Art of Defense: A Collaborative Handheld Augmented Reality Board Game

Duy-Nguyen Ta Huynh,* Karthik Raveendran,† Yan Xu,‡ Kimberly Spreen,§ Blair MacIntyre¶
School of Interactive Computing and GVU Center
Georgia Institute of Technology

Abstract

In this paper, we present Art of Defense (AoD), a cooperative handheld augmented reality (AR) game. AoD is an example of what we call an AR Board Game, a class of tabletop games that combine handheld computers (such as camera phones) with physical game pieces to create a merged physical/virtual game on the tabletop. This paper discusses the technical aspects of the game, the design rationale and process we followed, and the resulting player experience. The goal of this research is to explore the affordances and constraints of handheld AR interfaces for collaborative social games, and to create a game that leverages them as fully as possible. The results from the user study show that the game is fun to play, and that by tightly registering the virtual content with the tangible game pieces, tabletop AR games enable a kind of social play experience unlike non-AR computer games. We hope this research will inspire the creation of other handheld augmented reality games in the future, both on and off the tabletop.

CR Categories: K.8.0 [Personal Computing]: General—Games I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques

Keywords: augmented reality, collaborative game, handheld game

1 Introduction

Over the past few years, advances in the capabilities of 3D graphics, processing and display technologies in PDAs and camera phones have made these devices attractive platforms for mobile *augmented reality* (AR) games[1]. While AR has long been the subject of research labs and science fiction novels, the next generation of mobile devices will finally enable a huge number of people to experience applications and games that put graphics out in the world around them. The current consumer devices are barely up to the task, but have allowed a variety of games to be prototyped, both by academic researchers (e.g., [Henrysson et al. 2006], [Wagner et al. 2005], [Mulloni et al. 2008], [Xu et al. 2008]) and non-academics[2]. In many cases, these games have been created to demonstrate or test new technologies, especially software for tracking the mobile device relative to the world (e.g., marker trackers such as StbTracker[3]). Designing handheld AR games is challenging, both because the technology is difficult to work with, but also because we do not yet understand how to create effective AR play experiences. In our work, we have been focused not on advancing tracking and graphics technology, but rather on exploring the affordances of AR for play, especially as a vehicle for social gaming.

Figure 1: *The Art of Defense game. (Top) Views through the phones. (Bottom) Two views of the play space.*

In this paper, we present our work on *Art of Defense* (AoD), a cooperative handheld AR game. AoD is an example of what we have come to call *AR board games*, or games that use handheld computers (e.g., camera phones) and tangible props to combine the tangible elements of board games with the continuous simulation of a computer game. The handheld AR interface merges the physical and digital worlds, creating the illusion for the player that their device is a window on the hybrid world. Using computer vision techniques (e.g., in AoD we combine marker-based tracking, color recognition, and contour recognition) allows the players to control the game with tangible objects, using physical manipulation of the props (tiles and tokens) to control objects and actions in the game. In this paper, we present AoD as an example AR board game, as well as discuss the technical aspects of the game, the design rationale and process we followed, and the resulting player experience.

AoD is a strategy-based "Tower Defense" style game, in which two players work together to protect their central tower from waves of attacking enemies. Over the past year, the game has been tested informally and refined, including numerous demonstrations to visitors to our lab and attendees at conferences. Most recently, we ran a formal user study with six pairs of participants. The subjects found the game fun to play, and allowed us to confirm that the collaborative social experience created by this AR game encourages

*e-mail: duynguyen@gatech.edu
†e-mail: kraveendran@gatech.edu
‡e-mail: yxu7@gatech.edu
§kim.spreen@gatech.edu
¶e-mail: blair@gatech.edu

[1] In this paper, we use AR to refer to games where the graphics and/or sound are tightly *registered*, or aligned, with the players view of the physical world. AR is part of the larger category of *mixed reality* games or *alternate reality* games, which also include games that use the player location as input to the game, or make references to the world around the game.

[2] For example, see the prototypes at http://www.cellagames.com.

[3] http://handheldar.net/stbtracker.php

communication and social behaviors of the sort not seen in other kinds of collaborative computer games.

This project is part of a long-term effort aimed at understanding the physical and social interactions facilitated by handheld AR games. Our previous work on Bragfish [Xu et al. 2008], a co-located competitive handheld AR game, found that the AR interface enabled players to integrate their perception of each others physical movements, and to incorporate direct social interactions, into the game play. In contrast, AoD aims to create a cooperative experience. In Bragfish, we used a fixed game board; here we use tangible game board pieces and input tokens to increase the players sense of agency in the game. A tile-based game board is also much more portable (making it well suited to games that use mobile phones), and allows us to naturally limit the player's visibility (to areas near the hex tiles) as part of the game design, further encouraging interaction with the physical space.

Through the design and evaluation of AoD, we hoped to answer the following research questions:

- What are the affordances and constraints of handheld AR?
- How can we transfer our understanding of the technology into the game experience that we want to create?
- What kind of play experience does the game support?

The paper is organized as follows. First, we discuss how our work is related to, and distinguished from, prior literature. Second, we discuss the affordances and constraints of handheld AR technology, and how they impacted the design of the game. Third, we introduce the game design itself, and the design rationale we used. Last, we report the qualitative and quantitative findings from a user study with six groups of players.

