

TABLETOP DISPLAY MAGNETS

by

Mark Dennis

Submitted in partial fulfillment of the
requirements for the degree of
Bachelor of Computer Science, Honours

at

Dalhousie University
Halifax, Nova Scotia
April 2008

© Copyright by Mark Dennis, 2008

DALHOUSIE UNIVERSITY

FACULTY OF COMPUTER SCIENCE

The undersigned hereby certify that they have read and recommend to the Faculty of Computer Science for acceptance a thesis entitled “TABLETOP DISPLAY MAGNETS” by Mark Dennis in partial fulfillment of the requirements for the degree of Bachelor of Computer Science, Honours.

Dated: April 28, 2008

Supervisors:

Dr. Stephen Brooks

Dr. Kori Inkpen

DALHOUSIE UNIVERSITY

DATE: April 28, 2008

AUTHOR: Mark Dennis

TITLE: TABLETOP DISPLAY MAGNETS

DEPARTMENT OR SCHOOL: Faculty of Computer Science

DEGREE: B.C.Sc. (Honours) CONVOCATION: April YEAR: 2008

Permission is herewith granted to Dalhousie University to circulate and to have copied for non-commercial purposes, at its discretion, the above title upon the request of individuals or institutions.

Signature of Author

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

The author attests that permission has been obtained for the use of any copyrighted material appearing in the thesis (other than brief excerpts requiring only proper acknowledgement in scholarly writing) and that all such use is clearly acknowledged.

Table of Contents

List of Figures	vi
Abstract	ix
Acknowledgements	x
Chapter 1 Introduction	1
1.1 Motivation	1
1.2 Goal	2
1.3 Outline	2
Chapter 2 Related Work	3
2.1 Storing Items: Storage Bins	3
2.2 Reaching Items: Interface Currents	4
2.3 Moving and Displaying Items: Mobile Spatial Tools	5
Chapter 3 Design	7
3.1 Item Attraction	7
3.2 Item Sorting & Grouping	9
3.3 Moving Items with Magnets	13
3.4 Displaying & Aligning Items	16
3.5 Magnet Layout	19
Chapter 4 Implementation	21
4.1 Tabletop Displays	21
4.2 University of Calgary Large Display Framework	24
4.3 Magnet Strategy, Attracted Component, & Attraction Manager	27
4.3.1 Item Attraction & Repulsion	28
4.3.2 Item Sorting & Grouping	30
4.3.3 Moving Items with Magnets	34

4.3.4	Displaying & Aligning Items	35
4.3.5	Magnet Layout	38
Chapter 5	Results & Scenarios	40
5.1	Results	40
5.2	Scenarios	44
5.2.1	Scenario 1: Real Estate	44
5.2.2	Scenario 2: Criminal Identification	45
Chapter 6	Conclusions	48
Chapter 7	Future Work	49
Bibliography	51

List of Figures

Figure 3.1	Magnet attracting items in the workspace.	8
Figure 3.2	Multiple magnets attracting different subsets of items.	10
Figure 3.3	Magnet with threshold. Process of attracting an item to a specific magnet.	11
Figure 3.4	(a) Multiple magnets and items without edges drawn showing group memberships; (b) Multiple magnets and items with edges drawn to show group memberships.	12
Figure 3.5	Process for breaking the attraction between a magnet and an item. The blue edge represents the attraction between Magnet 1 and the item. The red interrupted line shows the path that a user would draw to remove the attraction between the magnet and the item.	13
Figure 3.6	Magnet showing the dial for increasing and decreasing the magnetic pull that a given magnet has on items in the workspace. Turning the dial counter-clockwise decreases the magnetic force, while turning the dial clockwise increases the magnetic force.	15
Figure 3.7	(a) Magnet showing a large force halo pulls items at with greater force; (b) Magnet with smaller force halo pulls items with a lesser force.	16
Figure 3.8	(a) Magnet showing a blue halo representing a positive magnetic force; (b) Magnet showing a red halo with the same magnitude but representing a negative magnetic force.	17
Figure 3.9	(a) Magnet organizing a group of items into a pile; (b) Magnet aligning a group of items radially; (c) Magnet organizing a group of items into a grid.	18
Figure 3.10	A more realistic rendering of what the design of the tabletop display magnet looks like. This shows the dial, threshold, force halo, and toggle buttons attributed to these magnets as discussed in the design.	20
Figure 4.1	Mitsubishi Engineering Research Laboratory DiamondTouch display features signal receivers for each user which are capable of detecting inputs from particular users.	22

Figure 4.2	Microsoft Surface features a multi-touch display that can recognize a variety of inputs from fingers to objects such as cell phones, digital cameras, and credit cards.	23
Figure 4.3	SMART Technologies prototype tabletop display at Dalhousie University has a 77-inch display which supports up to two simultaneous inputs.	24
Figure 4.4	Example of how an i-buffer is used to effect the display of visualized components. (a) Shows the values of the buffer that could control the colour of visualized components; (b) Shows the effects of such a buffer (a ‘0’ value turns a rectangle yellow; a ‘1’ value turns a rectangle blue) [14].	25
Figure 4.5	<i>Composition Tree</i> showing the background component at the root and storage bins that contain two and three children images	26
Figure 4.6	Representation of a magnet’s threshold boundary. The threshold defines a circle with a given radius, r . By finding the Euclidean distance, d , between the centre of the magnet, \mathbf{p} and the point at which the user released the item, \mathbf{q} , and it can be determined whether or not an item has been dropped past the magnet’s threshold.	31
Figure 4.7	(a) Initial alignment of edge endpoints \mathbf{p} and \mathbf{q} , and input points \mathbf{i} and \mathbf{j} ; (b) Alignment of edge endpoints \mathbf{p} and \mathbf{q} , and input points \mathbf{i} and \mathbf{j} after being rotated by $-\theta$ degrees	34
Figure 4.8	The angles between the items displayed in the radial alignment is $\frac{2\pi}{M}$, where M is the maximum number of items that can be attracted to a single magnet. In this case $M = 8$	36
Figure 4.9	The order in which nine items are organized into a grid arrangement.	37
Figure 4.10	Screenshot of the actual magnets, with the dial, centre texture, and toggle buttons labeled.	39
Figure 5.1	(a) Shows a halo that visualizes a small positive pull; (b) Shows a halo that visualizes a large positive pull; (c) Shows a negative pull, or a repelling magnet.	41
Figure 5.2	Screenshot of magnet with edges connecting items that it attracts.	42
Figure 5.3	Screenshot of two magnets with edges turned on, showing the memberships of each item relative to the magnets in the workspace.	42

Figure 5.4	(a) Screenshot of magnet with pile alignment turned on; (b) Screenshot of magnet with radial alignment turned on; (c) Screenshot of magnet with grid alignment turned on.	43
Figure 5.5	Screenshot of the tabletop display magnet system being used to find a particular residential property with a desired set of attributes. The magnet on the left attracts those properties that have pools, and the magnet on the right attracts properties that are between \$250,000 and \$300,000. The two magnets working together allow for the properties with both of these attributes to be identified easily.	45
Figure 5.6	Screenshot of the tabletop display magnet system being used to find a particular person with a known set of attributes. The magnet on the left attracts those people with facial hair, and the magnet on the right attracts those people who wear glasses. The two magnets working together allow for the people who wear glasses and have facial hair to be identified easily.	47

Abstract

Tabletop computing on large tabletop displays has existed for several years. However, the tabletop computer has yet to be established as a vital piece of technology that has significantly improved the way people work and interact with each other. Interactivity on tabletop displays has, in the past, been described as awkward and inefficient. Many tools for overcoming the limitations of tabletop displays have been developed. However, very few tools are capable of being used to accomplish more than one particular task, which in itself is a limitation. As a result there is a need for a small set of tools, capable of accomplishing multiple tasks on a tabletop display. This thesis describes the design and implementation of the tabletop display magnet, a multipurpose tool used to reach, sort, move, and display data items in an efficient and less awkward way on tabletop displays.

Acknowledgements

Firstly, I would like to thank Dr. Stephen Brooks, Dr. Kori Inkpen, and Mr. Derek Reilly, from the Faculty of Computer Science at Dalhousie University for their help, guidance, and support while working on this project. Secondly, I would like to thank Dr. Sheelagh Carpendale, Mr. Fabricio Anastasio, Ms. Julie Stromer, and Mr. Matthew Tobiasz, from the Innovations in Visualization Laboratory at the University of Calgary for their Large Display Framework, which has proved to be invaluable. I would also like to thank Mr. Chris Maxwell for his help with the tabletop display.

Chapter 1

Introduction

1.1 Motivation

Personal computers have gained significant popularity over the past few decades. Personal computers have not only made people work more efficiently, but have also provided new means of entertainment and interactivity. Although, much of the success of personal computing is due to the capability of building powerful personal computers at low costs, the tools that have been created to improve a user's experience on a personal computer should be given a great deal of credit as well. For example, in terms of hardware, the computer mouse, invented by Dr. Douglas C. Engelbart, has been a fundamental invention in improving the way people work with personal computers [1]. In terms of software, advances in graphical user interfaces such as the advent of windows, icons, menus, and pointing devices, have also been important for personal computing.

Tabletop computing on large tabletop displays has existed for several years. However, the tabletop computer has yet to be established as a vital piece of technology that has significantly improved the way people work, and interact with each other. This may be somewhat surprising since much of the work that people have done for millennia, especially collaborative work, has been done on physical tabletops. Work that has been done on tabletop computers has, in the past, proven to be inconvenient and awkward [18]. This is also fairly surprising since it might be assumed that the large size of tabletop displays would allow ample space for people to work very efficiently.

In order to leverage the possibilities of tabletop computing, many tools have been developed in order to reduce the awkwardness of working on large tabletop displays. However, presenting users with a large array of tools may complicate the tabletop working environment. On the other hand, presenting users with only a small subset of these tools could limit the tasks that can be accomplished around the tabletop

display. As a result, a small set of tools is needed to allow users to perform many basic tasks enabling them to work efficiently on a tabletop display without having to become familiarized with a large array of tool.

1.2 Goal

The goal of this thesis is to create a tool, called a tabletop display magnet, that can be used in a collaborative tabletop display environment. We propose that the tool may help users overcome some of the major limitations that are found in organizing data on a tabletop display. The first limitation that it could overcome, is the problem of users being incapable of reaching data items because of the typically large size of tabletop displays. It could improve the ways that users are able to organize data items on a tabletop display, by allowing them to create groups and subsets from the data. Furthermore, the tool could provide improved ways of moving and displaying data on the tabletop display.

1.3 Outline

In Chapter 2, this report will begin by discussing some of the work that has been previously conducted to overcome the limitations with large tabletop displays. In Chapter 3, the design of the tabletop display magnet will be outlined and discussed. Each feature of the magnet will be detailed. Chapter 4 discusses how the design was implemented. It will include some of the mathematics involved in creating the tabletop display magnet. The SmartBoard prototype tabletop display that the tabletop display magnets were specifically designed for will also be presented and compared to other available tabletop displays. Chapter 4 will also discuss the Large Display Framework, created by the Innovations in Visualization Laboratory, which has played a fundamental role in the development of Tabletop Display Magnets. Chapter 5 will look at the results of the research that has been done. Chapter 5 will also discuss some scenarios in which tabletop display magnets could prove to be useful. This paper will close with conclusions about the system in Chapter 6 and a discussion of some possible future work to be done in Chapter 7.

Chapter 2

Related Work

Tabletop display magnets bring together many notions and ideas that have been developed for tools used in collaborative tabletop settings. Therefore, a tabletop display magnet can be regarded as a multipurpose tool that solves many problems that individual tools have been developed in the past to solve. The problems that these magnets try to solve include storing and grouping data, moving and displaying data, and reaching data over large distances in a collaborative tabletop display setting. This chapter will give an overview of some previous work that has been conducted in order to solve these problems listed here.

2.1 Storing Items: Storage Bins

Using tools such as “folders” as they would appear in a single user graphical user interface (GUI) on a personal computer have limitations that could obstruct collaboration between users in a tabletop display environment. Collins, in his paper “Exploring Tabletop File System Interfaces”, explains how the hierarchical nature of traditional file systems are inefficient [4]. He states that collaboration around a tabletop is much more visible than traditional file systems.

In *Storage Bins: Mobile Storage for Collaborative Tabletop Displays*, Scott et al. [18], describe their design and implementation of storage bins that came about from studies they had conducted on tabletop collaboration. They state that people around a tabletop setting tend to loosely group items together. Furthermore, they state that, at times, users will decide to remove items from such groups or add items that previously did not exist into those groups. The storage bins devised by Scott et al., allow users to organize these items into such groups by simply dropping items into the storage bins. These storage bins are scalable, which means that they are able to accommodate a variable amount of items. Furthermore, when an item is added to a storage bin, the item shrinks in size. As a result this would allow for more of the

workspace to be used when unneeded objects are kept in the storage bins. Storage bins share this capability with ZoomScapes. ZoomScapes are designated areas that are arranged around the periphery of the display. Like storage bins, when items are brought within the area, the items shrink in size [6]. Scalable Fabric, a prototype desktop application from Microsoft, further explores the notion of items shrinking by a certain amount depending on their locations. It works by allowing the user to designate a space such that the size of each window on the desktop is determined by how far the window is from that space [16].

The storage bins can also be moved around within the workspace. Therefore, when the users move storage bins in front of or near themselves the items within those storage bins can be easily placed within their personal territory. Conversely, storage bins can be placed in more central locations, which in turn would put those items in the entire group's territory. This means that storage bins not only allow for items to be grouped together, but they also allow for the items to be easily moved between different territories in a tabletop workspace[18].

