

Project Zanzibar: A Portable and Flexible Tangible Interaction Platform

Nicolas Villar, Daniel Cletheroe, Greg Saul, Christian Holz, Tim Regan, Oscar Salandin, Misha Sra, Hui-Shyong Yeo, William Field, Haiyan Zhang

Microsoft Research

{nvillar, daclethe, grsaul, cholz, timregan, ossala, t-misra, t-huyeo, t-wife, hazhang}@microsoft.com




Figure 1. The Zanzibar Mat (b) tracks objects equipped with NFC tags (c), senses touch, detects hover gestures, and can be rolled up for portability (a). Different tag designs can be used to determine rotation (d), stack objects (e), provide I/O (f), and extend sensing outside the bounds of the Mat (g).

ABSTRACT

We present Project Zanzibar, a flexible mat that locates, uniquely identifies, and communicates with tangible objects placed on its surface, as well as senses a user's touch and hover gestures. We describe the underlying technical contributions: efficient and localised Near Field Communication (NFC) over a large surface area; object tracking combining NFC signal strength and capacitive footprint detection, and manufacturing techniques for a rollable device form-factor that enables portability, while providing a sizable interaction area when unrolled. In addition, we detail design patterns for tangibles of varying complexity and interactive capabilities, including the ability to sense orientation on the mat, harvest power, provide additional input and output, stack, or extend sensing outside the bounds of the mat. Capabilities and interaction modalities are illustrated with self-generated applications. Finally, we report on the experience of professional game developers building novel physical/digital experiences using the platform.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-5620-6/18/04...\$15.00

DOI: <https://doi.org/10.1145/3173574.3174089>

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords

Tangible user interfaces; near field communication; capacitive sensing; flexible devices; games; play; education.

INTRODUCTION

Tangible User Interfaces (TUIs) enable users to interact with digital content through direct manipulation of the physical environment and objects, and have been an area of active research for over two decades. Previous work has shown the application of tangible interfaces to areas such as gaming [30], productivity [42, 47] and education [51]. Studies have demonstrated how TUIs foster learning and understanding of interactions [40], improve collaboration between multiple users [4, 40], as well as increase creativity [31] and entertainment [18]. However, despite the extensive body of work, examples of TUIs in daily consumer and commercial use remain sparse.

Initially, the emergence of optical tabletop systems facilitated the development of tangible interfaces, such as the Microsoft Surface table and Reactable [18]. While users interacted with physical artifacts on and above the table's surface, the optical system below tracked the location and identity of objects, and displayed graphics via projection. During the past

decade, however, portable computing devices such as phones, tablets and laptops have taken over personal computing applications, including gaming, productivity, and education, which has shifted the context of use to mobile scenarios at home and out in the world [23]. Since tabletop systems require a stationary installation, they do not lend themselves to portable contexts, which more recently has limited their applicability to specialized scenarios, such as exhibits [15].

Portable touchscreen devices use capacitive sensing to detect touch input, which has been leveraged to support simple tangible interactions such as detecting toys on the screen [21]. Given that capacitive sensing on commodity devices is not designed to support tangible interaction, doing so can require workarounds such as necessitating actively powered objects [43]. Perhaps the biggest limitation of mobile devices for tangible interaction is the limited screen real-estate; though larger tablet screens measure up to 14 inches, reaching the limit of easy portability, they still remain comparably small for applications involving even a few tangible objects. Products have sought to address this using cameras for tracking [27], but these are prone to occlusion and illumination issues.

In this paper, we introduce Zanzibar, an easily portable user interface platform for tangible and direct touch interaction. Zanzibar achieves a portable form factor through a sensing surface made from layered flexible materials, allowing users to roll up the device, stow, and carry it. Rather than providing its own display, Zanzibar supports integrating existing devices for a range of interactive experiences that we demonstrate below, including audio-based apps, augmented-reality experiences, and scenarios that leverage the displays of existing devices.

THE ZANZIBAR PLATFORM

With the Zanzibar Platform, we aim to realize the promise of rich tangible interaction scenarios originally illustrated by systems such as optical tabletops, but in a portable and practical form factor. As shown in Figure 1, the *Zanzibar Mat* tracks a user’s touch and tagged objects placed on or near the Mat’s surface. The sensing area of the Mat, delimited by the hatched pattern in Figure 1b, is close to the size of an A3 sheet of paper (30 × 42 cm), and can be rolled into a compact form-factor for transport and storage (Figure 1a).

The Mat uniquely identifies and communicates with tagged objects through Near Field Communication (NFC), using a design that allows sections of the Mat to be selectively energized, and provides efficient performance across the entire sensing area. Simultaneously, the Mat senses touch and recognizes hover gestures through capacitive sensing. This unique integration of spatial NFC scanning and capacitive sensing makes the Zanzibar Mat distinct from previous approaches to provide accurate touch and object tracking that is robust to occlusion and lighting conditions.

