Garibaldi 2.0”. This name is a bit of a misnomer, however, as a central part of our goal is generalization to other hard-to-display artwork, including panoramas, tapestries, very large, and very small more traditionally formatted art. The work is progressing along several tracks:

High resolution

Our current solution to handling the large file sizes is to simply reduce them. This has obvious problems when zooming and scaling is allowed as you rapidly lose the crispness that you’d expect from the high resolution scans. We are incorporation Microsoft Deep Zoom technology onto the Surface to address this, allowing use of the full scan resolution.

User Experience

Currently there is a large amount of UI chrome on the Surface display to allow for all of the interaction possibilities. This tends to occlude parts of the artwork, which is undesirable. The UX is being redeveloped to reclaim lost screen space by reducing the always-visible chrome to a minimal subset

while still remaining as easy to use and discoverable as possible to allow for use by untrained walk-up visitors.

Additionally, many new interaction techniques are being explored and developed in an experimental setting – we wish to provide for a richer and deeper experience.

Additional Display and Interaction Devices

We are planning to utilize devices other than just the Surface, including laptops and Tablet PCs, and touch walls. In order to do this the software framework is being re-architected such that the UI is a drop-in component allowing for the different experiences that are needed for such vastly different devices.

Exploration and Search

The current version only allows for browsing, both by panning at various speeds and by jumping around the panorama. We plan to add search capability in order to allow users to find specific items that they may be looking for more easily, and additional modes of browsing, such as time-line or spatial browsing.

Garibaldi 1.0 Ongoing work

The display at the British Library provides for an amazing opportunity to gather usage data “in the wild”. The current version of the software is being instrumented so that we may learn how walk-up users interact with a touch display of this nature.