However, storage bins only allow for an item to be a member of one group, and therefore cannot belong to multiple groups at the same time. Consequently, when an item is removed from a particular storage bin and put into the common workspace, or another storage bin, it loses its membership with the group it had previously belonged to. Therefore, it might be a difficult task for a user to remember which storage bins an item belongs when an item is moved out of or between multiple storage bins. Time might be wasted having to re-sort items.

2.2 Reaching Items: Interface Currents

In general, tabletop displays have fairly large surfaces. The large size of tabletop displays often make it difficult for users to be able to interact with all items within a workspace simply because of not being able to reach far enough. Not only can this introduce problems for the usefulness of large display environments, it could also pose problems for collaboration on tabletop displays. Furthermore, Scott [17] mentions that even if a user is capable of reaching a specific item in the workspace, they may in turn be reaching over another user's workspace, which ultimately would hinder effective collaboration for the entire group.

Hinrichs et al. [9] have developed a way in which users working in a collaborative workspace can select items without having to physically reach over other users' personal workspaces, while providing better access to all of the items within the workspace. They have developed a tool called an Interface Current which acts as a type of a digital Lazy Susan that, instead of making food available in a kitchen, makes items in the workspace available to all of the users. To have an item circulate within the current it must first be placed within the current by one of the users working within the workspace. The items within an Interface Current flow in a particular path. The path and flow can be modified by the user. For example, the user can specify the Interface Current to be a stream that follows the periphery of the workspace or the user can specify an Interface Current to be more like a pool that appears at the centre of the workspace. Users can also control the speed and direction of the flow of the items within the current by simply swiping inside the borders of the current. The length and speed of the touch will determine how fast items within the flow circulate. The direction of the flow can also be changed by the swiping in the direction opposite to the current's flow.

What is important about this tool is that it provides a way for users to gain access to any item that is circulating within the Interface Current with very little effort. However, if an item is out of reach for a particular user, another user near the item must place the item in the Interface Current in order for it to be accessed by the user who wishes to work with that particular item. Such Interface Currents can also act as means of storage, just like the storage bins proposed by Scott et al. This means that Interface Currents can also be used to group items. However, this also means that items can either be placed in Interface Currents or in the workspace but not in both, which is a limitation of the system.

2.3 Moving and Displaying Items: Mobile Spatial Tools

Due to the typically large size of tabletop displays and the number of people working on a tabletop display, there can be many tasks that are being accomplished at the same time. As a result, tabletop displays may have many items that fill the workspace. Having a large number of items within a workspace may hinder work due to clutter. Therefore, there must be ways of reducing the clutter in a collaborative workspace.

Tools have been developed in order to move items around a workspace in order to clear space for the items that need to be worked on, or to organize objects so that they can be worked on in a more efficient manner.

Isenberg et al. [11] have introduced a set of tools for moving and organizing items within a tabletop workspace. These tools include: a motion tool, a piling tool, an alignment tool, a motion orientation tool, and a radial orientation tool. The first tool that Isenberg et al. [11] propose is a motion tool that allows users to push items around the workspace. This acts as a type of broom that is able to sweep the user's workspace clean so to avoid clutter.

The next tool that Isenberg et al. [11] discuss is a piling tool, which allows users to create piles of items that appear within a common location. Piles have been used extensively to extend the desktop metaphor of physical piles. Piles on a desktop can be used to create loosely organized groups of items to reduce clutter, as piles are space efficient.

A set of orientation and alignment tools have also been developed by Isenberg et al. [11]. An alignment tool has been designed to temporarily align items to a given shape. This would allow the user to see a group of items within a specific threshold to share the same alignment. They have also created a permanent alignment tool or orientation tool that aligns items within the given threshold permanently as the tool is passed over items. Finally, another orientation tool was developed called the radial orientation tool that acts in a similar way to the previously mentioned alignment and orientation tools but instead aligns the items in the workspace radially.

Chapter 3

Design

The goal of this thesis is to design and implement a tool that accomplishes the tasks of storing and grouping, reaching, moving, and displaying data items, such as images, within a tabletop display workspace. The benefit of having a separate tool for each task means that such tools are kept relatively simple and generally perform well in terms of usability. However, providing a user with too many separate tools might confuse users in that keeping track of what tool accomplishes what task is a task in itself. Therefore, it is beneficial to have a single tool that accomplishes a subset of key tasks such as the ones mentioned here. However, caution must be taken against designing a tool with too many features that would be confusing for users. As a result, the design of the tabletop display magnets brings together a limited number of powerful features that can be used to accomplish the tasks mentioned: grouping and storing, reaching, moving, and displaying data items on a tabletop display.

3.1 Item Attraction

The most important feature of the tabletop display magnets is its ability to attract items within the workspace. As shown in Figure 3.1, items are attracted towards the coordinates of the magnet. By being able to attract items towards a set location, users can store items in a common location. This could prove particularly beneficial when a workspace is becoming cluttered and, therefore, hard to work in.

Items accelerate towards the magnet as they become closer. This means that if an item is farther from the magnet it is being attracted to, it will move more slowly than another item that is attracted to the same magnet at a closer distance. Even though real magnets are not commonly used to organize data items in a workspace on a physical table, by having items accelerate towards a magnet as they become closer will increase the realism of the magnets. This is because the distance between an object and a magnet in the physical world is a determining variable for the force

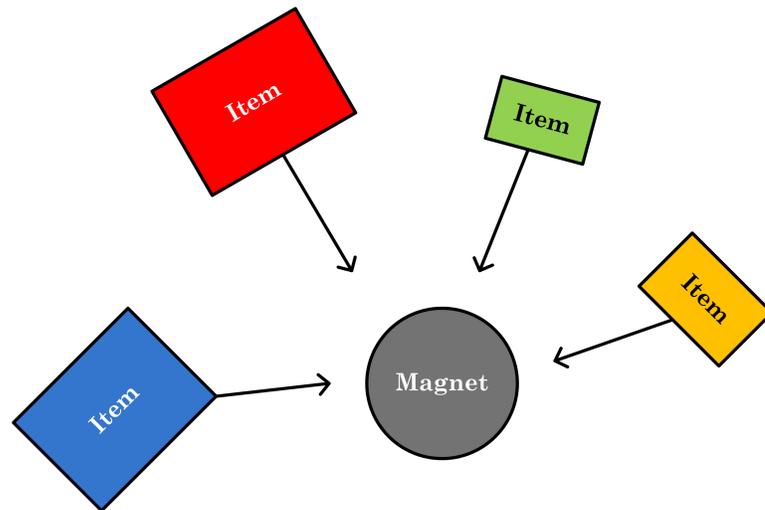


Figure 3.1: Magnet attracting items in the workspace.

of attraction that object has towards the magnet. This may in turn allow users to work more naturally with the magnets.

Magnets are also able to be displaced in the workspace. This means that users are able to change what location items are attracted to. When an item reaches the location of the magnet it will essentially become stuck to that magnet. As a result, when a magnet is moved around the workspace the items that it attracted move with it. Obviously, users will most likely want items to become “unstuck” from the magnet so that they are capable of being moved to whatever location is suitable with respect to the work that is being performed. Therefore, a button is provided to the user so that they can break the so-called magnetic force between the item and the magnet.

The attracting component of the tabletop display magnet can in itself be used to perform some of the tasks that multiple other tools have been developed to accomplish. Firstly, like the Interface Current tool that was described in Section 2.2, these magnets give users the ability to gain control of items that are not necessarily within their reach. This can be achieved by turning the magnet on so that items in the workspace are attracted towards the magnet. This means that if the magnet is in the reach of the user, then the other items in the workspace are also in reach for the user. The tool can also be used as something akin to the piling tool described in Section 2.3. Turning magnets on and allowing items to be attracted to the coordinates of those magnets, allows for the space, previously occupied by the items, to become free.

Consequently, users would be able to reduce clutter by bringing items to a common point.

3.2 Item Sorting & Grouping

One major limitation with the design of the tabletop display magnet, as described thus far, is that there is only one single magnet existing in the workspace. As mentioned at the end of Section 3.1, users only have the ability of accessing or reaching items in the workspace if the magnet is accessible. That is, if the magnet is not within the reach of the users then the users would not be capable of taking advantage of the capability of magnets being able to reach for items in the workspace. If there is only one magnet existing in the workspace, then there is a greater chance that it will be out of reach of some of the users working in the workspace. To overcome this problem, the user is able to dynamically create new magnets that are capable of existing and working together in the same workspace. This would allow each user to have a magnet that attracts the items in the workspace, in turn giving them the capability of reaching items that are not necessarily in their close proximity.

There is more potential for the use of multiple magnets than simply providing users with the ability to reach all items within the workspace. Some users may not wish to have every item attracted to the same magnet. Therefore, it could be beneficial to allow magnets to attract different subsets of items. For example, suppose a group of graphic artists working on a tabletop display wish to have some colours (like cool colours such as blue and green) be attracted to one magnet, and have warm colours (such as red, orange, and yellow) be attracted to another magnet. This is shown graphically in Figure 3.2.

Essentially, this allows for magnets to group items in the workspace together just as the Storage Bins discussed in Section 2.1 were able to group items together. Furthermore, to reiterate, when the items that are attracted to a particular magnet reach the magnet they are attracted to, they become stuck and move with the magnet when the magnet is displaced. Just like Storage Bins, magnets allow for groups of items to be moved throughout the workspace, which means that these groups can easily be moved between user territories and a common territory.

Just as users may want to make their own groups of items using Storage Bins

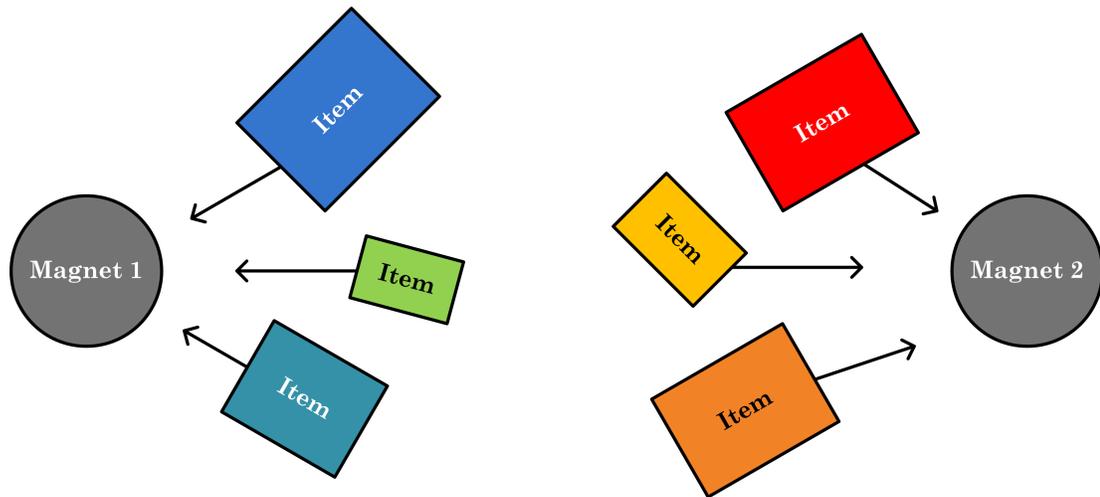


Figure 3.2: Multiple magnets attracting different subsets of items.

introduced by Scott et al. [18], users using Tabletop Display Magnets may also want to make their own groups of items, without the constraint of the groups being predetermined. In order to allow users to create and modify their own groups of items, the ability to make an object attracted to a magnet is given to the user. To accomplish this, a threshold is drawn around each magnet. Moving an item within a threshold creates a link that will allow that item to be attracted to the given magnet once the magnet's magnetic pull is turned on. Figure 3.3 visualizes this process. This method was chosen to provide users with a familiar drag and drop functionality that is used in many personal computing applications. Even the Storage Bins that Scott et al. [18] discuss use a drag and drop functionality to associate items with their parent bin. Of course, with these magnets, bringing an item outside of the threshold will not break the link that an item has with a magnet. This allows an item to become attracted to multiple magnets. This may prove to be useful for users, especially in applications where multiple attributes of items are being considered.

Once users have created groups for the items within the workspace, it may not be completely apparent to the users which items belong to which groups. With storage bins it is obvious to see which items belong to which storage bin since the items that belong to the storage bin appear directly within it. This is not the case with magnets, since an item that is attracted to a magnet does not show that it is attracted to such a magnet, and consequently, a part of a group. The only means of showing that an

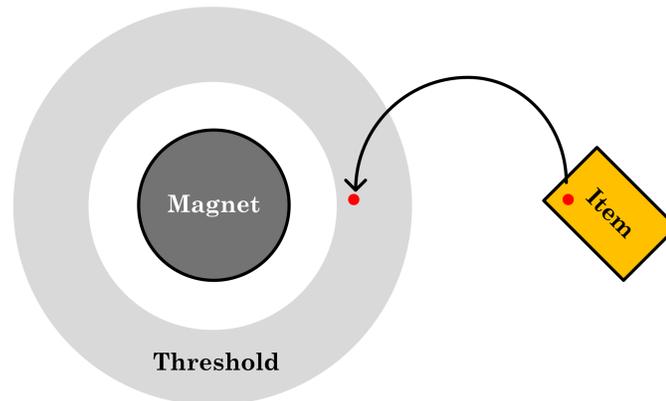


Figure 3.3: Magnet with threshold. Process of attracting an item to a specific magnet.

item is attracted to a particular magnet is by turning the magnet on and observing whether or not that particular item is being drawn towards the centre of the magnet. To overcome this problem, the ability to draw edges between magnets and the items that they are attracted to, is given to the user. What results are graphs that show the memberships items have in relation to the magnets in the workspace, where magnets and items are vertices.