2 Related Work

There has been a great deal of interest in making games with handheld devices, due to their size and mobility. Early outdoor mixed reality games on handheld devices (e.g., "Treasures" [Barkhuus et al. 2005], "Pirates!" [Björk et al. 2001]) used GPS for local positioning and placed virtual objects at locations in the physical world. However, those games did not use AR technology to enhance the player's perception of virtual objects. Without consistently aligning the virtual world with recognizable landmarks in the physical world, it is difficult for users to merge the two worlds in their heads in the same way they do with AoD.

Recent advancements in mobile phone hardware have enabled the creation of 3D augmented reality applications on handheld devices. Invisible Train [Wagner et al. 2005] was one of the first AR systems on a handheld device that could track fiducial markers and render 3D graphics in real time. However, it was not a fully realized game. AR Tennis [Henrysson et al. 2006] not only tracks the markers and renders the scene in 3D, but also utilizes the handheld device as an interaction tool. Nevertheless, AR Tennis suffers from the problem that the handheld device is not only an interaction tool but also the display system; the game is awkward because players cannot easily control their view of the virtual world while simultaneously using the device as a tennis racket. [Mulloni et al. 2008] and our previous game, Bragfish, explored handheld AR games in a social context but did not fully utilize the affordances of handheld AR, such as its potential for tangible interaction techniques, to more fully draw the player into the combined physical/virtual play space.

Tangible interaction techniques have long been a goal for AR system design, since the freedom of user movement afforded by AR creates the need for direct interaction with the physical elements of the merged world. Many techniques have been designed for AR applications based on input devices, such as paddles or cubes [Billinghurst et al. 2000] with markers on their visible surfaces. However, these markers can only define the spatial relationship between the physical and virtual worlds, so, virtual and physical elements are still separated even though they coexist on the same space. The naturalness of sketch-based interfaces in games[4] has inspired us to transform user sketches into meaningful virtual game objects. Furthermore, the use of physical objects in board games has also inspired us to design tangible game tokens as our game pieces.

A number of research projects have explored such hybrid physical/virtual games. For example, in "False Prophets", players use tangible pieces to move around a digital game board projected onto a touch-sensitive table [Mandryk and Maranan 2002]. In "Wizard's Apprentice", the physical board and digital display are shown to the player in parallel [Peitz et al. 2006]. In these games, the physical game board is wired to recognize player's actions. But in AoD, the mobile device is in charge of actively recognizing these actions instead. The hybrid world is shown via the displays on the players' handheld devices. The mobility of the interface allows the game to be played anywhere. Each of these different interface technologies and approaches will lead to different play experience.

3 Handheld AR Issues

Our main research goal is to explore game structures and mechanics that could leverage the advantages of AR on handheld devices. Mobile AR offers a range of game elements that designers can leverage, drawn from both digital games (e.g., 3D interactive graphics, sound and various interaction techniques) and traditional board games (e.g., physical game pieces and game boards). Players may have an increased awareness of their physical space, and can easily interact with other players and the environment [Xu et al. 2008].

Although recent advances in AR technologies have enabled AR to be used in commercial console games[5], game designers need to be aware of a range of sometimes severe technological limitations when designing such games for mobile devices. In this section, we discuss the potential as well as technical limitations of mobile AR games, and suggest some possible design options that can help overcome these limitations, or leverage them as part of the game design.

3.1 Display and Tracking

The appeal of merging 3D graphics with the physical world, and the possibility of viewing it naturally from any viewpoint, are two oft-cited advantages of AR over conventional computer games. By moving AR from head mounted displays (HMDs) to handheld devices, the user is freed from the ergonomics issues of current HMDs. Using mobile phones for AR also makes AR interfaces ubiquitous: people usually have their phones with them, and can therefore play the games anywhere and anytime.

A major problem of mobile AR is that the small viewing area of handheld devices limits the amount of the merged world a user can perceive at any one time. Fortunately, building on the idea of seamful design [Barkhuus et al. 2005], the small viewing area can be used as a design element. Since players cannot observe the whole hybrid play space at one time, game designs that purposely limit information about what is happening in some parts of the game world

[4]For example, Crayon Physics http://www.crayonphysics.com/

[5]For example, see Eye of Judgment: http://www.eyeofjudgment.com

may lessen the impact of the small visible area. For instance, numerous strategy games, (e.g., Warcraft III) use a "fog-of-war" that prevents players from viewing game activity in areas unoccupied by player units. Moreover, the ability to naturally control the viewing location allows the player to very easily move from one view to another, and from a farther-back overview to in-close detail views.

The second problem with mobile AR lies in the tracking technology for accurate and real-time registration of the virtual and physical worlds. Although recent advances in tracking technologies enable very accurate and robust 3D tracking in large scale environments, they are still computationally too expensive to execute on current smart phones. Wagner has demonstrated a natural feature tracker for mobile phones [Wagner et al. 2008b] that tracks off of 2D images, but it cannot yet track many objects at the same time, which severely limits the possibilities for user interaction with different physical game pieces. In this paper, our solution is based on a conventional marker tracking approach for AR [Wagner et al. 2008a], but the discussion is general enough to be applied to a natural feature tracking method when such technology is more mature.

