



i3 Lab

interaction | interfaces | invention

WaveWatch – Developing an Ambient Display of Web Traffic

**Ben Shelton
Keith Nesbitt**

ABOUT THIS SERIES

The i3 Lab is a cross-disciplinary, industry-focused research network currently based in the School of Electrical Engineering and Computing at the University of Newcastle, Australia. The principal aim of the i3 Lab Working Paper Series is to disseminate the research and/or technical outcomes of projects in an easily accessible format. The content of Working Papers generally falls into one of the following categories:

- *Preliminary findings or results, the release of which is intended to stimulate debate and/or discussion to assist in the further development of the research*
- *Technical reports associated with applied research that may be written in a less academic style than usually published in academic journals*
- *Extended versions of published works, often containing additional implementation/ application detail, figures and tables.*

The opinions or conclusions expressed in the Working Paper Series are those of the authors and do not necessarily reflect the views of the i3 Lab as a whole.



© Ben Shelton & Keith Nesbitt

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.



 Follow @i3_lab

Email i3Lab@newcastle.edu.au

Web <http://i3Lab.newcastle.edu.au>

WaveWatch – Developing an Ambient Display of Web Traffic

Ben Shelton

School of Design, Communication and IT
University of Newcastle,
Callaghan, 2308, NSW, Australia

Ben.Shelton@newcastle.edu.au

Keith Nesbitt

School of Design, Communication and IT
University of Newcastle,
Callaghan, 2308, NSW, Australia

Keith.Nesbitt@newcastle.edu.au

ABSTRACT

This article describes the iterative development of a novel Ambient Information System. The study was designed to investigate the use of 3D ocean waves as a metaphor for visualizing real-time web traffic. In this display twenty levels of web traffic were mapped to the frequency and turbulence of ocean waves. This report describes the various stages in the process used to develop the display and the various technical issues and risk confronted during development. It begins by describing the main business requirements, key design criteria and the data characterisation process used in the requirements gathering phase. It then describes the iterative prototyping stage of the project. This stage includes a description of the process followed for the 3D model design, the scene rendering, the evaluation of display aesthetics, the final data mappings, the system implementation and testing of three prototypes. Finally, the in-situ deployment and evaluation stage of the display is described. Using a pre-existing framework, the key issues associated with ambient display development are categorised and discussed. This discussion reports on key technology, deployment issues, and project management concerns. More general issues related to orchestrating such ubiquitous experiences and deployment of Ambient Displays in public spaces are also recorded. These findings provide a useful reference for other display developers wanting to understand some of the risks associated with design, development and deployment of such displays.

1 Introduction

Ambient Information Systems are intended as peripheral computing displays that try to visualize useful information seamlessly in the context of one's everyday environment. Intended as ubiquitous devices that might be present everywhere, Ambient Displays are typically designed to be aesthetically pleasing so they can be more easily integrated into a private or public environment and also function without interrupting the user's primary attention. In this sense they try to adhere to the concept of Calm Computing (Weiser and Brown, 1996).

Although a number of different Ambient Displays have been developed using a range of media there are still areas in the domain that require further investigation. This includes the difficult questions of how best to design and evaluate a system that is required to have both an aesthetic and informative function while not demanding the users undivided attention? In this paper we begin to address these questions in a case study that uses 3D graphics to visualize dynamic information in the user's periphery.

WaveWatch uses a novel wave metaphor, mapping the number of users on a webpage to the height and frequency of pre-rendered ocean waves in real-time. The metaphor is simple, if web traffic is low the ocean remains calm and as web traffic increases the ocean waves become more and more turbulent.

2 Determining Requirements

The aim in developing the *WaveWatch* display was to use aesthetically pleasing, high-definition, 3D rendered graphics in order to deliver useful but non-alarming information to users in a modern workplace. Although the deployed display device would be a large, 2D LCD screen, the use of 3D graphics was considered advantageous as it introduced more flexibility around the design of the visualization. We also believed that high-quality graphics would also match current user expectation around the fidelity of the imagery.

In this case study the aesthetic appeal of the imagery was considered to be equally important as the utility of the display to present information. The business goal of the display was not to create a visual data-mining tool but rather to invoke a general interest among staff in novel ways of monitoring website traffic.

Our initial design motivations introduced some immediate concerns. Firstly, the real-time rendering of 3D graphics is problematic as both wave modeling and photorealistic rendering are difficult to achieve quickly on current standards of computing hardware. Secondly, of more general concern to designers of Ambient Displays, was the potential conflict between designing for both aesthetics and utility. We wanted to create a pleasing relaxing ocean scene aligned to concepts of Informative Art (Redström et al, 2003) while still allowing users to faithfully recognize broad patterns of web usage.

To address the various design tradeoffs, we first determined the business requirements for the display and then adopted an iterative process of designing, implementing and testing prototypes (see Figure 1). After three iterations the final prototype was deployed in the work place for two weeks before being evaluated by staff using an existing heuristic evaluation approach (Mankoff et al, 2003) that addressed our design criteria.