For example, in Figure 3.4 (a), without prior knowledge, it is impossible to determine which items are attracted to which magnet. However, in Figure 3.4 (b), it is quite clear that items that are of primary colours (red, blue, and yellow) are attracted to Magnet 1, and items that are of secondary colours (green, purple, and orange) are attracted to Magnet 2. Consequently, two separate groups emerge. Of course, an item can be an endpoint for multiple edges from different magnets since, as previously mentioned, an item can be attracted to multiple magnets.

It may become necessary over time for users to modify groups by not only adding items but also removing items from groups. Referring back to the Storage Bins discussed by Scott et al. [18], in order to remove an item from a storage bin, users must drag an item outside of the storage bin. This of course is quite intuitive for storage bins. However, to recall, an item's membership to a group is not broken when it moved outside of a magnet's threshold. Therefore, the drag and drop method of breaking the association between a magnet and an item is not appropriate in this context. Thus, a different method for breaking the link between an item and a magnet must be established. This can be accomplished by augmenting the edge

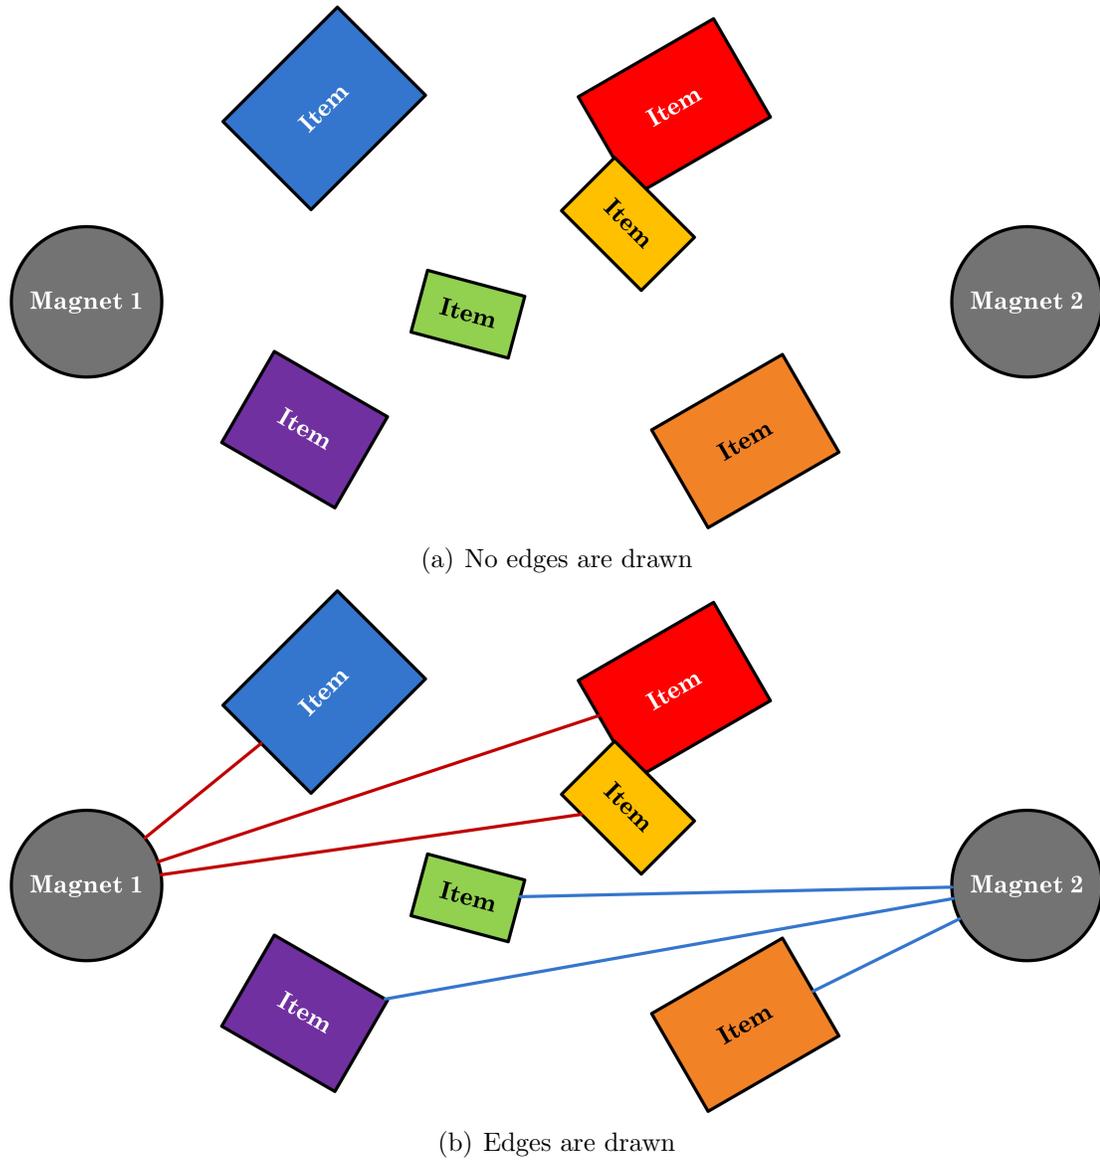


Figure 3.4: (a) Multiple magnets and items without edges drawn showing group memberships; (b) Multiple magnets and items with edges drawn to show group memberships.

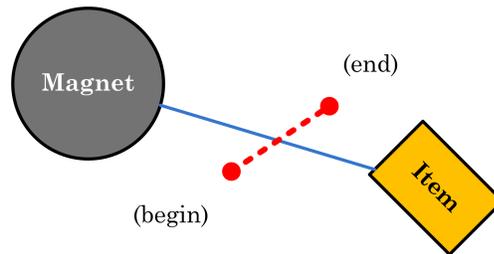


Figure 3.5: Process for breaking the attraction between a magnet and an item. The blue edge represents the attraction between Magnet 1 and the item. The red interrupted line shows the path that a user would draw to remove the attraction between the magnet and the item.

drawing feature that has been discussed. In order to break the association between a magnet and an item, the user must first ensure that the edges are shown between the magnet and all of the items it attracts. Once the edges are shown, users can strike over the edges that they wish removed. Figure 3.5 shows how the attraction between a magnet and an item can be broken.

Given the features that have been discussed in this section, users will be capable of building groups that allow items to be sorted and stored just as they can with Storage Bins discussed by Scott et al. [18]. However, what is beneficial about using magnets is that it allows items to not only belong to multiple groups but also be moved anywhere within the workspace and maintain memberships with any of the groups. Furthermore, as it has been discussed, items can be removed from a group by simply breaking the links, or edges, between magnets and items.

3.3 Moving Items with Magnets

In Information Visualization, the motion of objects in a perceived workspace is very important. Firstly, it has been found that motion can play a fundamental role in attracting the attention of users [20]. In the context of tabletop display magnets, users may wish to attract a group of items that have been defined by a magnet. By turning the magnet on, it would set the items, that are attracted to it, in motion. This would attract the attention of the users towards those items that are being moved. Also, through many studies it has been found that people are very sensitive to motion and movement within a perceived workspace. Furthermore, people are able to identify

emerging patterns in data that is set in motion. In fact people are able to identify patterns in motion equally well as, or better than static patterns that employ colour, size, rotation, and other static patterns as described by Gestalt Laws [20]. This is particularly helpful information with respect to Tabletop Display Magnets. When users turn on multiple magnets, the magnetic force that the magnets exert on the respective items in the workspace cause those items to be set in motion towards the magnets that they are attracted to. For example, if two magnets are turned on and are located in very different areas of the workspace, the items will be moved towards two separate, very different locations. As the items move towards these locations the users could potentially identify which items belong to which magnets simply by what direction the items are moving in, and thus mentally group the items before the items have even reached the magnet, particularly if only few magnets are turned on. Therefore, it can be seen that motion plays an important role for this tool. Consequently, this tool should allow the user to control the motion that magnets induce in the workspace.

It may become necessary for the user to move items in the workspace at different paces. As it was mentioned in Section 3.1 the magnetic force that pulls an item towards a magnet is greater as the distance between the item and the magnet decreases. Therefore, the acceleration of an item will increase as an item becomes closer to a magnet. Therefore, an item that is far from the magnet will begin accelerating at a slower rate than an item that is much closer. Suppose a user wishes to access items in the workspace that are at the opposite end of the tabletop display. The user can do so by turning on a corresponding magnet that will attract the desired items towards the user. However, since the items that the user wishes to access are at the opposite end of the table, the items may move very slowly towards the magnet. Conversely, a user may wish to gain access of items that are close to the magnet by turning on a magnet. However, since the items that the user wishes to access are close to the magnet, the items may move too fast for the user's liking. These two situations pose a problem for the user. There must be a way for the user to control the magnetic force that pulls items towards magnets.

To accomplish this, a dial is provided for each magnet that will allow the user to increase or decrease the magnetic force that pulls items towards it. This is shown in

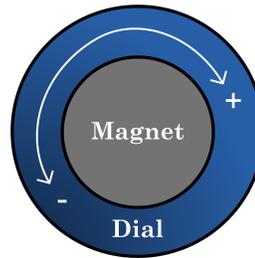


Figure 3.6: Magnet showing the dial for increasing and decreasing the magnetic pull that a given magnet has on items in the workspace. Turning the dial counter-clockwise decreases the magnetic force, while turning the dial clockwise increases the magnetic force.

Figure 3.6. The user can turn the dial on a magnet by moving their finger around the designated part of the magnet. Moving their finger counter-clockwise lessens the magnetic force that a magnet exerts on the items that it attracts. Conversely, moving their finger clockwise increases the the magnetic force.

Users may also want to know how high or how low the magnetic force is for a particular magnet. Therefore, there must be a way to visualize the magnetic force a magnet exerts on its associated items. To do this a cloud or halo appears around the magnet. For example, Figure 3.7 shows this. As it can be seen the halo around the magnet in Figure 3.7 (a) is larger than the halo around the magnet in Figure 3.7 (b). As a result, the force acting on the item in Figure 3.7 (a) is greater than the force acting on the item in Figure 3.7 (b).

Just as users may want to attract items within a workspace, they may also want to push items around. Just like the motion tool discussed in Section 2.3, tabletop display magnets can be used to push items around the workspace. Being able to push items around the workspace allows the user to move unneeded items away from the space that the user is working in on the tabletop display, reducing clutter. At first this may seem counter intuitive to the idea of a magnet, as magnets are usually associated with attracting items as opposed to repelling them. However, this can be considered a fairly natural extension to the design of these magnets. So far, these magnets have the capability of decreasing and increasing the magnetic force of attraction between magnets and items. Instead of decreasing the magnetic force of attraction until the magnetic force reaches zero, the magnetic force can become negative. As a result, the negative magnetic force will push items from the magnet as oppose to attract items

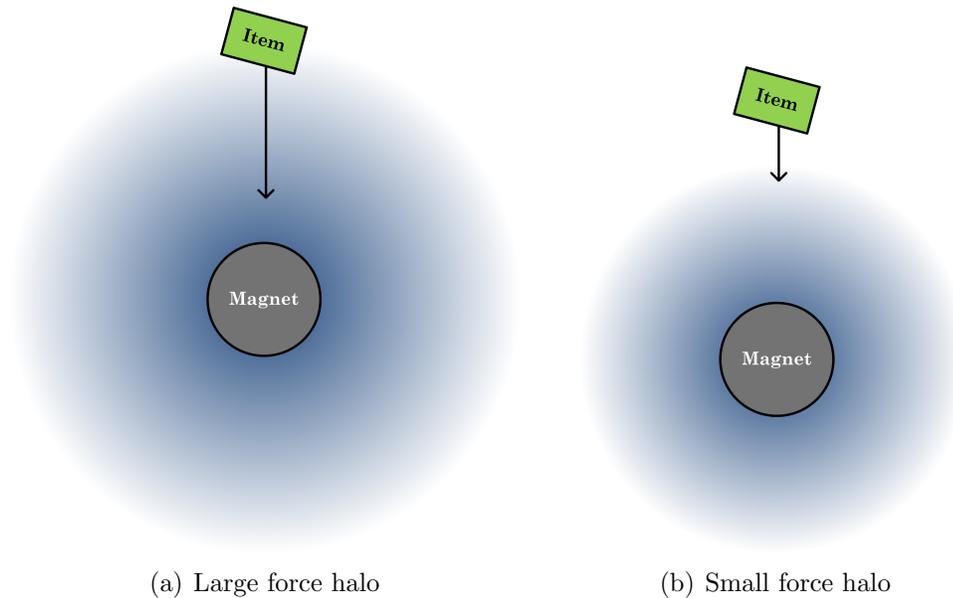


Figure 3.7: (a) Magnet showing a large force halo pulls items at with greater force; (b) Magnet with smaller force halo pulls items with a lesser force.

to the magnet. Instead of being visualized by a blue halo, a negative magnetic force can be visualized by a red halo.