The current limits of processing capability and camera resolution put constraints on the range of possible distances between the handheld device and the tracked object. On one hand, the captured image resolution has to be small enough to guarantee fast image processing on the phone and real-time tracking speed. On the other hand, the area of the marker viewed in the captured image has to be large enough for the tracking algorithm to calculate a stable and accurate 3D camera pose. Consequently, the camera resolution limits the player's range of motion; the camera needs to be far enough from the game space to get an entire marker in view, but not so far as to lose tracking. While the whole game space could be large, this constraint again limits the amount of information the player can perceive from the AR world.

Another problem facing mobile AR is the tradeoff between portability and fixed game boards. Multi-marker tracking techniques, which use a large board with many pre-located markers to provide a large play space, are very common in AR. However, a large marker board conflicts with our portability goal, since it is inconvenient for players to carry a large game board with them. Advanced methods that track physical objects with no prior knowledge (e.g., [Klein and Murray 2007]) could address this issue in the future, but for now, mobile AR game designers have to address this constraint in order to make their games portable.

Instead of using a fixed large marker board, we build on the tile-laying game mechanic made popular by games such as Klaus Teuber's Settlers of Catan[6]. In AoD, each hex-shaped tile has a trackable marker on it. One hex tile serves as the center of the world and during the game, the player puts new markers next to the existing ones to build up the game map. By limiting the number of tiles, players are required to remove existing tiles and put them at new locations to explore the world. As long as the map configuration of those tiles is known and consistent, the system can steadily track the game space and register the virtual content. By leveraging marker tiles, or similarly compact game pieces, handheld AR games can become truly portable using current technology, because the player can easily carry the tiles with them everywhere.

A dynamic marker-tile game board is a compelling mechanic for AR game design. Aside from making the game portable, the marker tiles also require tangible interactions with the game pieces. The fact that the AR world can only be seen where the markers exist reinforces the limited visibility of the handheld, as discussed above, allowing for hidden information or "fog-of-war" effects.

[6]http://en.wikipedia.org/wiki/Settlers_of_Catan

3.2 Tangible Interaction

As discussed in the previous section, the dynamic tile-based map building mechanic allows players to physically interact with the game pieces. However, the marker tile map is only the ground on which the augmented content sits in the physical world. We would also like to allow players to interact with virtual game objects.

There have been examples of tangible interaction in mobile AR (e.g., [Henrysson et al. 2005], [Henrysson et al. 2007]). In the simplest form, one can use buttons on the phone to interact with the virtual world. To enhance tangible interaction, one could use a tracked paddle, cup, or cube (e.g., by putting with markers on them, as in [Billinghurst et al. 2000]). One could also detect the user's gestures, or sense the motion of the phone [Henrysson et al. 2007]. We do not explore the last two options in AoD, although newer devices such as Apple's iPhone point toward such techniques being practical in the future.

Using objects with markers on them to interact with the virtual world creates a well-defined relationship between the physical and virtual worlds: the markers define the location and meaning of the virtual objects in the physical world. The biggest hurdle is that in AR, the markers typically lack meaning to humans, and simply serve as placeholders for virtual objects. Even if one can put some meaningful picture inside the marker, the number of AR objects is bounded by the number of marker patterns and their locations in the physical space. Therefore, one objective of our design was to seek out interaction techniques that go beyond simple spatial relationships between physical and virtual worlds.

Early in our design, we decided to use markers that have a hollowed out interior [Wagner et al. 2008a], allowing the user to draw a sketch that can be analyzed by the game. This sketch can be used to control the game, creating units and changing their properties. Parameters such as accuracy of the sketch could be used to enhance the resulting virtual unit, create variations, and provide greater feedback to the user by aligning the appearance of the marker tile to the virtual content. In this way, the sketch techniques could help to bridge the semantic gap between the physical and virtual world. Although the sketch interaction is compelling, in a mobile AR game it is difficult for people to hold the phone, steady the marker tile and sketch at the same time. Sketching also takes time, which is not suitable for fast paced games.

An alternative to sketching is to use pre-designed tokens as place holders for virtual objects, akin to game pieces in board games. Such a token needs to be carefully designed so that, on one hand, it could convey some meaningful concept to the player about the object it represents, but, on the other hand, remain simple enough to be recognized by the phone accurately and quickly. Compared to sketches, tokens are more tangible and easier to use in the game, but they are fixed in property and meaning.

4 Game Description

Art of Defense is a networked two player cooperative game inspired by the "Tower Defense" genre. The objective of the game is to survive as many waves of enemies as possible by preventing them from reaching the base (see Figure 1). Enemies approach along predefined paths that converge at the base and deal damage proportional to their strength. The game ends when the health of the base tower falls to zero. Players begin with a common pool of money that can be used to create new towers or upgrade existing ones. These towers attack enemies and earn money for each killed enemy unit.

Players start with a single hexagonal tile (corresponding to the base) placed on the table. No tokens may be placed on this tile. There are 15 map tiles excluding the base (much fewer than the current 10×11 game map). The map can be explored by placing additional tiles adjacent to existing ones (and adding them by viewing both tiles through the phone). Adding a new tile reveals the map corresponding to that region (including enemy units) and also shows the terrain in the neighboring tiles. Tiles can be removed and reused (towers can only be created on tiles that are in place in the world).