To supplement the Mat, we have developed hardware design patterns for tangible objects that build on NFC to support interface elements of varying complexity and interactive capabilities, and which leverage and extend the capabilities of the Mat. An object becomes trackable by simply tagging it with a commodity NFC sticker of one of two standard sizes (Fig-

ure 1c) to give the object the distinctive capacitive footprint and unique NFC ID used by the tracking system. Optional metallic features around the tag enable the Mat to detect tag orientations (Figure 1d). In addition to commodity tags, we have developed custom tags that harvest power from the NFC field for additional input/output capabilities, such as detecting button presses and controlling LEDs or actuators in a tangible object (Figure 1f). Finally, we have designed stackable tags (Figure 1e), as well as tags that extend sensing outside the bounds of the Mat surface (Figure 1g).

To complement our hardware contribution, we have developed the Zanzibar Software Development Kit (SDK): libraries for communication with Zanzibar Mats, along with an Application Programming Interface (API). The API provides high-level interaction events (e.g., object or touch detection) and accepts control commands for active tags on the Mat (e.g., to toggle LEDs). Our SDK integrates with the Unity game engine, chosen for its cross-platform development capabilities, and because we consider physical/digital games to be a particularly rich application area for tangible user interfaces.

We validate our platform design and SDK by reporting on a deployment with two professional game development studios: the first developing an original physical/digital board game for Zanzibar, and the second adapting an existing game to work with the platform.

RELATED WORK

Zanzibar is related to tangible interaction in research, games, and education, building on previous touch technologies, object tracking, and interactive RFID techniques.

Tangibles in Games and Education

Tangible user interfaces (TUIs) use physical objects to make digital information directly manipulable with our hands and perceptible through physical embodiment [16]. Traditional boardgames have been played for centuries, providing challenges, skill learning, and social interaction. Digital games offer visual and audio feedback, save and resume options, and the ability to play remotely with others. Tangible interface designers have worked on educational games [7], and applied the Montessori methods of hands-on learning through physical manipulation [51], particularly for engaging children in playful learning [32].

Tangible play combines the tactile benefits of traditional gaming with those of digital games for enhancing learning outcomes. Papert’s theory of constructionism suggests that the best learning experiences occur when learners are actively engaged in building things [28]. Tangibles often include passive [31] or active [22] blocks, tiles [35], robots [37] as part of board games, tabletop games, or mobile device games. Common research themes relate to solo or collaborative learning through play (e.g., programming [44], electronics [2], math and science [9, 36]) and creativity (e.g., music [25] or 3D modeling [41]).

Tangibles in Research

Ishii et al.’s conceptual framework explores the possible types of coupling between material and virtual representations for

tangible interaction [17]. Graspable User Interfaces, precursors to tangible interfaces, described methods to directly control digital objects through physical handles [8]. In addition to the research efforts on tangible interaction, related projects have also examined the conceptual understandings of tangible interaction [1].

Tracking Tangibles, Touch Input, and Hover Gestures

Vision-based tracking: Systems have frequently tracked tangible objects using an overhead camera paired with a projector to provide a co-located display (e.g., [48]). Since camera-based tracking faces challenges from occlusion and lighting conditions, a reliable setup can become complicated, requiring additional tracking in the infrastructure [34]. Optical tabletops thus often mount cameras and projectors below the touch surface [6], tracking tangible objects using fiducial markers [18] or by analysing their footprints [6, 46]. While optical tabletops enable rich capabilities, their form factors is necessarily large and not easily portable.

Capacitive sensing: Capacitive sensors are low-cost, low-power, thin, and support multi-point sensing, making them prevalent in today’s touch-enabled devices. Grosse-Puppenthal et al.’s survey presents a comprehensive overview of capacitive sensing and its applications [10]. To track objects on touchscreens, related efforts have obtained the ‘raw’ sensor images on capacitive commodity sensors to detect the capacitive footprints of touches and objects in contact (e.g., [13, 14]) as well as used conductive markers attached to objects (e.g., [33, 49]). The challenge in tracking objects this way is the limited resolution of the sensor combined with the number of unique marker patterns for reliable tracking. Switching the capacitive sensing mode to self-capacitance, researchers have also used touchscreens to detect hover gestures (e.g., [11]).

Magnetic tracking: Sensetable [29] adapted the magnetic tracking system in Zowie Playsets [50] for tangible interaction, which used individually tuned resonant coils to track and differentiate objects. This technique is a low-cost and robust way of tracking unpowered objects, but is limited by the number of coils that can be uniquely identified. Others have explored the use of permanent magnets embedded in objects together with the magnetic sensors in phones to detect the location and orientation of objects placed over the device [3].