For example, Figure 3.8 shows a magnet with positive magnetic force (a) and a magnet with negative magnetic force (b). Both magnets have the same magnitude and therefore exert the same amount of force on the item that each act upon that happen to be separated by the same distance. However, the signs of the forces are different. Therefore, the positive force pulls the item towards the magnet, whereas the negative force pushes the item away. It should be noted that the item, regardless of whether it is being attracted or repelled, moves along the same line that passes through the centre of the magnet and the item. Items do not return to the positions they were in before being acted upon by tabletop magnets when the magnetic pull is negative.

3.4 Displaying & Aligning Items

Up to this point, not much has been discussed about how items are displayed once they reach a magnet. What has been discussed is that magnets can be used to create piles of items just as the piling tool that Isenberg et al. [11] created to reduce clutter

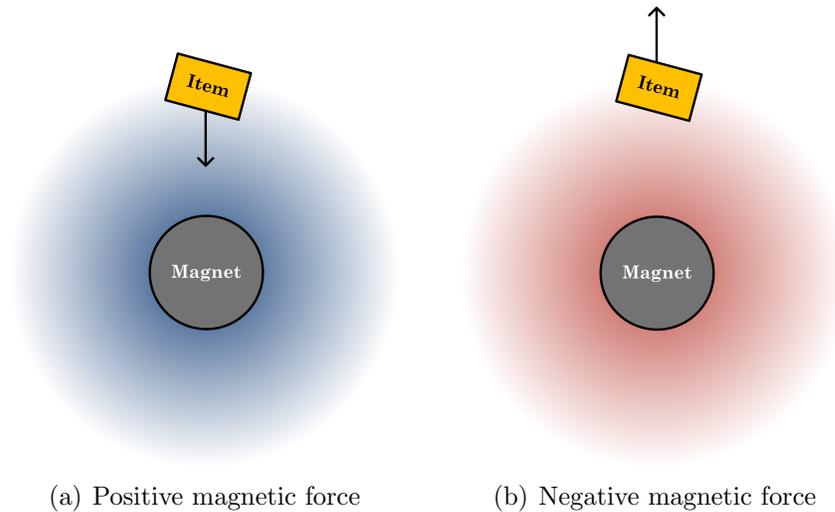
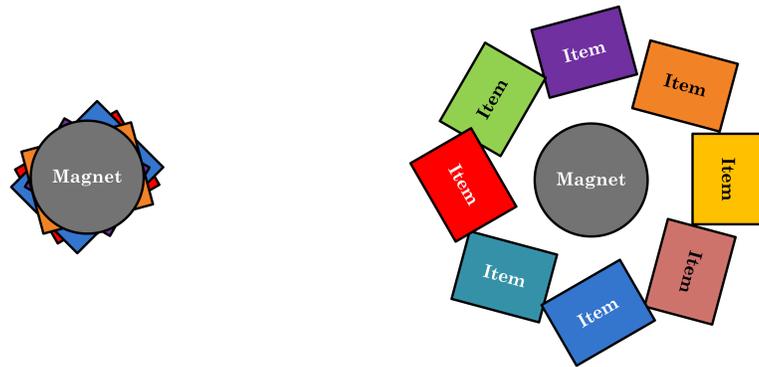


Figure 3.8: (a) Magnet showing a blue halo representing a positive magnetic force; (b) Magnet showing a red halo with the same magnitude but representing a negative magnetic force.

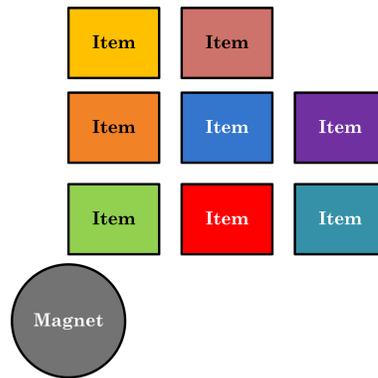
by concentrating a group of items at a common coordinate as defined by the location of a magnet. This may be the best way for items to be aligned in order to conserve space, but it does not allow users to view the items in the group once they have reached the magnet. This is because, not only would the magnet be overlapping the items, but other items may overlap the other items in the group.

As it can be seen in Figure 3.9 (a), items are tightly piled on top of one another, conserving space. However, it is not entirely clear how many items are in this pile, or more importantly, what items are in this pile. Therefore, the option to change how the items are presented is given to the user. In Figure 3.9 (b), the items that are attracted to the magnet are now distributed around the magnet in a radial alignment. Having items displayed radially allows for the items belonging to the group to be better recognized. Furthermore, it is easier to count the number of items that appear in the group. This is because the items in the group do not overlap as much as they would in a pile. However, the downside of displaying items in a radial alignment is that it consumes more space than a pile. Also, the radial alignment is also limited in the number of items that can appear in the group. The more items that are in the group, the more overlap that will exist. As shown in a pile, overlap can make identifying the items within a group a difficult task. Therefore, an additional means of aligning items belonging to a group in a grid is provided to the user. By arranging



(a) Pile

(b) Radial Alignment



(c) Grid Alignment

Figure 3.9: (a) Magnet organizing a group of items into a pile; (b) Magnet aligning a group of items radially; (c) Magnet organizing a group of items into a grid.

items in a grid, there is guarantee that items will not overlap as they would in a pile or in a radial alignment. This is shown in Figure 3.9 (c). This means that the items that are attracted to the magnet are equally viewable and nothing is obstructing the users' views of any of the items in the group. However, there is also a negative aspect with the grid alignment, since it takes up considerably more space than a radial alignment. This is especially the case if a magnet attracts many items. Since each of these alignments have their own negative and positive aspects, the ability to toggle between these alignment options is given to users so that they can use their discretion to decide which alignment is most appropriate for the work that is being done. It should also be noted that with the radial and grid alignments, the items are scaled to have roughly the same dimensions to avoid lopsided viewing of items avoiding additional misuse of workspace and clutter. When the magnetic pull is turned off, the items return to their original sizes that were attributed to them before being scaled by the magnet alignment.

3.5 Magnet Layout

Up to this point, the magnet has been depicted in a fairly general way. In the provided figures, the magnets have been shown as grey circles with very few features. However, many features have been attributed to the design of the magnet so far. The magnet is first of all able to be displaced around the workspace. Secondly, there is the capability of turning magnets on and off. Moreover, there is a threshold that is used to drag and drop items into a magnet for items to be attracted to the magnet. There is a dial that allows users to control the magnetic force that is exerted on items in the workspace. There are also halos that surround each magnet to quantify their magnetic pulls. The user can toggle edges on and off between magnets and the items that they attract (or repel). Finally, users are also able to switch between various ways of aligning items once they reach destination magnets.

The reader may be a bit reluctant to believe that the proposed magnet can accomplish these tasks given the way the magnets have been portrayed in the provided figures. Figure 3.10 shows a better approximation of what a tabletop display magnet looks like. To displace the magnet the user would select the middle part of the magnet (labeled "Magnet"). As mentioned previously in this chapter, to add an

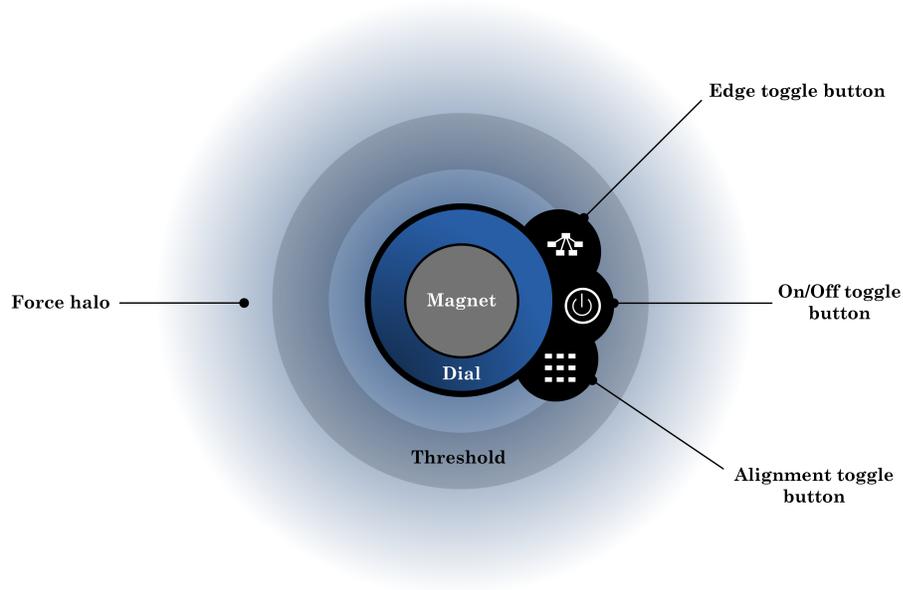


Figure 3.10: A more realistic rendering of what the design of the tabletop display magnet looks like. This shows the dial, threshold, force halo, and toggle buttons attributed to these magnets as discussed in the design.

item to the magnet, users must drag and drop items past the magnet’s threshold. The user can increase and decrease the force field by rotating the dial clockwise and counter-clockwise, respectively.

What has not been discussed have been the three buttons that are placed on the right side of the magnet. These buttons have been grafted to the side of the magnet since the centre of the magnet is already being used for magnet displacement. The first button is used to toggle the edges that are drawn between the magnet and the items that it attracts and repels. The next button is used to toggle the magnet on and off. Toggling the button on lets the magnet exert force on the items that are associated with it. Toggling the button off stops the exertion of force on the associated items. Finally, the last button is used to toggle between the alignments mentioned in Section 3.4: piles, radial alignment, and grid alignment.

Now that the design of the Tabletop Display Magnet has been discussed, an overview of the implementation can now be given. The next chapter discusses some of the key components and notions needed to implement this tool, along with information about how the features that belong to the design of this tool are implemented.

Chapter 4

Implementation

4.1 Tabletop Displays

When researchers begin researching topics relating to tabletop displays, they usually build or find a tabletop display that suits their specific needs. As a result there are many varieties of tabletop displays that are available today, each with their own capabilities and types of interactions. To begin, there are three different categories in which tabletop displays can be separated with respect to how they display their images [14]:

- An **overhead-projected display** has its image projected on its surfaces from one or more projectors that are positioned above its surface.
- A **rear-projected display** has its image projected by one or more projectors beneath its surface.
- A **self-illuminating display** has its image generated by itself without the image being projected onto the surface.

Along with how an image is displayed on a tabletop's surfaces, there are many methods that have been employed to interact with a tabletop display's surface.

- **Mice** can be used on tabletop displays just as they can on desktop computers. In this case the tabletop display would simply act as a large monitor that would be used with a personal computer.
- Tabletop display **pens and styli** have also provided familiar tools for interacting with a tabletop display, similar to markers on whiteboards. [12]
- Tabletop displays can also accept **direct touch** with the tabletop surface from fingertips or other hand positions.



Figure 4.1: Mitsubishi Engineering Research Laboratory DiamondTouch display features signal receivers for each user which are capable of detecting inputs from particular users.

Some tabletops also support different numbers of simultaneous inputs and users. Some tables can only support one input at a time, whether it be from a mouse, stylus, or direct touch. Other tables are capable of supporting two or more inputs simultaneously, which would in turn facilitate multi-user work. Some tabletop displays are even capable of detecting which user an input comes from.

These days, different tabletop displays have been developed by various researchers and companies such as the Mitsubishi Engineering Research Laboratory (MERL), Microsoft, and SMART Technologies.

To begin, the *MERL DiamondTouch* is an overhead-projected display that is capable of accepting multiple touches from multiple users [13]. It is also among the few tabletop displays that are able to detect inputs from specific users. This is done by the use of a receiver which receives signals, that are sent through a user's body, from the tabletop upon touching the surface of the display. It is also capable of detecting various hand gestures such as fists and open palms[19]. DiamondTouch displays come in two different sizes: 32-inches diagonal table and 42-inches diagonal table .

Last year, Microsoft introduced its tabletop display called *Microsoft Surface* that it intends to make commercially available. It features a multi-touch display that can recognize "fingers, hands, gestures and objects placed on the surface". The display is 30-inches in diagonal, smaller than the MERL DiamondTouch [5].



Figure 4.2: Microsoft Surface features a multi-touch display that can recognize a variety of inputs from fingers to objects such as cell phones, digital cameras, and credit cards.

Finally, SMART Technologies is a producer of interactive whiteboards, although various prototypes of tabletop displays have been developed. Some of their tabletop displays are simply their whiteboards that are positioned horizontally rather than up against a wall or on a stand. However, SMART Technologies have developed a rear-projection tabletop display prototype that supports a $2,800 \times 2,100$ resolution. It is roughly 77-inches in diagonal. It can also support up to two simultaneous inputs from either a mouse, or finger touch, or styli [14].

Tabletop display magnets have been implemented for the 77-inch SMART Technologies prototype tabletop display. Although, as it has been mentioned, there are various types of tabletop displays in existence, the research that has been performed on tabletop display magnets does not limit the use of these magnets to the tabletop prototype developed by SMART Technologies. The problems that tabletop display magnets solve are common among all tabletop displays that facilitate collaboration. As a result the ideas discussed in this section for implementing tabletop magnets could be applied to other tabletop displays and environments.



Figure 4.3: SMART Technologies prototype tabletop display at Dalhousie University has a 77-inch display which supports up to two simultaneous inputs.