Each player is assigned a unique color (either red or blue) that remains constant throughout the course of the game. The red player can only build red towers, which in turn can only damage the red enemies. Similarly, the blue player is restricted to blue towers, which solely attack blue enemies. Towers can be built on empty tiles on the map by placing a triangle token of the appropriate color and turning it to point the tip in the desired direction of fire. The player adds the tower to the game by viewing the tile through the phone and pressing a button. Each tower costs a fixed amount to construct but can be rotated at any time for no cost.

There are three upgrade tokens (magenta, green and black) corresponding to an increase in the damage dealt to a unit, rate of fire and range, respectively. Each may be applied up to a maximum of three times per tower. To upgrade a tower, the player must place the token on the tile containing the tower and confirm the action on the phone. The final score is based on the time since the start of the game that the base tower has remained standing.

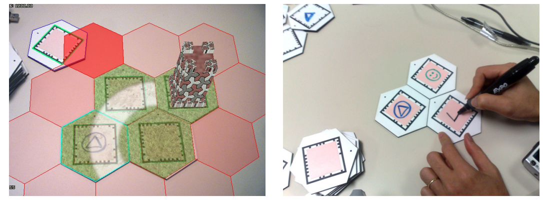

(a) (b)

Figure 2: *Old interface concepts. (a) First interface for map building. (b) Sketch-based tower creation.*

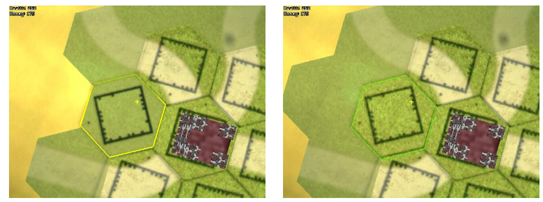

(a) (b)

Figure 3: *Current UI for adding tiles (a) New tile is highlighted in yellow (b) The added tile is highlighted in green*

5 Game Design Discussion

Picking the strategy genre allowed us to incorporate elements from board games (map building using hexagonal tiles, tangible game tokens) and their computer counterparts (limited visibility due to fog of war, real-time unit movement). Tower defense games typically need little to no micromanagement of player units, with the main player actions being the placement of towers (which automatically fire at enemies without player intervention) and upgrading of towers. Further, the pace of such games is moderate and does not require rapid button presses (clumsy on most mobile devices) or violent camera movements (disruptive for vision based tracking). Combining the gameplay of tower defense with the exploration aspect of other real-time strategy games (such as Blizzard Entertainment's Starcraft and Microsoft's Age of Empires) aligned perfectly with our vision of the player interacting with game pieces to deploy units and to also develop an understanding of a larger space than could be seen on the phone screen at a single instant of time.

AoD was developed iteratively over the past year, and we received feedback from a wide variety of people through the development cycle. We began with paper/physical prototypes to understand the interactions that we were designing, and then developed a version for a desktop computer to test the specific game mechanics with tracking and augmented graphics. Once we had ironed out all bugs, we ported the code to the Symbian platform. Building a desktop version first saved us a great deal of time and effort. However the powerful CPU made it very easy to overlook the bottlenecks that could exist on the lower-powered phone. On numerous occasions, we had to scale back or discard features after realizing that the phone could not support the level of complexity that we desired. For instance, advanced pathing for enemies, higher polygon models and detailed textures did not make it into the current version. In the following sections, we describe the game components and their iterative refinement through the development process.

5.1 Map Building

We opted to embed a hollow square marker in a hex game tile for several reasons. The shape permits the addition of tiles in six directions and avoids the ambiguity of diagonal movement that is commonly seen in games using square tiles. Furthermore, the distance covered is equal in all six directions. Finally, we found it easier to handle hex tiles than square ones because they could be gripped without obscuring the marker. We used a laser cutter to fashion a set of hex tiles and soon discovered that this shape had the added benefit of better inter-tile alignment and led to more stable tracking because pieces were less likely to move during gameplay. We experimented with different sizes for the tile such that the inner space was large enough to permit sketching without the user accidentally crossing the boundary. We finally settled on a hollow marker with a size of 5.5cm, with each side of the hexagon being 4.85cm.

To add a tile to the map, the player had to place the new tile next to an already added one and press a button to confirm this action. In our earliest prototype, only portions of the map with tiles under them were revealed to the user (see Figure 2(a)). This proved to be confusing because it was not clear which direction the players needed to explore to find a path, leading them to add and remove tiles frequently. In the next iteration, we added limited visibility near tiles, where spaces next to tiles would show a dimly lit version of the terrain, but not enemy units (see Figure 2). This visual cue was sufficient for players to add tiles in the areas that they believed were important (usually around paths leading to the tower).