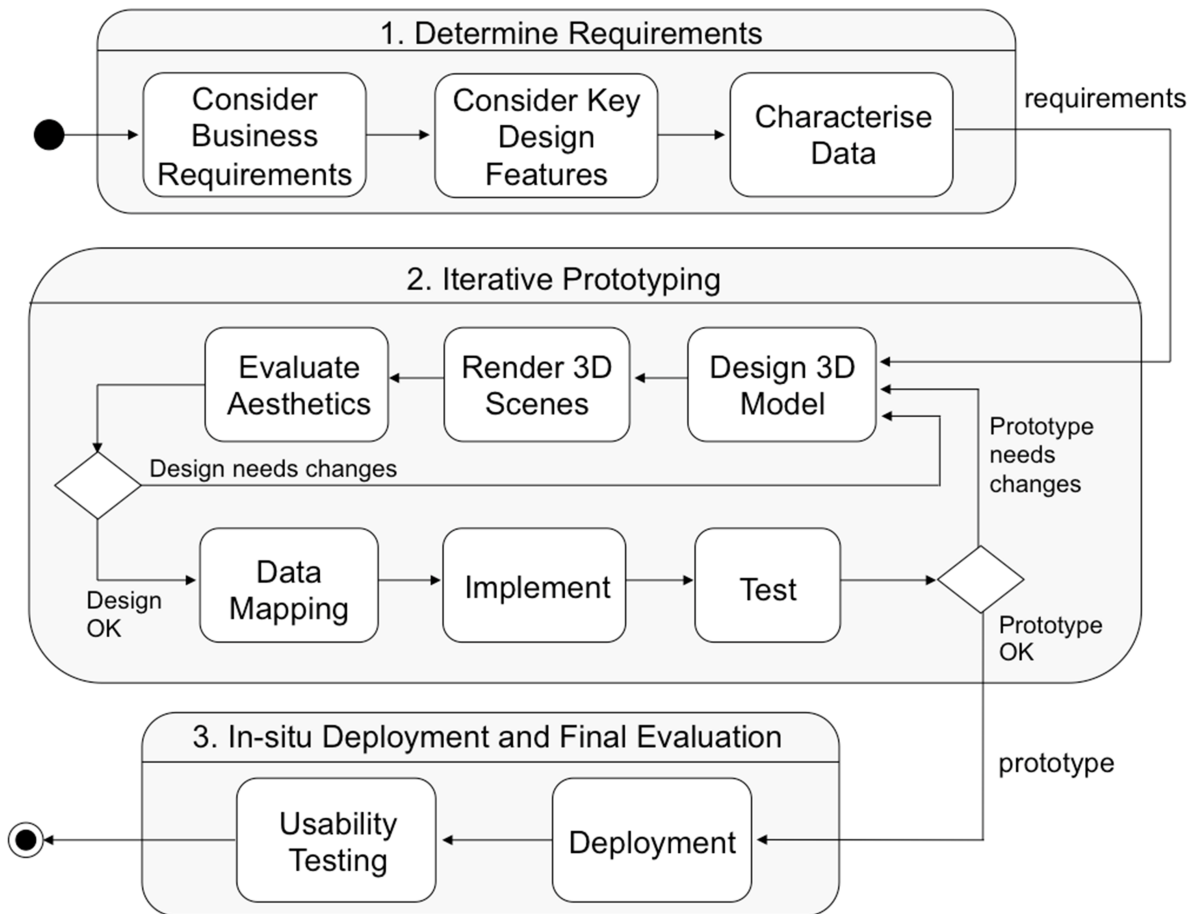


Figure 1: Development Process. A UML diagram showing the iterative process used to design and evaluate the WaveWatch display.

2.1 Business requirements

The WaveWatch display was sponsored by the Marketing and Public Relations department at The University of Newcastle (www.newcastle.edu.au). A key requirement for this group was that the display should visualize traffic levels for a specific webpage of interest. The display of “relevant data” to the proposed audience matched one of our own design goals. The group were also aware of the research nature of the project, and that the display was intended to raise awareness within the work group about patterns of web traffic. The group was also informed that the intention of the display was not to provide a diagnostic tool for formal analysis or monitoring of web traffic. It was not necessary for users to identify key patterns in the web traffic when using the display.

To be successful in a business sense, the display only needed to motivate interest and promote discussion on the topic of real-time monitoring of web traffic. This activity is especially applicable to the group during targeted web campaigns that are designed to increase web traffic to specific parts of the University’s website.

2.2 Key Design Features

The WaveWatch display was intended as a secondary or peripheral information source. For this reason, it needed to be assimilated into the environment and be aesthetically pleasing to look at. Thus the key design features of the display were to provide an aesthetically pleasing, non-alerting display that provided useful information to the workers in the deployment area. The display needed to run in “real-time” and adequately reflect current changes in web activity. The web site needed to be one of direct interest to the group. It was expected the display would motivate interest and discussion amongst workers without interfering with normal operational activities. Thus the display was primarily expected to serve a social function. It should stimulate a sense of discovery as well as active and passive engagement among the group (Memarovic et al, 2012).

2.3 Data Characterization

To understand the range of values expected on the website, traffic data was collected for a period of 14 days. A custom data scrapper, written using the Google Analytics live reporting API (<http://developers.google.com/analytics>) was used to determine the number of active users on the webpage at regular intervals. This data was then stored in a SQL Server database (<http://www.microsoft.com/en-au/server-cloud/products/sql-server/>) before being exported for analysis (see Figure 2, Figure 3).