4.2 University of Calgary Large Display Framework

The Innovation in Visualization Laboratory at the University of Calgary, or InnoVis, have conducted much research concerning large interactive displays such as interactive walls, interactive whiteboards, and tabletop displays. As a result they have built numerous applications for large displays that require responsive interactions. Providing responsive interactions on large displays is a difficult task to accomplish. Since tabletop displays are generally quite large, they possess enough surface to potentially display many high-resolution objects at the same time, many more than a typical computer monitor. Furthermore, to maintain the table workspace metaphor, a tabletop display should provide immediate responses to any interaction users have with the tabletop display. For example, if a user wishes to move an object from one end of the tabletop display to another, the tabletop display should respond immediately by moving the object across the table as the user desires. Moreover, because of their collaborative nature, tabletop displays should be able to support responsive interactions from multiple users [10].

To overcome limitations for developing software for large displays, InnoVis have developed a framework that provides developers with a “foundation to build applications for tabletop displays” [10]. Above all, the Large Display Framework (LDF), which was recently publicly released in late-2007, seeks to allow responsive interactions for large display applications. The framework they have developed also acts as

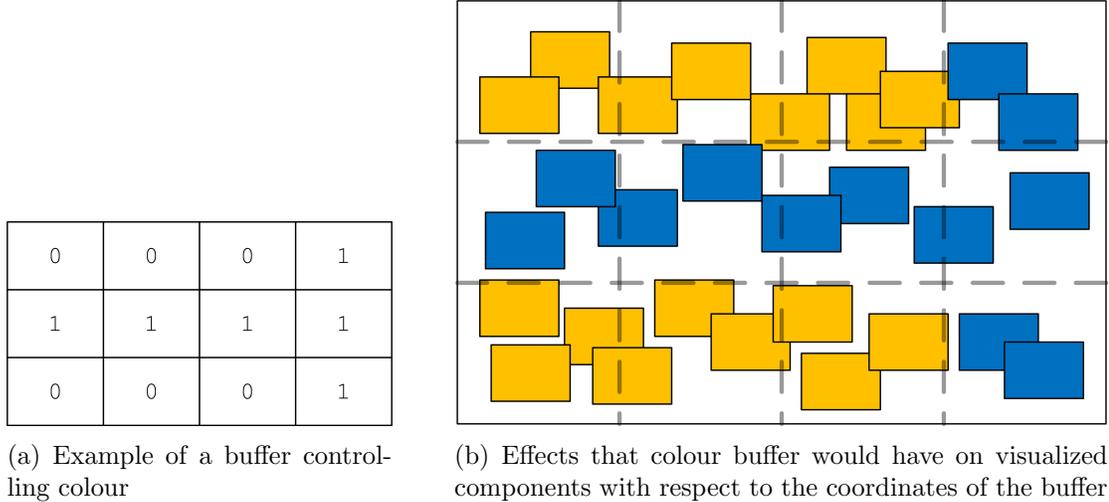


Figure 4.4: Example of how an i-buffer is used to effect the display of visualized components. (a) Shows the values of the buffer that could control the colour of visualized components; (b) Shows the effects of such a buffer (a ‘0’ value turns a rectangle yellow; a ‘1’ value turns a rectangle blue) [14].

a kind of base application, which developers can use to implement their own tabletop applications. The LDF played a fundamental role in the implementation of the tabletop display magnet tool.

The LDF uses the idea of interaction buffers, or *i-buffers*, and the composition of *visual components*. To begin, an i-buffer is a data structure that is used to control the visual appearance of a visual component. For example, i-buffers are used to set the colour, scale, angle of rotation, and displacement of a visual component. In this case a visual component is anything that is visualized on the tabletop display. For example, a visual component could be a square, circle, or even an image, drawn to the display. Figure 4.4 shows an example InnoVis uses to demonstrate how i-buffers are used to set attributes of visual components [2].

There are various types of visual components. For example, a *container* is a type of visual component that contains other visual components. A Storage Bin as described by Scott et al. could be an example of a container, since a Storage Bin is capable of containing other visual components, like images. The relationship of which visual component belongs to which container visual component is described by means of a *composition tree*. At the root of a composition tree is the background component of the workspace, since each component in the workspace appears within

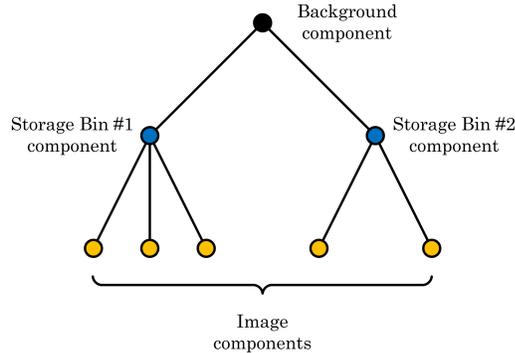


Figure 4.5: *Composition Tree* showing the background component at the root and storage bins that contain two and three children images

the background component’s space. Furthermore, there could be a storage bin component that contains three images, and another containing two images. The storage bins would be represented as parents of the images, and the background would be the parent of both storage bins, in the composition tree.

Strategies are classes that describe the behaviour of a visual component. A strategy could be developed to describe the shape of a visual component or how a visual component moves or is interacted with. For example, a strategy could be developed to describe how a visual component can be rotated. Developers can give a visual component many strategies, and essentially create a group of behaviours for a visual component. Typically, a visual component would have a strategy attributed to it for describing what is displayed to the user, and another strategy allowing a visual component to be moved around the surface of the tabletop display by the user. These strategies interact with the previously mentioned i-buffers. They also relate to event handling. Visual component strategies can be reused between different visual components.

The LDF does not only provide the means to create these strategies for creating and controlling visual components that appear in large display applications, but it also provides a toolkit that contains strategies that have already been built to accomplish tasks that are fairly common across most large display applications. One such strategy provides the capabilities as described by “*Rotate n’ Translate*” [7], which handles the rotation, translation, and “tossability” of visual components that appear in a large display application. Other strategies provide the ability to load and display images, to

resize objects, and to apply borders to objects. As a result, fairly involved applications can be assembled quite easily with the use of the LDF.

For the sake of building a simple environment containing multiple images that can be rotated, translated, and tossed around a tabletop workspace, these strategies that are provided with the LDF have been employed. Given a simple environment with multiple images within it, only the tabletop display magnets as described in Chapter 3 need to be implemented. To do this, three new strategies are created: a magnet strategy, an attracted component strategy, and an attraction manager strategy. Section 4.3 will discuss the implementation of these strategies.

Before, the implementation of those strategies are discussed it should be noted that the LDF has been tailored to work with OpenGL, a well established two-dimensional and three-dimensional application programming interface (API), GLUT, a utility library used with OpenGL, Qt, a graphical user interface (GUI) toolkit, and SMART Board SDK, a software development kit used for large display interactions on large displays manufactured specifically by SMART Technologies. It has also been tailored to be worked on using the Microsoft Visual Studio 2005 integrated development environment (IDE) with visual C++.

4.3 Magnet Strategy, Attracted Component, & Attraction Manager

As it has been mentioned, three new strategies have been created to implement the behaviours of the tabletop display magnet as described in Chapter 3.

The **Magnet** strategy is complex in that instead of dispersing its functionality over several strategies, it accomplishes all of the behaviours that are specific to magnets. Therefore, the magnet strategy is responsible for, first of all, displaying a magnet, as it is shown in Figure 3.10. It is also responsible for handling interactions from users by determining what part of the magnet is being interacted with and how it should respond to those received interactions. The magnet strategy is also involved in applying its forces to the items that it is attracting or repelling.

To designate an item as being able to be attracted to magnets, a second strategy was created called the **Attracted Component** strategy. The attracted component strategy is a simple strategy that was created to allow visual components to keep track of what magnets they are attracted to, and which magnet the visual component is

stuck to, if any. It also keeps track of the state of attraction. That is, whether the item is at an equilibrium and is stuck to a particular magnet, or the item is not at an equilibrium and the forces acting on the item are setting it into motion. The most fundamental task of the attracted component strategy, is to update the coordinates of the visual component given the sum of the forces that are acting on the item from the magnets.

Lastly, an **Attraction Manager** strategy was created to orchestrate the interactions between magnets and attracted components. It does this by keeping track of all the attracted components that exist in the environment. By keeping track of the attracted components, it is easy to allow multiple magnets to attract the same object. Conversely, it is easy to allow an item to be attracted by multiple magnets, since the attraction manager strategy controls when an attracted component updates the visual component’s coordinates.

Having introduced the three main strategies that have been created in order to implement the functionality as described in the design outlined in Chapter 3 of tabletop display magnets, the remainder of this section will iterate through the design of the system once more explaining how each aspect of the design is implemented.

4.3.1 Item Attraction & Repulsion

The attraction of items within a workspace towards the centre coordinates of a magnet is arguably the most important feature of the tabletop display magnet. In order for a magnet to attract or repel an item, it must “know” which items it attracts or repels. Therefore, the Magnet strategy keeps track of which items it attracts by using an array of Attracted Components. For each of the items that a magnet attracts it must “tell” each item how much to move in both the x and y directions in the workspace.

For now, suppose that magnets only attract items at constant rates. Therefore, what must first be computed is a unit vector in the direction towards the magnet’s centre, $\mathbf{p} = (p_1, p_2)$ from the item’s current location, $\mathbf{q} = (q_1, q_2)$. To do this we must find the angle of this unit vector with respect to the x -axis of the workspace. Using simple trigonometry this angle, θ can be computed as follows:

$$\theta = \arctan\left(\frac{q_2 - p_2}{q_1 - p_1}\right)$$

Once this angle is computed, we can find the x and y components of the unit vector by finding the cosine and the sine of the angle, θ . Therefore, assuming that the item updates at a constant rate of λ the following vector \mathbf{v} is passed on to the (attracted component) being attracted to the magnet, describing how much to update the coordinates of the item:

$$\mathbf{v} = \begin{bmatrix} \lambda \cdot \cos \theta \\ \lambda \cdot \sin \theta \end{bmatrix}$$

However, the design stated in section 3.1 that items are not to be attracted to magnets at constant rates, but should accelerate towards the centres of magnets in a natural fashion. To accomplish this the equation used to describe the gravitational pull between two objects was employed. Therefore, instead of these magnets acting as actual magnets, they in fact act more like two-dimensional planets (masses) that attract objects of some lower mass. The gravitational force, F between two objects with masses m_1 and m_2 that are separated by a distance r is described as follows:

$$F = G \frac{m_1 m_2}{r^2}$$

G is the universal gravitation constant [8]. However, we will ignore this constant since we are not entirely interested in representing the the exact gravitational force one object would impose on another. Therefore, we have the following proportionality:

$$F \sim \frac{m_1 m_2}{r^2}$$

Furthermore, we are not interested in knowing the force at which a magnet pulls an object towards it. Instead, we are interested in the acceleration of the item towards the magnet. The force F applied to an item with mass m and with acceleration a is defined by $F = ma$. Therefore, suppose that m_1 is the mass of an item being pulled towards the magnet at an acceleration a_1 , then:

$$m_1 a_1 \sim \frac{m_1 m_2}{r^2}$$

Therefore, dividing both sides of this proportionality by m_1 , the acceleration a_1

of the item towards the centre of the magnet, is described by the following proportionality:

$$a_1 \sim \frac{m_2}{r^2}$$

To simplify the task of finding the acceleration of the item we convert this proportionality into an equality giving:

$$a_1 = \frac{m_2}{r^2}$$

The mass of the magnet, m_2 , is simply arbitrarily chosen keeping in mind that the larger the mass of the magnet, the greater the item attracted to that magnet accelerate towards it.

The value of r is found simply by finding the Euclidean distance between the magnet, with coordinates $\mathbf{p} = (p_1, p_2)$ and the item being attracted, with coordinates $\mathbf{q} = (q_1, q_2)$. This value can be computed as follows:

$$r = \sqrt{(p_1 - q_1)^2 + (p_2 - q_2)^2}$$

Therefore, we now have a value that describes the acceleration, a_1 of the item towards the magnet. We therefore replace the previously mentioned eigenvalue λ , which represents a constant velocity and replace it with $a = a_1$. Since we are only concerned with this acceleration a for one single time period, it acts as a velocity. Therefore, we now pass the following vector, \mathbf{v} , to the item to describe how much the item's coordinates should be updated:

$$\mathbf{v} = \begin{bmatrix} a \cdot \cos \theta \\ a \cdot \sin \theta \end{bmatrix}$$

4.3.2 Item Sorting & Grouping

Section 3.2 discusses how multiple tabletop display magnets can co-exist within a workspace, but also how these should allow users to sort and group items. This section will describe how the sorting mechanisms were implemented. That is, this

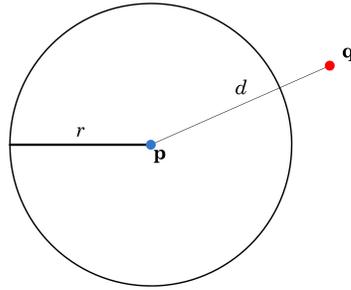


Figure 4.6: Representation of a magnet’s threshold boundary. The threshold defines a circle with a given radius, r . By finding the Euclidean distance, d , between the centre of the magnet, \mathbf{p} and the point at which the user released the item, \mathbf{q} , and it can be determined whether or not an item has been dropped past the magnet’s threshold.

section will describe how the system was implemented to allow items to be dragged and dropped past a magnet’s threshold to make an item attracted to a magnet, and how it can be released from a magnet by striking the edge that exists between a magnet and an item.

To begin, it should be reiterated that a magnet will attract those items that it “knows” it should attract. Those items that it “knows” it should attract are the items that are referenced by an array that each magnet has. Therefore, if the array of a magnet contains a reference to a certain item, then that magnet will attract that item. Conversely, if the array of a magnet does not contain a reference to that particular item, then that magnet will not attract that item. Therefore, these arrays that reference the attracted components, essentially form the groups that sort the items in the workspace.