Adding a tile requires both the to-be-added and one of the existing tiles to be in view. However, with the limited field of view of the phone's camera, it is possible to lose sight of all of the existing tiles. Designing an intuitive interface on the phone for supporting this map-building feature proved to be challenging. Our earliest prototype displayed the possible positions, at which a new tile could be added, with the position closest to the new tile highlighted in a darker shade. However, it was not obvious to the user if both tiles were being tracked at the moment of pressing the button to confirm the action. This led to unnecessary key presses and frustration in poor lighting conditions. We remodeled this system to highlight the existing tracked tiles in green, while new tiles are highlighted in yellow if they are close enough and are being tracked, or in red if they are too far away (see Figure 3). Once added, the new tile is highlighted in green as well. While this interface does not handle

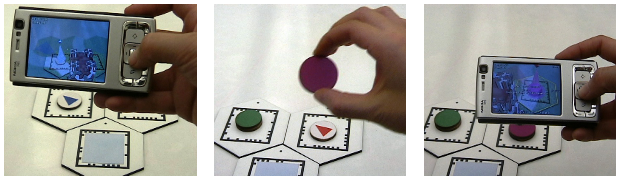

(a) (b) (c)

Figure 4: *Building and upgrading a tower. (a) Triangular token creates and orients tower. (b) Power-up tokens. (c) Applying a power-up to a tower.*

lost tracking either, we have found that players tend to look for the yellow signal as an indication to confirm the action and this made the tile addition process work smoothly under reasonable lighting conditions.

5.2 Sketching

Our original design for AoD allowed players to sketch shapes (corresponding to different units) in the empty space inside the hollow marker (see Figure 2(b)). Our first prototype supported lines (walls), squares (traps), triangles (directional towers) and circles (omnidirectional towers).The players had to look at the tile with the phone to confirm the correct shape was detected. The size and position of the hand drawn shape, as well as the accuracy of the shape, were used as parts of the game.

However integrating the sketch recognition effectively into AoD proved to be difficult. Even after we found suitable erasable materials for the center of the tiles, implementing a robust, fast sketch recognizer was not possible on top of everything else the mobile phone needed to do during each rendering cycle. Most sketch-based games (on a computer or Nintendo DS) use stylus based input that is highly accurate because they receive a densely sampled version of the shape. In our case, recognition had to be performed on a single image at a fairly low resolution with considerable error (non-closed shapes, faded segments due to varying pressure applied by the user). While more expensive thresholding algorithms alleviated some of these issues, sketching also forced the user to place the phone down or risk moving the other game tiles while performing the action. Ultimately, we decided against adding this feature to the final version of the game. While sketching could provide a unique means of user expression, our prototype was unable to utilize this fully due to technology limitations.

5.3 Game Tokens

As an alternative to sketching as a means of tangible interaction, we again looked to board games for inspiration. Physical pieces are appealing because they lower the learning curve for the player and create a mapping from the physical to virtual world. Further, the detection algorithm for sketching could be reused for these pieces, without many of the problems created by hand-drawn sketches. For the current game, we use exactly one kind of player unit: a directional tower that is represented by a triangle (see Figure 4(a)). Players could rotate the piece and look at it through the phone to have it shoot in the desired direction. While a more elaborate game could easily use more types of tokens, we wanted to refine the interactions and keep the game relatively simple to avoid overwhelming the casual player.

Switching to game tokens also allowed us to do away with menu screens for unit upgrades. Using three uniquely colored tokens for the various upgrades (see figure 4) meant that players simply had to place the desired token on a tower to upgrade it. To support this operation, we added a color detection phase to our existing algorithm and calibrated the colors according to the lighting conditions and the phone's camera when the game starts.

5.4 Cooperative Play

One of our goals for AoD was to create a collaborative experience and test our intuitions about how AR would foster social interaction. Redesigning a primarily single player game concept (tower defense) into a cooperative game that fostered the desired interactions required more than just making the game harder. We considered several ways of dividing responsibilities amongst the players to encourage cooperation (for instance, one option was to allow one of the players to observe and explore the map while the other added towers), but we found such unequal divisions of labor to be disruptive to the player experience and to the balance of the game. Instead, we decided to allow both players to explore the map but each to only add towers of a certain color. Towers of one color (e.g., blue) could inflict damage on enemies of the same color, but not on other (e.g., red). This helped us to avoid the scenario where players would split up the space physically and focus solely on their halves. Furthermore, the pool of money was shared between the players, forcing them to make strategic decisions together. An advantage of this color-based division was that it allowed us to design levels that could test the extent of cooperation between the players by varying the proportion of red and blue enemies in each wave. For example, if players did not communicate with each other, a path that was poorly guarded against blue enemies could be easily overrun.

6 Technology

Our game targets the Nokia N95 running the Symbian OS. The N95 contains an ARM-11 based Texas Instrument OMAP2 processor running at 330MHz, and a 5 megapixel camera that can capture video at 640×480 up to 30 frames per second (fps). The graphics pipeline on the N95 is designed for OpenGLES 1.1[7], a compact version of OpenGL for embedded devices.

After receiving an image from the camera for processing, we need to display it on the screen by uploading it to the texture memory. OpenGLES only accepts textures with size $2^m \times 2^n$. Unlike OpenGL on the PC, updating part of the texture using glTexSubImage2D on the phone with OpenGLES is much slower than uploading a whole new texture with glTexImage2D. Moreover, pixel-by-pixel copy between different image buffers for different purposes (texture mapping and tracking) is not cheap on the phone. For all of these reasons, there is an important tradeoff in the image resolution for fast processing speed, tracking accuracy and viewing quality; we needed to carefully choose a good resolution that balances those factors. Our experiments found that a 256×192 video image offered a good balance of viewing quality, processing speed and tracking accuracy. Images at 320×240 are also a good choice, trading off reduced speed for more stable tracking and registration.