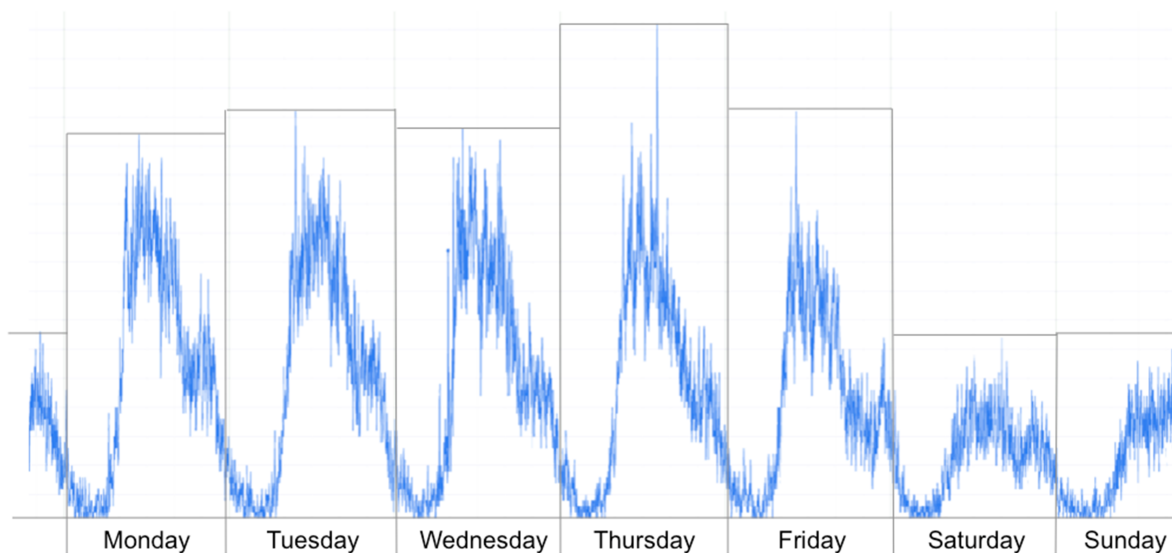


Figure 2: Weekly Web Traffic. Number of users of the website studied showing similar daily variations in traffic. Note the drop off on weekends.

The analysis showed a regular pattern of activity. During weekdays there was with quite a low number of users from midnight to around 8am in the morning when a rapid rise would occur. The traffic would remain high between 8am and 5pm, although large variations could occur during this period. After 5pm traffic would typically drop off again with another spike around 8pm before it began to fall away. Traffic on weekends followed a similar pattern although with lower volumes.

This preliminary analysis would be used during the data-mapping step to help decide the type of activity that needed to be reflected in the display. Since the final display would be deployed in a typical working environment we planned to customize the design for normal operational hours from 8am to 5pm, from Monday to Friday.

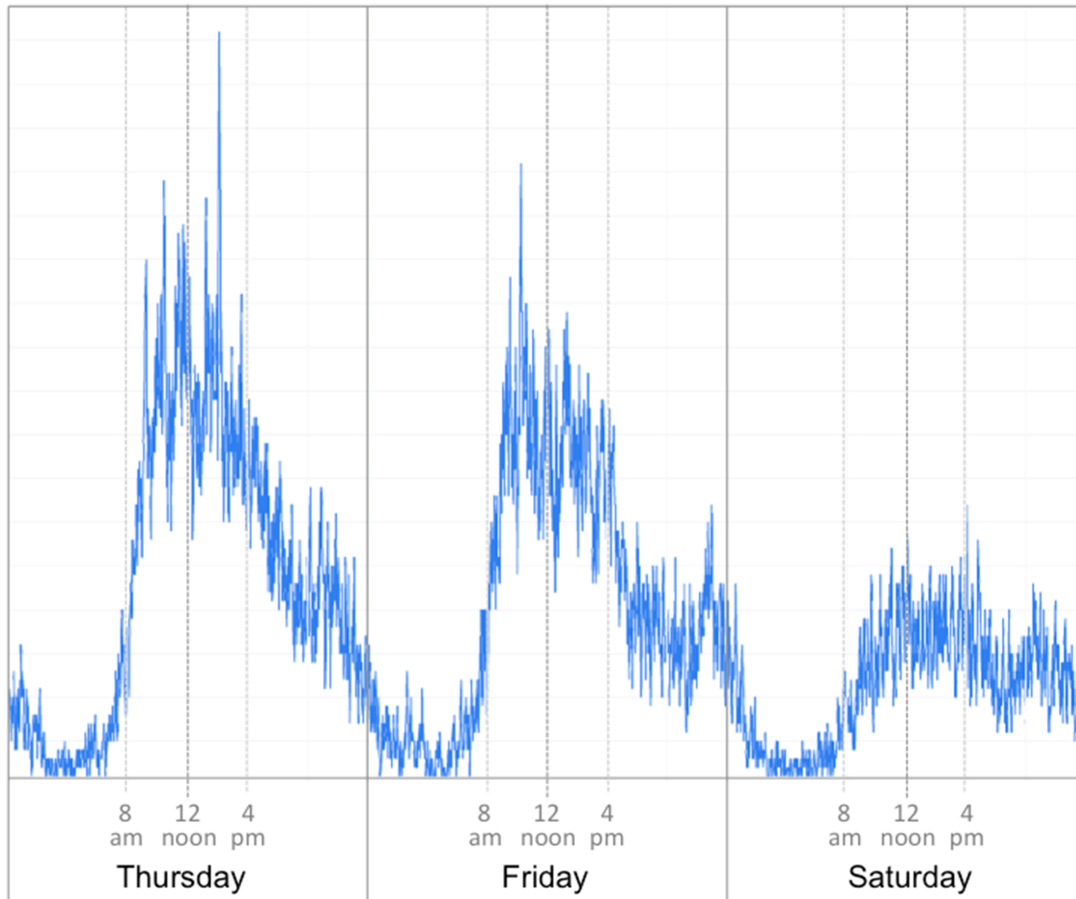


Figure 3: Intraday Traffic Variations. Typical variations in web traffic showing a rapid increase in activity occurs around 8am and a slower decline after 5pm. Also of note are the large variations that can occur during the day.