In order to make an item attracted to a magnet, a user must drag and release the item passed the magnet’s threshold. When this is done, a reference to that item is added to the magnet’s array. To accomplish this, when a user selects an item, drags the item around and then releases the item somewhere in the workspace, each magnet is polled to see whether or not that item was released within the magnet’s threshold. To determine whether or not an item has been dropped past the threshold of a magnet, a simple algorithm is used.

To determine whether or not an item has been dropped past the magnet’s threshold we must first compute the Euclidean distance, d , between the centre of the magnet, $\mathbf{p} = (p_1, p_2)$, and the point at which the user released the item, $\mathbf{q} = (q_1, q_2)$. This

distance is calculated as follows:

$$d = \sqrt{(p_1 - q_1)^2 + (p_2 - q_2)^2}.$$

Once the distance between these two points is calculated, it can be determined whether or not the item was drop past the magnet's threshold. If the distance is less than the threshold's radius, r , then the item was dropped past the magnet's threshold. Otherwise, the item was drop outside of the magnet's threshold. This is shown in Figure 4.6.

Using the foregoing algorithm, we are able to determine whether or not a release point is within a magnet's threshold.

The next issue that must be addressed is how to remove an item from a group by removing the attraction a magnet has on that item. As mentioned in Section 3.2, removing this attraction is done by striking the edge that shows the connection between a magnet and an item. Firstly, the edges between magnets and items must be drawn. This is easily achieved using `GL_LINES` in OpenGL, using the coordinates of the magnet and the coordinates of the item as the vertices that are used to draw the line. To recall, these lines are used to show the memberships belonging to the items.

When the user drags their finger across the surface without having any item selected, the system will pass two points, \mathbf{i} and \mathbf{j} to every magnet that has their edges turned on. The first point, \mathbf{i} represents the coordinate of the user's finger when it first touches the surface, the second point, \mathbf{j} represents the coordinates of the user's finger when it is released from the surface of the display. For each edge that is drawn by the magnet, it is determined whether or not the user has dragged their finger over the edge. If it determines that the user's finger has been dragged over the edge, the edge is removed, and consequently, the item is removed from the corresponding magnet's attracted component array.

To determine whether the user dragged their finger over the edge, we must first find the rotation angle of the edge given its two endpoints, the coordinates of the magnet, and the coordinates of the item, which we will denote, $\mathbf{p} = (p_1, p_2)$ and $\mathbf{q} = (q_1, q_2)$ respectively, relative to the x -axis. As it has been shown before, this is done as follows:

$$\theta = \arctan\left(\frac{p_2 - q_2}{p_1 - q_1}\right)$$

Using this angle, we rotate the edge by $-\theta$ so to align the edge with the x -axis. We can rotate this edge using the following matrix:

$$\mathbf{M} = \begin{bmatrix} \cos(-\theta) & -\sin(-\theta) \\ \sin(-\theta) & \cos(-\theta) \end{bmatrix}$$

Therefore, we have new coordinates for the endpoints of the edge: $\mathbf{p}' = \mathbf{M}\mathbf{p}$ and $\mathbf{q}' = \mathbf{M}\mathbf{q}$. This transformation is shown in Figure 4.7. We also multiply the coordinates \mathbf{i} and \mathbf{j} by \mathbf{M} to get \mathbf{i}' and \mathbf{j}' respectively. These new coordinates are also shown in Figure 4.7. To determine whether or not the edge was cut, we first determine which side the transformed input coordinates are on. This is done by comparing the y -coordinates of \mathbf{i}' and \mathbf{j}' with the y -coordinate of either \mathbf{p}' or \mathbf{q}' . (The y -coordinates will be the same since the segment is rotated by $-\theta$). If the y -coordinate of \mathbf{i}' or \mathbf{j}' is greater than the y -coordinate of \mathbf{p}' or \mathbf{q}' , then the points are on one side, otherwise, they are on the other side. If each \mathbf{i}' and \mathbf{j}' are on opposite sides then the user potentially dragged their finger through the edge. However, before concluding whether or not this is the case, it must be determined whether the line that passes through \mathbf{i}' and \mathbf{j}' intersects the segment between \mathbf{p}' and \mathbf{q}' . To do this we find a point on the line through \mathbf{i}' and \mathbf{j}' that has the same y -component as \mathbf{p}' and \mathbf{q}' . If the x -coordinate of the point is between the x -coordinates of \mathbf{p}' and \mathbf{q}' , then we know that this point intersects the segment between \mathbf{p}' and \mathbf{q}' . If the two segments intersect, then it can be concluded that the user dragged their finger through an edge causing the edge to be deleted.

The algorithms discussed in this section are implemented under the Magnet strategy. Given these algorithms the tabletop display magnet system is capable of having items being added to groups dynamically, thereby making the item attracted to the groups' corresponding magnet. It is also able to dynamically remove items from groups by having the user pass their finger over the edge representing the link between a magnet and an item.

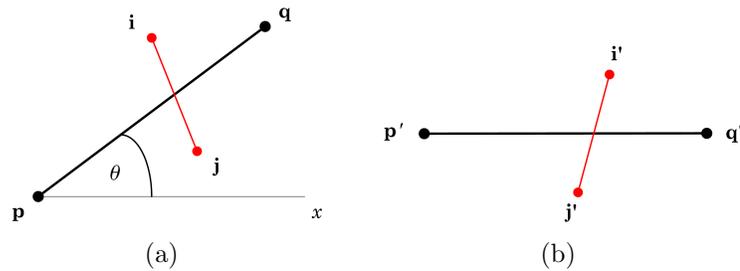


Figure 4.7: (a) Initial alignment of edge endpoints \mathbf{p} and \mathbf{q} , and input points \mathbf{i} and \mathbf{j} ; (b) Alignment of edge endpoints \mathbf{p} and \mathbf{q} , and input points \mathbf{i} and \mathbf{j} after being rotated by $-\theta$ degrees

4.3.3 Moving Items with Magnets

Tabletop display magnets should also be able to control how they move items they attract. Furthermore, magnets should be able to repel items in order to help clear clutter in a particular space. To allow the user to control how magnets attract and repel items, a dial was designed. By having the user turn the dial clockwise, the magnetic pull increases, causing the attracted items to accelerate towards the magnet at a faster rate. Conversely, turning the dial counterclockwise decreases the magnetic pull, causing the items to accelerate towards the magnet at a slower rate, or begin moving away from the magnet at an increasing rate if the magnetic pull is made negative.

Therefore, to accomplish this the dial must be implemented. To implement this dial, a ring is drawn around the magnet to designate the boundaries of the dial. The magnet, must then determine if the dial is being interacted with by determining that the input from the user occurs within the boundaries of the dial. This is done in the same way it is determined that an item was dropped within the threshold of a magnet as described in subsection 4.3.2. As the user moves their finger around the dial, the coordinates of the user's finger are polled. From these coordinates, the angle between these coordinates and the centre of the magnet with respect to the x -axis is determined. This allows the change in angle to be identified. As the change in angle decreases, the magnetic pull increases. Conversely, as the change in angle increases, the magnetic pull decreases. Therefore, the change in angle is negatively correlated to how the magnetic pull is updated.

To visualize the magnetic pull, a halo is drawn around the magnet. The radius

of the halo is proportional to the magnitude of the magnetic pull. However, if the magnetic pull is negative, the halo is drawn in red, and if it is positive it is drawn in blue.

In subsection 4.3.1, it was found that the acceleration, a_1 , of the item being attracted towards a magnet is defined as follows:

$$a_1 = \frac{m_2}{r^2}$$

The variable m_2 is the arbitrarily chosen mass of the magnet, and r^2 is the distance between the item and the magnet. In order for the acceleration of the item to be correlated to the magnetic pull that the dial controls, we simply set the arbitrarily chosen mass of the magnet equal to the magnetic pull. Therefore, decreasing the magnetic pull would decrease the mass of the magnet, which as it can be seen, would decrease the acceleration of the item. If the magnetic pull is negative, then the direction of the acceleration will also be negative, which would cause the vector \mathbf{v} to be flipped. Therefore, regardless of whether or not an item is being attracted or repelled by a magnet, the item moves along the same line that passes through the centre of the magnet and the centre of the item. The item moves away from the magnet, along the line, when the item is being repelled, but towards the magnet, along the line, when the item is being attracted.

Therefore, it has been shown how the dial for updating the magnetic pull has been implemented, and how it effects the magnetic pull of the magnet, which in turn effects how the items interact with the magnet. These concepts have all been implemented in the Magnet strategy.

4.3.4 Displaying & Aligning Items

In Section 3.4 three different methods (radial, pile, and grid) were discussed to align the items that are attracted to a magnet. This section explains how these three alignments were implemented.

At the time of creating a new magnet, three buffers are created to store alignment values for all of the items that could potentially be attracted to the magnet for each of the three alignments. The alignment values are computed relative to the centre of

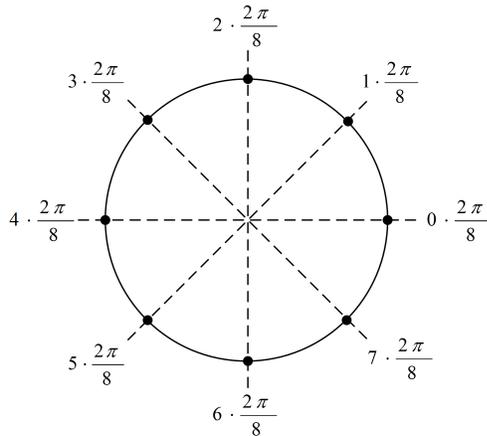


Figure 4.8: The angles between the items displayed in the radial alignment is $\frac{2\pi}{M}$, where M is the maximum number of items that can be attracted to a single magnet. In this case $M = 8$.

the magnet.

The radial alignment coordinates are computed in a similar way as a circle is drawn in two-dimensional computer graphics, since the radial alignment creates a circle of items around the magnet. If there can be a maximum of N items that can be attracted to a magnet, then, to evenly distribute the items in a radial alignment, the items must be spaced $\frac{2\pi}{M}$ radians apart. Therefore, the angle between the first alignment position and the centre of the magnet is 0 radians. The next angle between the second alignment and the centre of the magnet is $\frac{2\pi}{M}$, and so on. Such angles are shown in Figure 4.8. As a result, the coordinates for i^{th} radial alignment point are:

$$\left(r \cdot \cos \left(i \cdot \frac{2\pi}{N} \right), r \cdot \sin \left(i \cdot \frac{2\pi}{N} \right) \right)$$

These points have the centre of the magnet as their origin, and r is the radius of the circle on to which the radial alignment is projected.

Most likely, the maximum number of items attracted to a magnet will not be attracted to a magnet. Instead there is some number $N < M$ of items attracted to a magnet. However, it is desirable to have these items, when being displayed radially, to have the maximum amount of space between each other. As a result, the i^{th} item will have the $\left(i \frac{M}{N} \right)^{\text{th}}$ radial alignment coordinates, when displayed in a radial alignment around the magnet.

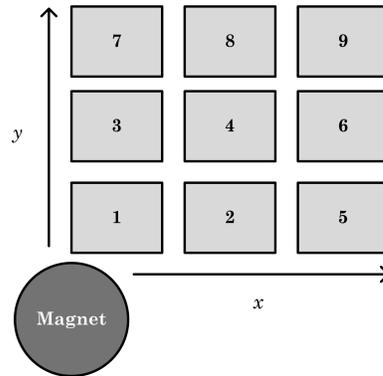


Figure 4.9: The order in which nine items are organized into a grid arrangement.

The pile alignment is computed in exactly the same way as the radial alignment. The difference, is that the radius of the circle on which the alignment is project is considerably smaller, but not zero. Having a non-zero radius allows the items to be dispersed equally when the magnet is set to repel the items it has attracted.

The grid alignment coordinates are computed in a space efficient way. Figure 4.9 shows the order in which nine items are organized into a grid arrangement. As it can be seen, the first item is the closest to the magnet. The second item is directly adjacent to the first item, and so on. The following pseudo-code shows how the grid alignment coordinates are computed.

```

COMPUTE-GRID-ALIGNMENTS()
1   $M \leftarrow$  Max. items
2   $count \leftarrow 1$ 
3   $coords[] \leftarrow$  ARRAY( $M$ )
4  for  $i \leftarrow 1$  to  $\sqrt{M}$ 
5  do for  $j \leftarrow 1$  to  $i - 1$ 
6      do  $coords[count] \leftarrow (xSpace \cdot i, ySpace \cdot j)$ 
7           $count \leftarrow count + 1$ 
8      for  $j \leftarrow 1$  to  $i$ 
9          do  $coords[count] \leftarrow (xSpace \cdot j, ySpace \cdot i)$ 
10              $count \leftarrow count + 1$ 
11 return  $coords[]$ 

```

In this pseudo-code, M is the maximum number of items that a magnet can attract,

coords is an array that holds the computed coordinates, and *xSpace* and *ySpace* refer to the horizontal and vertical spacing between the items in the grid alignment, respectively. That is, each item will have at least *xSpace* amount of space between it and the items that lie directly to its left and right, and *ySpace* amount of space between it and the items that lie directly above and below the item. This algorithm constructs the coordinates of the grid alignment by gradually building columns and rows, until a grid with dimensions $(\sqrt{M}) \times (\sqrt{M})$. The i^{th} row has one more item than the i^{th} column so to ensure that a square grid shape is maintained. Therefore, on the first iteration, a column of size zero is created, and a row of size one is created. On the second iteration, a column of size one and a row of size two are created. In general, on the i^{th} iteration a column of size $i - 1$ and a row of size i are created.