We use the Edgelib cross-platform game engine[8] to develop our game. Edgelib version 3.3 has some rendering problems when we combined it with native OpenGLES (which we need to do to set the projection and model-view matrices to match the camera view). We eventually replaced most of Edgelib's graphics engine by our own, based on the Milkshape model format and optimized with fixed-point operations for fast computation on N95 hardware. Despite the rendering problems, Edgelib provides us a good framework for

[7]http://www.khronos.org/opengles
[8]http://www.edgelib.com

cross-platform development with easy input button mapping, sound (using the Hekkus library[9]) and networking.

For marker tracking, we use StbTracker library[10] and manage the data structures for dynamic multi-marker tracking. By default, StbTracker uses all observed correspondences of marker corners in the map to optimize the camera pose. Although this approach is accurate and stable for fixed markers in a large multi-marker board, it is not well suited to our situation where the marker tiles are not rigidly fixed relative to each other. The gaps and improper alignment of the marker tiles will cause the optimization process to converge to some incorrect local minimum, leading to errors and flickering in registration. Consequently, we simply use the 4 corners of the largest marker observed in the current frame for pose optimization. Due to the fact that the biggest marker does not change frequently, this simple approach helps reduce the flickering problem, is robust to sloppy tile placement, and also speeds up the optimization process significantly.

As noted above, recognition of hand-drawn shapes is a difficult problem. We have tested different image processing techniques (e.g., edge finding, Hough transforms) to extract simple geometric primitives from the sketch and then apply a hierarchical classification system, but most are too expensive to operate on the phone in real time. We ended up using a variant of the approach used by most AR marker trackers to quickly recognize simple geometric shapes such as lines, triangles, rectangles or circles. The algorithm works by first thresholding the image, extracting the contours of any blobs inside a marker, searching for extreme points on the contours, and determining the contour shapes by the number of extreme points. In order to enhance the recognition accuracy, we use an adaptive thresholding technique [Bradley and Roth 2007] to cope with the uneven local lighting properties on different parts of the sketches, and guarantee that the sketch contour is continuous after being thresholded. For color-filled geometric shapes on the tokens, we use the same recognition algorithm, but without using the adaptive thresholding technique since the filled interior of the shape already guarantees the boundary is continuous.

In order to calibrate the camera for color recognition, we place all color tokens to be recognized at specific places relative to the base marker (the position of each token is highlighted by rendering a circle of that token's color in the alpha channel). After that, we calculate the mean color of the camera pixels in a 5×5 region at the center of each circle. During game play, the sketch or token color is classified by the closest Euclidian distance of its average pixel colors to one of those mean colors.

7 User Study

7.1 Goals

To understand the play experience and improve the game design, we conducted a small user study with twelve participants. Specifically, we hoped to understand how the augmented reality interface affected game play, and the kinds of cooperative play that occurred. These research questions were not merely designed to gather the feedback to improve AoD. Rather, we were interested in finding the patterns of behavior related to game interface and cooperative game design to inform the design of future handheld AR games.

[9]http://www.shlzero.com
[10]http://handheldar.net/stbtracker.php

7.2 Participants

We recruited twelve participants on campus. We recruited each person separately and ensured that each pair of participants were not friends before the study. All of the participants were undergraduate and graduate students from the computer science and computational media majors, aged 21-26. Two of them were female. The participants were from six different countries. Six participants did not use English as their native language. All but two had not played augmented reality games before the user study. All but two had the experience of playing with strangers in online games or sports. In the following sections, the participants are referred to using a convention like G1-A (the number is the group id, and the letter is the player id in the group.)

7.3 Procedure and Setting

The user study included three parts. First, the players learned the game interface. The participants were given a phone and game pieces, without the game running. A researcher gave an introduction to the game controls by following a pre-written script and showing a pre-recorded instructional video. Second, the participants played the game together. Each team was given at most three times to play. The video of the game sessions and player behavior was recorded. One observer took notes while the game was played. Third, the participants filled out the questionnaire independently after playing the game. Afterwards, a semi-structured interview was conducted to understand more about the incentives and experience.

The user study was set up in a research lab. The two players sat on each side of the table, facing each other. The physical game tokens and tiles were piled at one end of the table. An instruction sheet showing how to use the phone buttons was provided to each player, for reference. To capture and synchronize the video of play behavior and game events, we recorded the game screens and top-down video for both players using a video quadbox (see Figure 5).

Figure 5: *User Study Setting*

7.4 Measures

7.4.1 Recorded Video and Observation

The recorded video synchronized the game screens with the top-down videos that captured the physical movements of both players. With this data, we were able to reconstruct what went on during game play, how the players moved and the problems they encountered. In addition, one researcher took observation notes during the game play, to record interesting moments of play and to generate related questions for the interview.

7.4.2 Questionnaire

The questionnaire was designed based on previous studies (e.g., our own work, [Mandryk and Maranan 2002]), and was divided into three sections: feedback about the game in general, evaluation of the game interface, and feedback about social interaction. We used cross-checking questions to address the problem of misreporting.

7.4.3 Interview

At the end of the user study, we conducted a semi-structured interview. While the observation and video focused on finding the behavioral patterns, the interview focused on understanding the subjective experience of the players, including the strategies they adopted, their awareness of the other player, and the comparison between this game and other games they had played before.