3 Iterative Prototyping

Having decided on our business goals, key design features of the display and completed an initial analysis of expected web traffic data, we next commenced the iterative design of the display. The key design vision was to provide an aesthetically pleasing, non-alerting display, based on a slightly abstract seascape showing a range of activity using ocean waves. While we report on each phase of the iterative development in a linear fashion, the actual process was more ad hoc, with considerable interplay between steps. Overall, the process was required to be agile, as we had only three months to develop the display. The general process aligns well with previous described for the rapid prototyping of pervasive displays (Fitton et al 2005).

Having decided on the broad design, namely a wave, the final scene composition was explored through three main design prototypes (see Figure 4). In the first two iterations the focus was on the 3D modelling, scene rendering and assessing the aesthetic qualities of the scene. By the third prototyping phase, the focus shifted to the data mapping, implementation and prototype testing activities.

3.1 Design 3D Model

Our initial intention was to use a real-time wave model that could generate 3D ocean waves based on simple parameters related to the amount of web traffic. However, the complexity of fluid dynamic models and the time required to render high quality animations made them problematic for use in our

real time display. Our pragmatic solution was to pre-render the wave scenes and leverage pre-existing tools for generating the waves.




Figure 4: WaveWatch Prototypes. The three alternative designs we developed, from the first (top) too final (bottom) version. Note we moved away from abstract representation to a more ‘natural’ seascape design.

The 3D wave models were created using the software Cinema4D (<http://www.maxon.net>). Cinema 4D is a flexible 3D graphics modelling tool that allows both the modelling and rendering of animated scenes. The background elements in the various prototypes such as the beach and palm tree, various rocks and shrubs were sourced as pre-existing models from Cinema4DDepot (<http://c4depot.com>).

The animated ocean of the display was created through a combination of two tools in Cinema 4D. The first of these, the noise tool was used to create the ocean chop of the scene. These chop effects were simulated by applying a Displacer object to a plane and modifying its Noise Shader attributes.

While creating the chop of the ocean was relatively simple, creating a cresting and breaking wave was more difficult. We tried a number of different methods before the final method was chosen. The early models utilized a plane with a cloth collider tag applied. Once animated a non-visible cylinder would ruffle the simulated cloth into a cresting wave shape. This method was ultimately abandoned due to its unrealistic look.

In place of this method a Cinema 4D plugin called Wave Deformer (Kollender, 2006) was used. This allowed us to create more realistic waves in Cinema 4D. Written in C++ the object generated by the plugin is applied to a plane that is then deformed into the shape of a wave. The plugin is powerful in that a number of parameters can be edited to change the appearance of the wave including twist amount, width, falloff, texture, texture offset and bend (see Table 1).

Parameter	Description
Twist Amount	Defines the overall size of the wave
Width	The overall width of the wave
Falloff	The distance from which the wave begins to dissipate along its horizontal axis
Texture	A texture that can be used to influence the wave's twist amount
Texture Offset	This parameter can be used to animate the applied texture to create more realistic waves
Bend	This parameter can be used to create curved waves.

Table 1: Wave Parameters. A description of parameters used in the of the Wave Deformer plugin (Kollender, 2015).

3.2 Render 3D Scenes

Our design goal was to generate animated scenes that would be displayed on a high definition LCD screen. Although we wanted a real time display it took some time for each wave to play out in the screen. Through trial and error, we decided on an animation length of least thirty seconds for each scene. This gave adequate time for a single wave (10 seconds) and also reduced unnecessary motion due to simply changing scenes too frequently.

Each of the thirty second scenes for the WaveWatch were rendered in full high definition (1920 x 1028) at thirty frames per second. Emerging as an issue throughout this process was the time to render each of the photorealistic scenes. This rendering process was lengthened by the high resolution of the scene, the high quality models, the high number of plane subdivisions required to animate the scene's ocean and the high quality horizon lighting. Although each of these elements increased the time needed to render each scene they each enhanced the overall aesthetics of the rendered output.

Overall, it took two months to compose and render the final 20 scenes we required using the consumer grade computer we had available. This was problematic when we wished to make even slight tweaks to the scenery, camera position or lighting conditions. Even subtle changes required considerable rendering time and it was not always clear if the final rendered outcomes would provide the aesthetic improvements we were aiming for.

3.3 Evaluate Aesthetics

The decision to use the natural metaphor of an ocean wave was not completely arbitrary. Anecdotally, it is not unusual for people to find themselves engaged in the activity of simply watching waves by the shoreline. Waves naturally seem to lend themselves to an aesthetic representation that is appreciated by a broad audience. Furthermore, the calming aspect of an ocean scene was considered to match well with the concept of Calm Technology (Weiser and Brown, 1996).

One of the developers involved in the project is also an artist and the initial design vision was broadly based on one of their previous paintings (see Figure 5). This is a similar approach used in early work described as Informative Art (Redström et al, 2000). In this project each successive design was evaluated by the same team member, their aim was to approach the style of the painting shown in Figure 5. This raises issues of how to best judge the aesthetic qualities of such displays as it involves subjective criteria that are difficult to formalize.

Even with a fairly clear aesthetic goal it was found that some pragmatic decisions equally influenced the final look we achieved. These decisions related to integrating the wave generator, reusing pre-existing tools and models, providing discrete levels of ocean activity and slow rendering times.

3.4 Data Mapping

When designing the data mapping we needed to consider the amount of variation we wanted to capture, given the typical characteristics of the web traffic and the desire to have noticeable variations occurring during the office hours when most workers would be present. By the same token we did not want to see large variations that might create distracting levels of movement in the periphery. We also had to consider the range of parameters, such as chop, wave height and frequency that were available to us in our modelling tools. Again after some trial and error we decided on 20 levels of activity that mapped from key levels of web traffic to a range of ocean turbulence (see Table 2, Figure 6, Figure 7).