Therefore, the implementation of each alignment has been discussed. It should further be noted that, when items are attracted towards a magnet, they must first reach the centre of the magnet before they are relocated to their respective alignment positions. Once they have reached their respective alignment positions, they are locked there until the magnet is turned off, or the magnetic pull of the magnet becomes negative. When the magnetic pull of the magnet becomes negative, the magnet pushes the items away along directed lines that both the magnet and the items lie on. It should be noted that items do not return to the previous positions they held before being acted upon by a magnet. These algorithms used to compute the coordinates of these three alignments have been implemented in the Magnet strategy.

4.3.5 Magnet Layout

The layout of the magnet is implemented using fairly basic computer graphics. The positioning of the buttons grafted along the side of the magnets has been programmed manually. Textures have been applied to the magnets. At the centre of the magnet, which is used to move the magnet, has a fingerprint texture, to suggest a grip on the magnet. The dial is textured with a gradient which suggests rotation. The buttons are labeled with icons to help the user identify what each button does. An actual screenshot of what the magnet looks like in the system is shown in Figure 4.10.

The Magnet strategy is also responsible for handling events concerning the magnet. For each event that the magnet receives, it must determine where the event occurred,

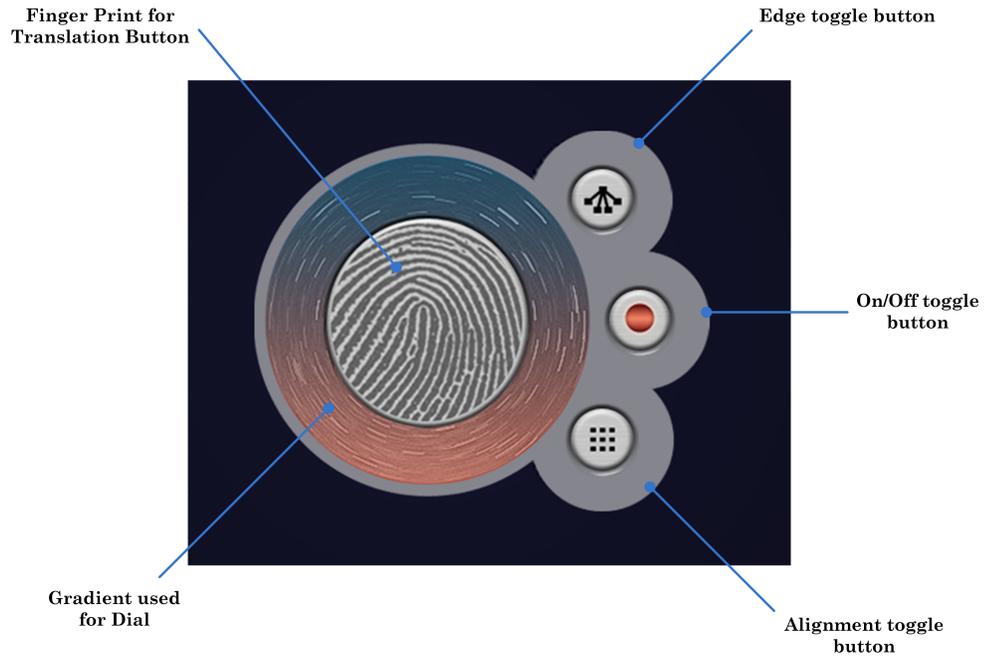


Figure 4.10: Screenshot of the actual magnets, with the dial, centre texture, and toggle buttons labeled.

to determine which part of the magnet is being interacted with. The way this is done is similar to the method for determining whether an item was dropped in the magnet's threshold, since the buttons, the dial, and the centre of the magnets are all circular.

Chapter 5

Results & Scenarios

This chapter will begin by providing an overview of the results of the system that was developed to demonstrate the capabilities of tabletop display magnets. The chapter will then provide scenarios in which these tabletop display magnets could be employed in order to give a better understand the benefits of tabletop displays.

5.1 Results

The primary result of this thesis is a multipurpose tabletop display tool, used to help users reach, group, move, and display data. The tabletop display magnet tool has been implemented, using the Large Display Framework built by the University of Calgary, to work on a SMART Technologies tabletop display. The following shows screenshots of the system used to demonstrate the capabilities of the tabletop display magnet.

The implemented tabletop display magnets are capable of attracting and repelling items within the system's workspace. The dial is successful in being able to modify the magnitude of the magnetic pull. The halo that surrounds the magnet successfully updates to reflect the magnitude of the magnetic pull. Figure 5.1 (a) and (b) show the halo of the magnet with a weak positive pull and a strong positive pull, respectively. The dial is also successful in reducing the magnetic pull so to make it negative. The system, successfully shows this by drawing the halo red, which is shown in Figure 5.1 (c).

Along with the halos, it can be seen in Figure 5.1 that the magnets possess a grey ring that surround the magnet, which represents the magnet's threshold used when the users decide to drag and drop items into the magnet. The threshold has been implemented successfully. Therefore, it is possible for users to drag and drop items into the magnet to create dynamic groups. Once, a magnet has items that it attracts, it's edge toggle button can be pressed to turn on the edges between the magnet and its

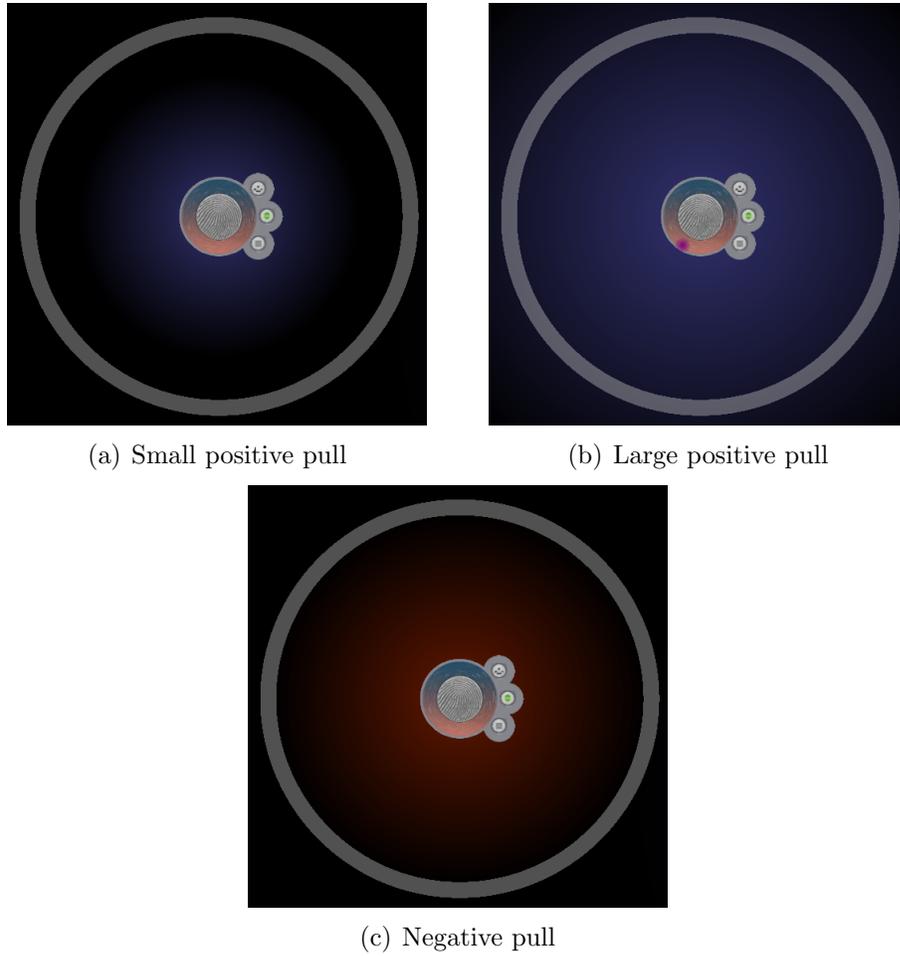


Figure 5.1: (a) Shows a halo that visualizes a small positive pull; (b) Shows a halo that visualizes a large positive pull; (c) Shows a negative pull, or a repelling magnet.

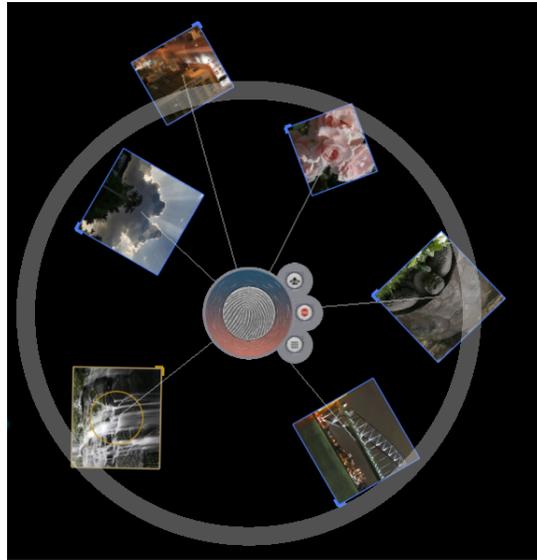


Figure 5.2: Screenshot of magnet with edges connecting items that it attracts.



Figure 5.3: Screenshot of two magnets with edges turned on, showing the memberships of each item relative to the magnets in the workspace.

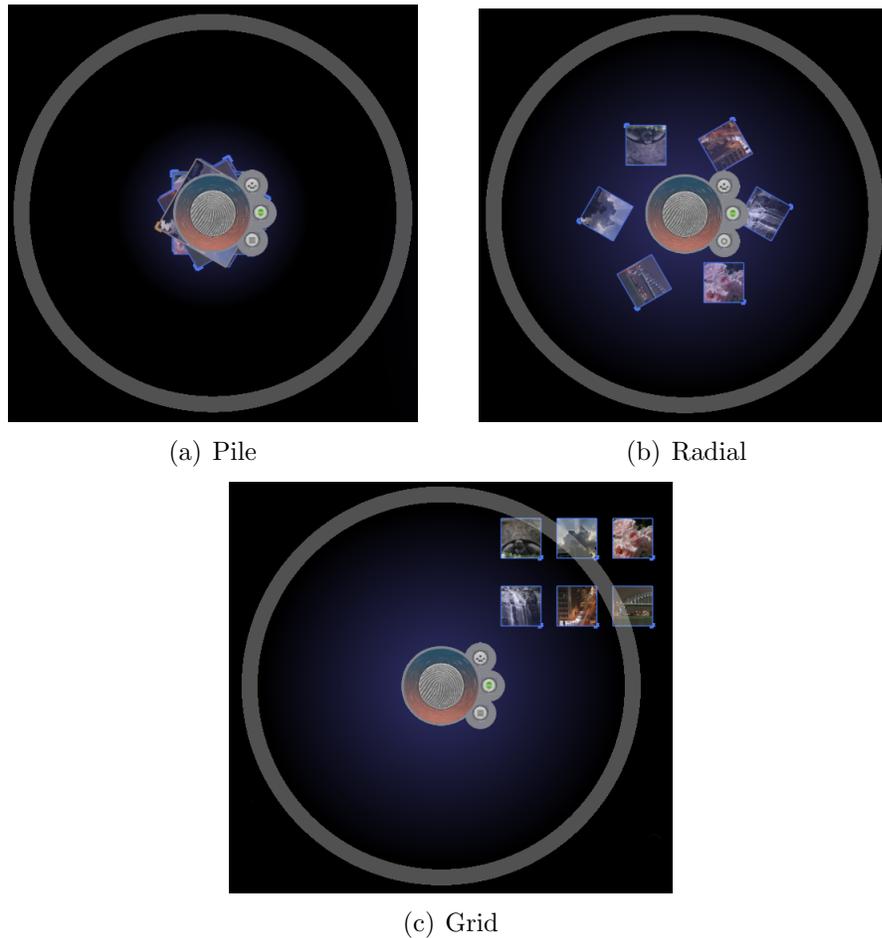


Figure 5.4: (a) Screenshot of magnet with pile alignment turned on; (b) Screenshot of magnet with radial alignment turned on; (c) Screenshot of magnet with grid alignment turned on.

corresponding items. This is shown in Figure 5.2. The system also makes it possible for items to belong to multiple magnets. This can be seen in Figure 5.3, where both magnets have edges drawn between the same items. The edges between magnets can also be removed using the method described in Chapter 3.

The system has also successfully implemented the capability of changing the alignment of the items that are attracted to a magnet when the items reach the centre of the magnet when the magnetic pull is turned on. The alignments can be changed by using the alignment toggle button. The icon on the toggle button reflects which alignment type is selected. The three alignments (pile, radial, and grid) are shown in Figure 5.4.

Finally, the resulting magnets are also capable of being turned on and off by means

of the toggle button grafted to its side. When a magnet is toggled off, the halo that surrounds the magnet is dimmed.

5.2 Scenarios

Having discussed the capabilities of the tabletop display magnets, some scenarios in which these tabletop display magnets can be used can briefly be presented. Of course, tabletop display magnets can be used effectively to manage images among a group of people, the following scenarios will show how tabletop display magnets can be used to accomplish other problems.

5.2.1 Scenario 1: Real Estate

The first scenario involves real estate. The job of a real-estate agent is to sell homes. They do this by listing homes in various real-estate listings in the newspaper, on television, or on the Internet. Also, buyers wanting to buy a house with specific amenities that suit their lives will contact real estate agents, informing them of what they wish to buy. As a result, real-estate agents present the buyers with a set of possible properties according to their desires.