7.5 Findings and Discussion

7.5.1 Feedback About The Game

Through the user study, the participants gave us positive feedback about AoD (see Figure 6). We interviewed the participants about what they liked and disliked about the game. "Cooperative play", "Tangible interface", and "Game concept" were most frequently mentioned as pros of the game. "The graphics went off when I moved my phone" (registration problem), "blocking the other one's view", "need a better overview of the bigger world" were most frequently mentioned as cons.

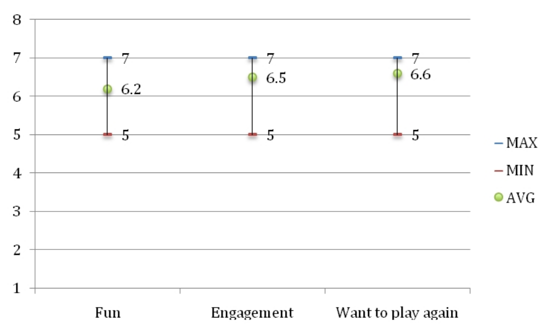

Figure 6: *Feedback Scores*

7.5.2 Feedback About the AR Interface

Interestingly, the handheld AR interface is related to both what the participants liked and disliked, as mentioned above. While people agree that the interface brings more fun to the game (4.3/5), the score for "the interface gave enough feedback" is just a little bit above neutral (3.3/5). Below are some more specific observations.

Handheld AR interface as the lens to the hybrid world: The handheld AR interface works like a lens onto the merged physical-digital game world. Like many games that have a fog-of-war effect, the players cannot see further than the spaces nearby. In AoD, the players need to physically move to explore the game space. Social interaction and collaboration is encouraged because of this, as discussed in the next section. However, this limited view also introduces some problems. For example, many players want to have a better sense of what is happening on the rest of the map. The lack of information about the enemy attack makes the game "hard to plan ahead" (G3-A, G2-B, G5-B).

Another issue is occlusion, meaning that the handheld is "blocking the way of the other player". The players worked out their own ways around this issue after they realized it would be a problem. They either reached out to point to the marker from another direction, or waited for the other person to finish the action first. Interestingly, in our previous study of Bragfish, we found that players intentionally blocked the way of each other to occupy the better position in the competitive play. In both cases, the players understand how the interface works, however, the design needs to consider these kinds of space-related interactions.

Emergent play: Trust, reference and communication: We found that the three teams (G1, G3, G4) who got the highest scores shared a similar kind of strategy. One player would be in charge of two of the paths (there were four paths in the game level we used), and guide the other player to the location where they need to perform an action. Through observation and recorded video, we found that this process could be broken down into three major components: trust, reference and communication. All of these are closely related to the design of the interface.

When players were using the handheld interface as a lens onto the game world, the limited area shown on the screen significantly affected game play. The strategy of moving around the whole map and making sure there are enough towers in the right places is inefficient with this limited view. The successful players realized that they could rely on the other player to cover half of the map and pay attention to only those places indicated by the other player. These groups mentioned that they needed to trust each other to make this strategy work (G1, G3, G4). As mentioned by G4, the immediate feedback assured by face-to-face interaction was necessary for creating trust between players who did not know each other.

In this game, how to refer to a certain location on the map became a key element for the collaborative process. One group found it hard to develop a successful reference protocol in the game (G6), while most other groups used pointing gestures assisted with language (G1-G5). Two groups also used physical tokens that they would place at specific locations to refer to later (G3, G4). The two basic protocols: pointing (directing-to) and placing (placing-for) can efficiently draw the attention of the other player. The tangible pieces facilitate using physical space as a reference system, and become an efficient form of communication mediation.

7.6 Social interaction

It is common to use games as ice-breaking media to bootstrap communication between strangers. This user study had some initial findings in the space of co-located cooperative computer games. In the questionnaire, the users reported that they feel comfortable playing with a stranger (4.2/5); and they have a neutral opinion about whether they think the play experience will be better with a friend (3.2/5). They enjoyed talking with the other player (4.4/5), while they believed that the other player was also willing to communicate (4.1/5). They preferred to play the game with a real person instead of a computer (4/5). The communication includes verbal conversation, hand gestures and body movements.

We also observed that all the groups started to talk more after the first round. As participants played longer, their conversation was initiated by and centered around the game play. In some cases, we saw that the two players did not talk at all for as long as several minutes during the first round (G2, G6); but they started talking after the first round finished and reported that face-to-face communication contributes to their game play in the interview. The conversations that happened between rounds were typically reflections on the previous game, the "theories" about which strategy may work better, and the plan for the next round. The conversation during the game was concentrated on passing knowledge about the game to the other player, informing the other player about the game status and asking for specific cooperation.

Prior research showed that co-located handheld gaming does not foster social interaction between players during gameplay [Szentgyorgyi et al. 2008]. The user study of AoD showed a good amount of player interaction that leveraged social cues, including non-verbal communication.

8 Conclusions and Future Work

We have discussed the potential and limitations of mobile AR games, and analyzed them in the context of a specific game design. AoD demonstrates specific mechanics that leverage and balance the game elements with the technology constraints. We also show that dynamic marker-tile-based map building and tangible interaction techniques, with either sketches or game tokens, allow AR games to be tightly integrated with the players perception of the world, creating the illusion that the game really "is" on the table top.