Figure 5: Original Artwork. This painting, called “Sky and Sea” set the original direction for the look of the final display (www.knesbitt.com/Paintings/sky_sea.htm)

	Web Activity Level	Chop Height (cm)	Wave Height (cm)	Wave Frequency (30 seconds)
Low	1	0.70	0	0
	2	0.95		
	3	1.10		
	4	1.20		
	5	1.35		
	6	1.60		
Medium	7	1.60	1	1
	8		2	
	9		3	
	10		4	
	11		5	
	12		6	
High	13	1.60	7	3
	14		8	
	15		9	
	16		10	
	17		11	
	18		12	
	19		15	
	20		18	

Table 2: Display Levels. There are 20 levels of ocean activity. These are distinguished by changing chop, wave height and wave frequency over the 30 second period. Note the scale, in centimetres, reflects the underlying 3D model's coordinate system.

3.5 Implementation of Prototypes

Once the data mapping was decided upon and each of the 20 animated levels rendered as mp4s, the focus shifted to more careful development of the software required to run the display. This software needed to track the number of users on a specific web page over a predefined period of time and then select the appropriate scene, one of the 20, and display it. We wanted this process to loop continuously in a seamless fashion.

The final prototype was created in Microsoft Visual Studio using the Windows Presentation Foundation (WPF) in C# (Msdn.microsoft.com). It required around 200 hours to implement and test. The Google Analytics Real Time Reporting API (<https://developers.google.com/analytics/>) was used to build the real-time data scrapper. This allowed us to gather a number of traffic related website metrics including the number of active users on a specific web page.

The application utilizes a multithreaded model. One thread collects the live web traffic data over a 60 second period. The second thread analyses the data, selects the appropriate wave animation and displays it (See Figure 8). This second thread checks every 30 seconds for the first data collection thread to complete. If the data collection thread is not complete it replays the current animation, otherwise it analyses the collected data and chooses the appropriate animation level for display. Note that this has the effect of always repeating the same 30 second animation at least twice once it is selected.