Using a tabletop display and tabletop display magnets, the real-estate agent could sort a variety of houses, represented by image, based on their attributes. For example, one tabletop magnet might group all of the properties that are between \$250,000 and \$300,000, while another magnet could group the properties that have swimming pools. Of course, these two groups may overlap. That is, there are probably some properties that are between \$250,000 and \$300,000 that also have swimming pools. Therefore, for the sake of this scenario, we suppose that a buyer is in fact interested in a property that is between \$250,000 and \$300,000 and has a swimming pool. The real-estate agent could then turn on the magnet grouping those properties that are between \$250,000 and \$300,000. This would show the buyer all of the properties within the desired price range. Then by turning on the edges between the magnet that groups properties based on whether or not they have a pool can be toggled on. Edges drawn from any of the properties grouped by the first magnet means that that property has a pool. As a result, the buyer is presented with a group of properties within the specific price range that also have a swimming pool. However,



Figure 5.5: Screenshot of the tabletop display magnet system being used to find a particular residential property with a desired set of attributes. The magnet on the left attracts those properties that have pools, and the magnet on the right attracts properties that are between \$250,000 and \$300,000. The two magnets working together allow for the properties with both of these attributes to be identified easily.

by visualizing the memberships of the properties to specific magnets, the buyer may find a property within the specific price range but does not have a swimming pool, that is more appealing than the others. Therefore, in this case, the magnets would not only help users organize their data appropriately, but also present new ways of observing their data. A screenshot of the system being used to find desired properties is shown in Figure 5.5. The images in this screenshot were retrieved from the Halifax Real Estate & Halifax Homes Properties website. They are owned by the Dominion Enterprises [15].

5.2.2 Scenario 2: Criminal Identification

The second scenario involves identifying a criminal. Suppose, a person enters a police station reporting a crime and give the police a description of the person who committed the crime. After some period of time, the police collect a series of suspects and take their pictures. The person who initially reported the crime is then called back in to identify the person who committed the crime. Instead of traditional methods of identifying the criminal, the police would present photos of the criminal suspects to person who reported the crime on a tabletop display. Before hand the police would

organize the photos according to attributes possessed by each of the suspects using tabletop display magnets. For example, one magnet could be used to group those suspects with facial hair. Another magnet could be used to group suspects that are balding, and another could be used to group the suspects who wear glasses. Now suppose that the person who reported the crime claims that the person who committed the crime wore glasses and had facial hair. To do this, the police could turn the magnet attracting those people with facial hair on, and toggle the edges on for the magnet that attracts those people with glasses. What would result is that all of the people with facial hair would be grouped together, and some of those people would have an edge drawn from them to the magnet that attracts people who wear glasses. This in turn would narrow the identifier's search. An example of how this scenario would be displayed is shown in Figure 5.6. The images in this screenshot are provided by the AT&T Laboratories Cambridge [3].

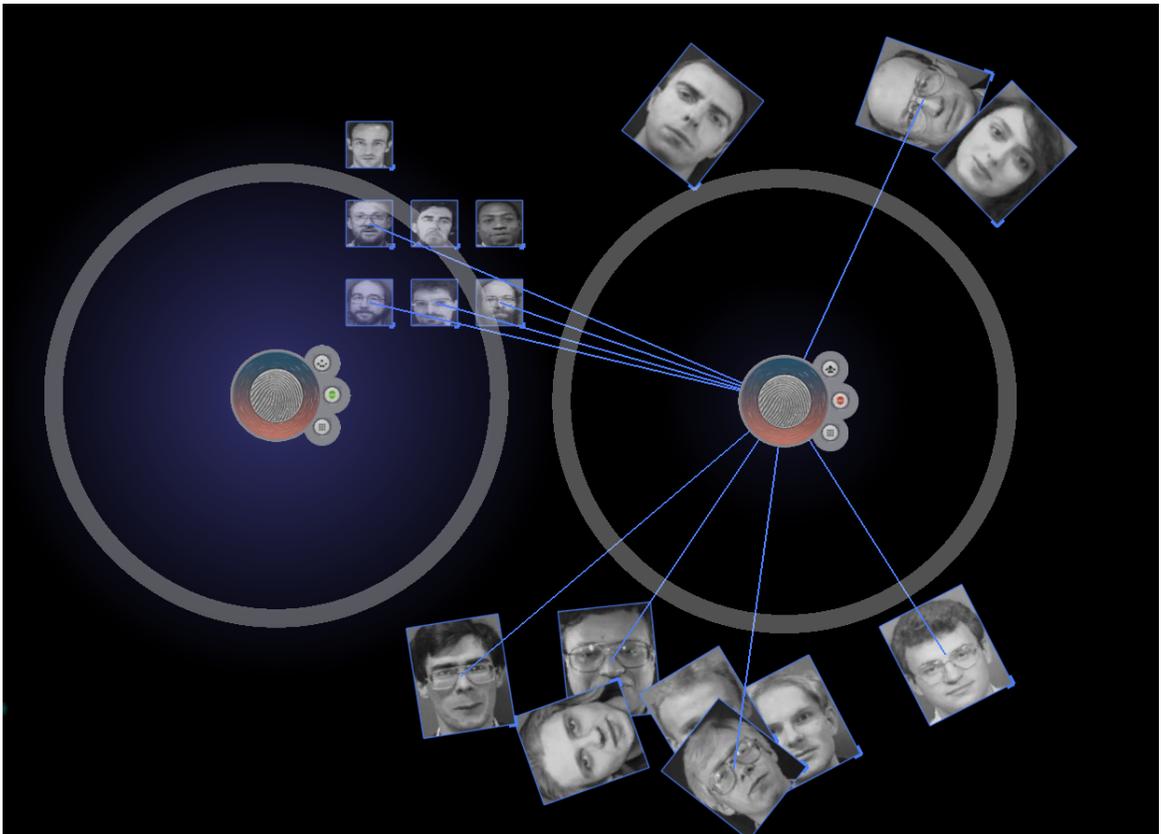


Figure 5.6: Screenshot of the tabletop display magnet system being used to find a particular person with a known set of attributes. The magnet on the left attracts those people with facial hair, and the magnet on the right attracts those people who wear glasses. The two magnets working together allow for the people who wear glasses and have facial hair to be identified easily.

Chapter 6

Conclusions

The foregoing thesis describes the design and implementation of a simple multipurpose tool, called a tabletop display magnet, used to overcome some of the major limitations that are faced when using tabletop displays. The tool strives to help tabletop display users to reach items that are hard to reach due to the large dimensions of most tabletop displays. The tool has also been designed to help users organize and group data item within a workspace. The tool also provides ways for the user to deal with clutter that might develop in a tabletop display's workspace, such as being able to push data items around, and different ways of aligning items.

The system has been implemented using the LDF that was built by the Innovation in Visualization Lab at the University of Calgary, which in turn uses a series of useful libraries and toolkits, such as OpenGL, Qt, and the SMART SDK. The LDF has provided the means of creating an application that provides responsive interactions. What resulted was a working implementation of tabletop display magnets that can be used to show how they can be used for reaching, grouping, moving, and organizing data within a tabletop display workspace. The tabletop display magnets as a result can accomplish many tasks that only several individual tools together are capable of accomplishing.

The research and exploration discussed in this thesis is an attempt to help provide the means for reducing the number of tools needed when interacting with a tabletop display. Despite their relative unpopularity in comparison to desktop computers, tabletop displays have much potential to support vast arrays of applications, especially those involving collaboration.

Chapter 7

Future Work

The research that has been conducted only provides an introduction to tabletop display magnets. Also, due to the limited amount of time, some of the features that were initially planned were not implemented. As a result, there remains much work that could be conducted in the future in regard to tabletop display magnets.

Firstly, focus groups need to be organized and conducted. Such focus groups should consist of 8 to 10 participants each, who would first be briefly introduced to the idea and implementation of tabletop display magnets. The focus groups would then be presented with various scenarios, such as the ones mentioned in Section 5.2. The groups would then begin brainstorming new functionality that would improve the performance of the tabletop display magnet system. Hopefully, a list of limitations and areas of improvement would result from these focus groups, which allow the magnets to be made more refined.

Once these focus groups have been conducted, and the magnets are updated accordingly, a usability study should be conducted in order to evaluate the performance of the magnets relative to other similar tools designed for the tabletop display. The usability study would involve multiple participants using the system to accomplish some set task in a predefined scenario, similar to those mentioned in Section 5.2. The results of the usability study would provide useful indication of whether tabletop display magnets are successful in providing a multipurpose tool that allow people to work more effectively on tabletop displays used for collaborative activities.

There are also other features that were initially planned to be implemented, but time did not permit. One feature involved using the magnets to attract data items based on preset attributes. Without the user having to physically construct the groups that data items belong to, the magnets would be given a condition based on the attributes that the data items possess. For example, in a real estate scenario, a magnet might be set to attract and repel houses that are priced between \$250,000

and \$300,000, and within 10 minutes from the city centre, where price and time from city centre are attributes given to the properties in the data set. Having this functionality would save users from having to form groups by hand, especially in cases where the datasets are very large. However, care would have to be taken in order to implement this functionality so that conditions could easily be set by users unfamiliar with computer programming, and that special consideration be taken for orientation on tabletop displays.

Another feature that could be explored is giving magnets the ability to return data items to the positions they were in before being acted upon by tabletop magnets. This functionality would be particularly useful in scenarios in which the coordinates of data items are important, such as layouts of images for magazine covers.

Bibliography

- [1] Bbc news — science/nature — mouse inventor strives for more. <http://news.bbc.co.uk/2/hi/science/nature/1633972.stm>, 2001.
- [2] Fabricio Anastasio. Innovis - ldf tutorial 1. <http://innovis.cpsc.ucalgary.ca/Software/LDFTutorial1>, 2007.
- [3] AT&T Laboratories Cambridge. The database of faces. <http://www.cl.cam.ac.uk/research/dtg/attarchive/facedatabase.html>, 2002.
- [4] Anthony Eugene Collins. Exploring tabletop file system interaction. In *CHI '07: CHI '07 extended abstracts on Human factors in computing systems*, pages 2171–2176, New York, NY, USA, 2007. ACM.
- [5] Microsoft Corporation. Microsoft surface. <http://www.microsoft.com/surface/>, 2007.
- [6] F. Guimbretire, M. Stone, and T. Winograd. Fluid interaction with high-resolution wall-size displays, 2001.
- [7] Mark S. Hancock, Frederic D. Vernier, Daniel Wigdor, Sheelagh Carpendale, and Chia Shen. Rotation and translation mechanisms for tabletop interaction. In *Proceedings of IEEE International Workshop on Horizontal Interactive Human-Computer Systems (TableTop2006)*, pages 79–86. IEEE Computer Society, 2006.
- [8] Tom Henderson. Newton’s law of universal gravitation. <http://www.glenbrook.k12.il.us/GBSSCI/PHYS/Class/circles/u6l3c.html>, 2007.
- [9] Uta Hinrichs, Sheelagh Carpendale, Stacey D. Scott, and Eric Pattison. Interface currents: Supporting fluent collaboration on tabletop displays. In *Proceedings of the 5th Symposium on Smart Graphics*, pages 185–197. Springer Verlag, 2005.
- [10] InnoVis. Innovis - ldf. <http://innovis.cpsc.ucalgary.ca/Software/LDF>, 2007.
- [11] Tobias Isenberg, Simon Nix, Martin Schwarz, Andr Miede, Stacey D. Scott, and Sheelagh Carpendale. Mobile Spatial Tools for Fluid Interaction. Technical Report 2007-872-24, Department of Computer Science, University of Calgary, Canada, July 2007.
- [12] Gordon Kurtenbach and George Fitzmaurice. Guest editors’ introduction: Applications of large displays. *IEEE Comput. Graph. Appl.*, 25(4):22–23, 2005.
- [13] Mistubishi Engineering Research Laboratory. Merl - diamondtouch. <http://www.merl.com/projects/DiamondTouch/>, 2008.

- [14] Andr Miede. Realizing responsive interaction for tabletop interaction metaphors. Master's thesis, Otto-von-Guericke-University of Magdeburg, Germany, 2006.
- [15] Halifax Real Estate & Halifax Homes Properties. Featured listings. <http://www.halifaxmetrohomes.com/content/featuredlistings.html>, 2008.
- [16] George Robertson, Eric Horvitz, Mary Czerwinski, Patrick Baudisch, Dugald Ralph Hutchings, Brian Meyers, Daniel Robbins, and Greg Smith. Scalable fabric: flexible task management. In *AVI '04: Proceedings of the working conference on Advanced visual interfaces*, pages 85–89, New York, NY, USA, 2004. ACM.
- [17] Stacey D. Scott. Territory-based interaction techniques for tabletop collaboration. doctoral symposium presentation. In *Conference Supplement of the ACM Symposium on User Interface Software and Technology (UIST)*, pages 17–20. ACM Press, 2003.
- [18] Stacey D. Scott, M. Sheelagh T. Carpendale, and Stefan Habelski. Storage bins: Mobile storage for collaborative tabletop displays. *Computer Graphics and Applications*, 25(4):58–65, 2005.
- [19] Edward Tse, Saul Greenberg, Chia Shen, and Clifton Forlines. Multimodal multiplayer tabletop gaming. *Comput. Entertain.*, 5(2):12, 2007.
- [20] Colin Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2004.