By giving each player their own view of the hybrid space, handheld AR games allow us to make use of the core features of traditional board games. Our study shows that AoD, with its tangible interaction and board game elements, is fun to play and encourages social interaction and communication between players. As AR technology and mobile phone hardware improves, we believe that AR board games will become increasingly popular.

As we move forward, we hope to leverage the rapidly improving sensing and display technology on mobile devices (e.g., natural feature tracking, multitouch screens and orientation sensors) to create even more compelling games. In the long run, we hope to explore the social potential of a much wider class of handheld AR games, especially those in large-scale outdoor environments. We believe that, with full mobility and immersion in the combined physical/virtual world, both tabletop and outdoor handheld AR games will transform the way people interact, collaborate and play games with each other.

Acknowledgements

We are grateful to Turner Broadcasting and Nokia for support, Imagination GmbH and Graz University of Technology for providing access to StbTracker, and Elements Interactive Mobile B.V. for access to EDGELIB. We would like to acknowledge Kris Kritmanorote and Iesha Hazel for game artwork and 3D models, and Timothy White and Michael Gorbsky for their design input.

References

BARKHUUS, L., CHALMERS, M., TENNENT, P., HALL, M., BELL, M., SHERWOOD, S., AND BROWN, B. 2005. Picking pockets on the lawn: The development of tactics and strategies in a mobile game. In *Proceedings of UbiComp 2005*, Springer, 358–374.

BILLINGHURST, M., IMAMOTO, K., KATO, H., AND TACHIBANA, K. 2000. Magic paddle: A tangible augmented reality interface for object manipulation. In *Proceedings of the IEEE and ACM International Symposium on Augmented Reality*, 111–119.

BJÖRK, S., FALK, J., HANSSON, R., AND LJUNGSTR, P. 2001. Pirates! using the physical world as a game board. In *Proceedings of Interact 2001*, 9–13.

BRADLEY, D., AND ROTH, G. 2007. Adaptive thresholding using the integral image. *Journal of Graphics Tools 12*, 2, 13–21.

HENRYSSON, A., BILLINGHURST, M., AND OLLILA, M. 2005. Virtual object manipulation using a mobile phone. In *ICAT '05: Proceedings of the 2005 International Conference on Augmented Tele-existence*, ACM, New York, NY, USA, 164–171.

HENRYSSON, A., BILLINGHURST, M., AND OLLILA, M. 2006. Ar tennis. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Sketches*, ACM, New York, NY, USA, 13.

HENRYSSON, A., MARSHALL, J., AND BILLINGHURST, M. 2007. Experiments in 3d interaction for mobile phone ar. In *GRAPHITE '07: Proceedings of the 5th International Conference on Computer graphics and interactive techniques in Australia and Southeast Asia*, ACM, New York, NY, USA, 187–194.

KLEIN, G., AND MURRAY, D. 2007. Parallel tracking and mapping for small ar workspaces. In *ISMAR '07: Proceedings of the 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality*, IEEE Computer Society, Washington, DC, USA, 1–10.

MANDRYK, R. L., AND MARANAN, D. S. 2002. False prophets: exploring hybrid board/video games. In *CHI '02: CHI '02 extended abstracts on Human factors in computing systems*, ACM, New York, NY, USA, 640–641.

MULLONI, A., WAGNER, D., AND SCHMALSTIEG, D. 2008. Mobility and social interaction as core gameplay elements in multiplayer augmented reality. In *DIMEA '08: Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts*, ACM, New York, NY, USA, 472–478.

PEITZ, J., BJÖRK, S., AND JÄPPINEN, A. 2006. Wizard's apprentice gameplay-oriented design of a computer-augmented board game. In *ACE '06: Proceedings of the 2006 ACM SIGCHI International Conference on Advances in Computer Entertainment technology*, ACM, New York, NY, USA, 79.

SZENTGYORGYI, C., TERRY, M., AND LANK, E. 2008. Renegade gaming: practices surrounding social use of the nintendo ds handheld gaming system. 1463–1472.

WAGNER, D., PINTARIC, T., AND LEDERMANN, FLORIANAND SCHMALSTIEG, D. 2005. Towards massively multiuser augmented reality on handheld devices. In *Proceedings of the Third International Conference on Pervasive Computing*, 208–219.

WAGNER, D., LANGLOTZ, T., AND SCHMALSTIEG, D. 2008. Robust and unobtrusive marker tracking on mobile phones. In *ISMAR 2008: Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality, 2008.*, 121–124.

WAGNER, D., REITMAYR, G., MULLONI, A., DRUMMOND, T., AND SCHMALSTIEG, D. 2008. Pose tracking from natural features on mobile phones. *ISMAR 2008: Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality, 2008.* (Sept.), 125–134.

XU, Y., GANDY, M., DEEN, S., SCHRANK, B., SPREEN, K., GORBSKY, M., WHITE, T., BARBA, E., RADU, I., BOLTER, J., AND MACINTYRE, B. 2008. Bragfish: exploring physical and social interaction in co-located handheld augmented reality games. In *ACE '08: Proceedings of the 2008 International Conference on Advances in Computer Entertainment Technology*, ACM, New York, NY, USA, 276–283.