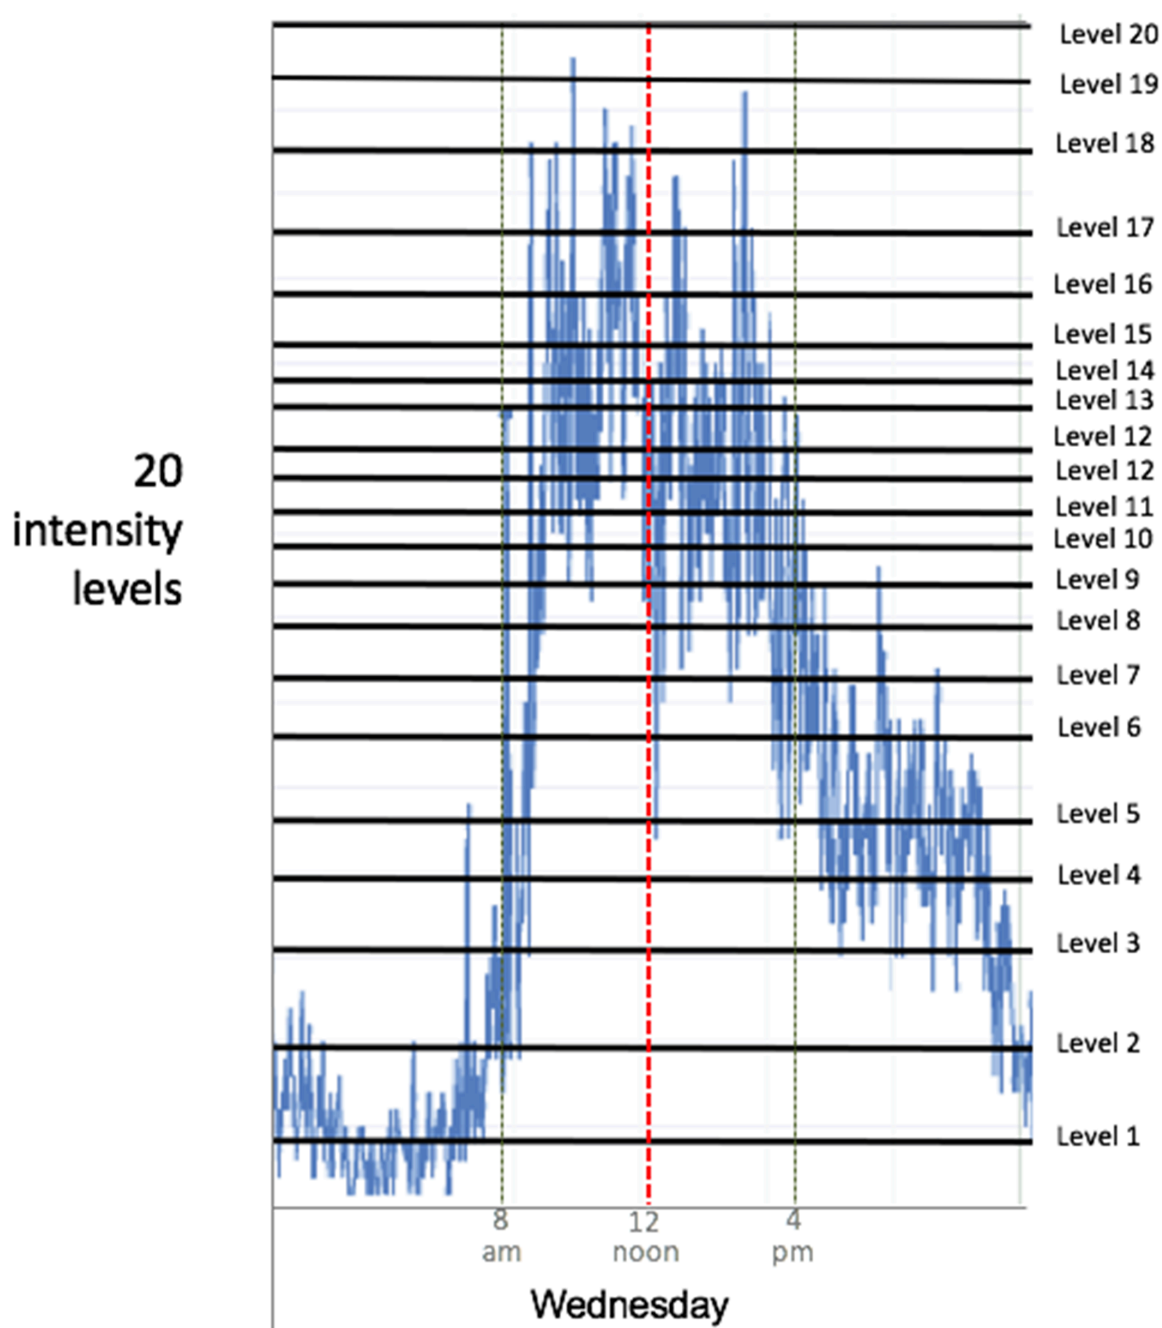


Figure 6: Activity levels. Note the 20 levels of web traffic are not linearly distributed. The levels (4-17) are closer together to capture more subtle variations expected during office hours.




Figure 7: Wave Activity. Shown in this figure are three representative levels of ocean turbulence. At the top is a low amount of web traffic (level 3). In the middle is a more medium range of web traffic (level 10). The bottom image shows high traffic (level 18).

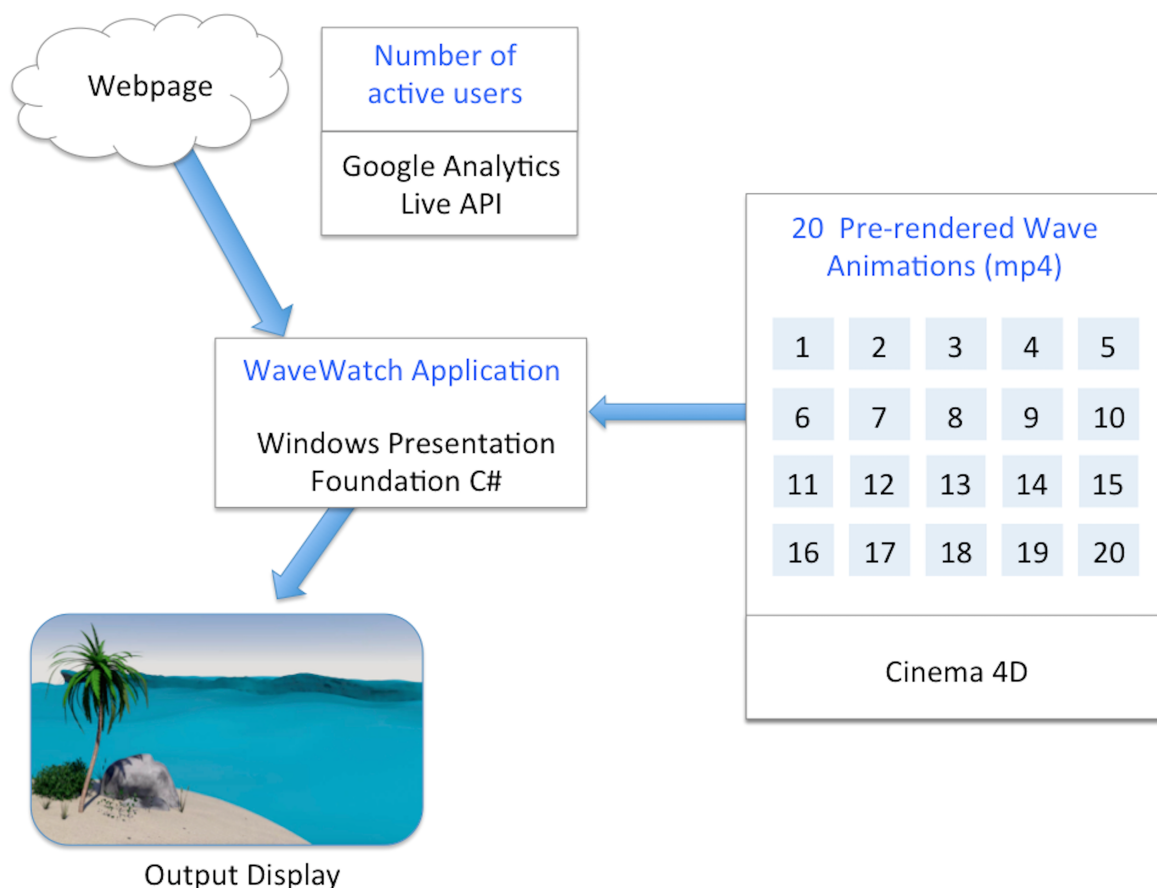


Figure 8: WaveWatch Software. WaveWatch implements a multi-threaded loop. One thread scraps data to determine the number of current active users. The second thread analyses this data then selects, loads and plays the appropriate animation level.

3.6 Testing Prototypes

An important issue identified at this stage was the difficulty of creating a seamless transition between the various animation levels. For example, when transitioning between level 9 and level 10, differing chop heights on the last frame of level 9 and the first frame of level 10 would not match. To overcome this, we standardized chop for three seconds at the beginning and end of each animation. This small modification was important to avoid distracting users with sudden discontinuities in the ocean scene.

We encountered a further problem when trying to load and play different scene levels using the WFP C# application. The initial method we used was to simply add a video to a WFP application using the WFP MediaElement (Msdn.microsoft.com). However, we encountered problems when trying to switch from one activity level to the next. Intuitively, we simply changed the source of the MediaElement whenever a new animation level was needed. However, the MediaElement accomplishes this by first removing its currently playing video and then loading the next video before displaying it. This resulted in a disconcerting blank screen for about half a second while the new video was loaded.

Eventually, we overcame this problem by a work around using two stacked MediaElements in a single WFP application. These two MediaElements were placed one on top of the other in the application's interface. When a new animation level was needed we would load it in the

MediaElement that was hidden. When ready this second player would then be switched so it was in front of the previously active player. This gave the effect of a completely seamless cut between the two animations.

During this stage we also tested various time periods for switching between scene levels. Initially we planned to collect data for thirty seconds matching the scene change to each 30 second animation. However, when inspecting the display, it was found that shifting the wave level every 30 seconds created unintentional noise in the display. The ocean tended to appear unnaturally chaotic. This was because we often encountered spikes in activity when averaging over 30 second periods. To overcome this problem, we averaged the number of users over 60 seconds. This helped remove the spikes and settle the ocean down. Choosing a 60 second period also allowed us to simply replay the same 30 second animation twice before shifting to the next scene level. These changes created a slightly calmer display. Subjectively, it also seemed to tune the display dynamics, giving the performance a more natural ocean motion.

4 Deployment and Evaluation

The WaveWatch had been designed for deployment into the Marketing and Public Relations (M&PR) office at The University of Newcastle, Australia. The local area managers supported all phases of the design, development, deployment and evaluation of the display. The research was also approved through the official research ethics process used by the University.

4.1 Deployment of WaveWatch

WaveWatch was installed on a consumer grade computer connected to the standard University network. Official permission and special access rights were obtained to access the data related to traffic information. The output was displayed on a pre-existing, 42-inch LCD display. The monitor had been wall-mounted to afford visibility to the workers in the open plan office. Although the display was pre-existing it was infrequently used before the deployment of WaveWatch. Workers in the area were informed one month before the deployment, when they were given a brief description of the display and its intention. There was also an official launch of the display on the first morning of operation when workers were invited to ask questions.

Following the launch, the display remained continuously active for two weeks. The display operated continuously even though most workers in the area would only see the display during normal work hours. During brief periods of network maintenance, no new data was collected so the animation simply looped on the current level until the network became active again. During office hours the display was regularly inspected for problems but there were no disruptions to its operation over the two-week deployment period.

4.2 Usability Testing

After the two-week deployment of WaveWatch a follow up email was sent to the staff giving them the chance to anonymously complete an evaluation questionnaire. The questionnaire comprised a number of Likert scale and open-ended questions. These questions were specially developed from a heuristic evaluation designed for assessing the usability of Ambient Displays (Mankoff et al, 2003). The response rate for completing this evaluation was 65% (n=15/23).

Generally, most workers in the area found the WaveWatch to be useful and easy to use. With 87% of respondents answering positively when asked to reflect on the utility of the data. Of the responders, 93% rated the display as both unobtrusive and non-interrupting to their everyday work. These

responses suggest that the display met our expectations of Calm Computing (Weiser and Brown, 1996).

Despite this, 87% of respondents also agreed that the WaveWatch was able to attract the attention of others, particularly when large waves occurred. In reviewing the responses to the related open-ended question, it was found that social factors were also an important reason for the display gaining attention. The display was located near the office entrance and acted as a zone of both passive and active engagement while also enticing the curiosity for discovery in workers (see Memarovic et al for a discussion of these social factors in pervasive displays).

In all 93.3% of participants responded that the WaveWatch did not display too much information. While this was also part of our design goal, many participants commented that they would like to see a lot of additional data. This included more traditional visualization devices such as scales and temporal markers and even personalized interactivity. Further research will be required to determine the optimum number of data sources to find an acceptable mix between the amount of information and the calmness of such a display.

It was found that the use of an ocean wave metaphor to convey real time data was easily interpreted, with 80% of respondents agreeing that they were easily able to understand the information presented. Indeed 93% of respondents reported being able to interpret the display at a glance. Furthermore, 93% of respondents also reported they thought the display was aesthetically pleasing. Several participants made further positive comments related to the aesthetic and pleasant nature of the display.

5 Discussion

In general, the deployment of the WaveWatch display was successful, with our key business and design goals being met. In this discussion we use a previous list of 13 lessons learned from developing pervasive displays (Storz, 2006) to reflect on our own project experience.

5.1 Technology and deployment

“1. Never underestimate the effort involved in creating real system deployments” (Storz et al, 2006). We found this lesson to be most relevant. The deployment required unexpected hurdles such as reviewing the project for occupational health and safety concerns related to dynamic visual content. Although we used a pre-existing monitor for display purposes, in the end we needed to provide our own carefully configured computer to drive the display. We also underestimated the time required to gain signoff on using potentially sensitive web traffic data, which also raised possible security concerns when accessing the secure network.

“2. Never underestimate the impact of environmental factors on a deployment” (Storz et al, 2006). As our deployment occurred in doors on a previously existing display we experienced no major problems in this area. However, one design feature we had considered, namely an auditory overlay of ocean sounds was abandoned quite early in the project due to concerns that it would create too much noise.

“3. Deploy for maintainability and change” (Storz et al, 2006). We had planned only for a two-week deployment and ignored at least one significant issue in our prototype design. We assumed relatively constant levels of web traffic during the deployment period. We would need to investigate more automated approaches to scaling the data for the display if the deployment was ongoing. As the display was in a University environment it would generally be important to also consider seasonal variations in traffic, for example, variations that might occur in, and out of, teaching periods.

“4. Anticipate and plan for regulatory compliance issues” (Storz et al, 2006). Many of these issues were addressed early on in the project design. As a matter of course we involved all relevant workers in an open communication process. This helped to gain important management input and vital support for the project. A formal ethics approval process requiring management signoff also helped address potential compliance issues.

5.2 Monitoring and management

“5. Ensure that it’s possible to remotely monitor the system’s output as the user perceives it” (Storz et al, 2006). This was not a concern for our project as one of the researchers was situated near the display during the deployment phase.

“6. Provide tools and abstractions that let individuals reason about the system’s internal state” (Storz et al, 2006). This was indeed a key design criterion in our project and we had considerable knowledge in the research team related to visualization design. We also evaluated the deployment using this criterion. Even so, a few workers in the area commented that they would like to use the tool in a more formal analytical way by seeing scales and actual numbers as part of the display.

5.3 Content creation and management

“7. Never underestimate the importance of content” (Storz et al, 2006). As we have noted, the look and feel of the content, the aesthetics of the display was a primary concern during our design, development and evaluation.

“8. Set aside adequate resources for content creation” (Storz et al, 2006). Despite our strong focus on content we could definitely have used more time and the input of experienced designers and modelers to help refine the content. We believe we only achieved about 70-80-% of our initial vision in the final prototype. Time and resource pressures definitely resulted in some ‘pragmatic’ rather than ‘ideal’ design decisions.

“9. Managing content for pervasive computing is a major task for which existing tools are poorly suited” (Storz et al, 2006). Although we managed to prototype a reasonably robust solution in the end it did require us to overcome a number of problems in regard to seamlessly looping and switching videos. We had expected this functionality would be provided as a matter of course in the media API.

5.4 Orchestrating ubiquitous computing experiences

“10. Ensure that public deployments can support orchestrated performances (Storz et al, 2006).” We had little control over the actual performance of the display. However, we did make a design change to slow down activity on the display. This in effect asserted a subtle influence on the performance by making the dynamics more naturalistic.

“11. Ensure that users don’t experience partial or inconsistent changes in system state” (Storz, 2006). We tried to address this concern by testing a live system for one week before Storz et al. Fortunately, no unexpected changes or inconsistent states occurred during the deployment phase. We had anticipated unexpected network outages by ensuring the display simply looped the current animation when network problems were detected.

5.5 Working in public spaces

“12. Managing user expectations is crucial in public ubicomp deployments” (Storz et al, 2006). Having worked in many previous software deployments, we had realized the importance of this when

the project was conceived. We therefore addressed this issue at an early phase of the project by seeking input from workers in the area as soon as the project began. We believe this provided much needed support to ensure the deployment occurred as planned.

“13. Prepare yourself, your team, and your work for public scrutiny” (Storz et al, 2006). We were fortunate in that the display was deployed amongst “friends”. The workers in the area had been involved in the project from early on the development. Even so a number of critical comments were received. Many of these were insightful and useful, while others focused on individual preferences for the content.

6 Conclusion

The outcomes from the project were generally successful. For example, the department has since been granted further funding to extend the display to a broader University audience. We would also like to extend the work to address further research questions.

We would like to investigate the best way to automate the data characterization process where seasonal variations can be expected. This might be by way of tracking moving averages in the data and devising a more dynamic mapping to the scene levels. While pre-rendered scenes worked well in this case study we would still like to investigate how to directly model and render 3D breaking waves using algorithms from fluid dynamics. This is expected to introduce numerous challenges if a real-time display is to be maintained. We would also like to adapt the wave metaphor used in this display for other types of dynamic data including stock market activity, real-time twitter feeds and relevant performance indicators. Finally, we would like to trial some objective, reaction time tests in the laboratory to better understand the cognitive workload associated with peripheral displays such as WaveWatch.

7 References

- D. Fitton, K. Cheverst, C. Kray, A. Dix, M. Rouncefield and G. Saslis-Lagoudakis “Rapid prototyping and user-centered design of interactive display-based systems”, *Pervasive Computing*, IEEE, vol. 4, no. 4, pp. 58--66, 2005.
- J. Kollender, 'Wave Deformer for C4D', Vimeo, 2014. [Online]. Available: <https://vimeo.com/91515753>. [Accessed: 30- Sep- 2015].
- J. Mankoff, A. Dey, G. Hsieh, J. Kientz, S. Lederer and M. Ames, “Heuristic evaluation of ambient displays”, in *SIGCHI conference on Human factors in computing systems*, 2003, pp. 169--176.
- N. Memarovic, M. Langheinrich, F. Alt, I. Elhart, S. Hosio and E. Rubegni, “Using public displays to stimulate passive engagement, active engagement, and discovery in public space”, in *4th Media Architecture Biennale Conference*, 2012, pp. 55--64.
- J. Redström, T. Skog and L. Hallnäs, “Informative art: using amplified artworks as information displays”, in *DARE 2000 on Designing augmented reality environments*, 2000, pp. 103--114.
- O. Storz, A. Friday, N. Davies, J. Finney, C. Sas and J. Sheridan, “Public ubiquitous computing systems: Lessons from the e-campus display deployments”, *Pervasive Computing*, IEEE, vol. 5, no. 3, pp. 40--47, 2006.
- M. Weiser and J. Brown, “Designing calm technology”, *PowerGrid Journal*, vol. 1, no. 1, pp. 75--85, Jul, 1996.