Interactive RFID Applications

Radio frequency identification (RFID) technologies have been used to associate digital content with physical objects on tabletops (e.g., [45]). RFID has also been used with optical tracking as a means to uniquely identify objects [26]. Work including WISP [5], RapID [39] and IDSense [20] has demonstrated the ability to track objects tagged with parasitically-powered UHF RFID tags at room-scale distances, but such systems require a relatively large and high-powered reader antenna to operate.

While Near Field Communication (NFC), a subset of RFID, is commonly used in mobile devices in applications such as contactless payments, it has also been adopted by the toys-to-life product category (e.g., Activision Skylanders [38], Nintendo Amiibo [24] and Lego Dimensions [19]). These products use NFC-tagged toy figurines to unlock digital content in



Figure 2. Zanzibar Mat components, from left to right: Stretchable NFC antenna array, flexible antenna interconnect, flexible capacitive electrode matrix, rigid PCBs housed in the baton enclosure.

videogames by placing them on a NFC reader “portal”. Interaction with the physical toy is limited to placing and removing figurines on a portal; most of the game is played through touchscreen or controller input.

IMPLEMENTATION

We now detail the implementation of the Zanzibar Mat, which includes: scalable NFC coverage that can localise tags; capacitive sensing to detect object footprints, touch, and hover input; a tracking system that fuses NFC-based and capacitive-based tracking information to report interaction events; and fabrication techniques for a rollable Mat surface.

In addition, we describe hardware design patterns for tangible objects in order to: track rotation of objects on the Mat, harvest power and use it to enable on-object input/output capabilities, extend sensing outside the bounds of the Mat surface, and detect the presence and order of stacked objects. Finally, we discuss the Zanzibar software development kit and the features it provides to tangible interface designers.

Zanzibar Mat

The Zanzibar Mat is designed for self-contained operation—all sensing and processing happens in the Mat itself; high-level interaction events are sent to connected devices via USB or Bluetooth.

We implemented the Mat as two physically distinct parts: a flexible area that provides an interaction surface of 30×42 cm (roughly ISO A3 format), attached to an extruded aluminium baton of $35 \times 2.4 \times 3.8$ cm. The flexible area is a laminate that includes two sensing layers: a NFC antenna array and a capacitive electrode matrix (Figure 2). The top and bottom external layers, as well as internal separation between sensing layers, are made from fabric materials. The baton houses circuit boards and electronic components for capacitive sensing and NFC subsystems, which include analogue front-ends, digitizers, and RF switching components. A central ARM Cortex-M4 100MHz 32-bit processor (ST Micro STM32F412) performs control, signal processing, tracking and communication tasks. Two rechargeable 400 mAh batteries ($3.5 \times 2.5 \times 0.5$ cm) power the system for over 2 hours of active user interaction at full frame rate.

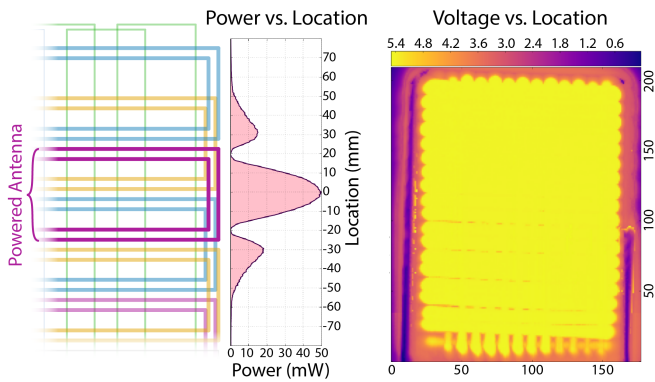


Figure 3. Coupling response for a single double-loop column antenna (left). Overlapping antennas minimize nulls and deliver consistent voltage to tags across the Mat (right).

Scalable and Localized NFC Coverage

The NFC subsystem is able to energize, localize and communicate with multiple tags up to 30 mm above the surface of the Mat. To do this efficiently, we implement a row/column NFC antenna pattern that extends across the flexible surface of the Mat (Figure 2). Each antenna is an electromagnetic coil that can be electronically switched into a single NFC transceiver (TI TRF7970A) to couple with tags along its length. The system has been optimized to work with the ISO14443 standard.

Every NFC tag has a globally unique identifier (UID). The NFC standard has an anti-collision system that allows multiple tag UIDs to be resolved by one reader, therefore each row or column may report the presence of multiple tags. Each antenna has a coupling response that is strongest when the tag is in between the conductors, drops to zero when the tag centre is over the conductor, then increases when it is outside, and finally decays again as shown in Figure 3 (left). To compensate for the nulls present on either side of each antenna, there is overlap between neighbouring antennas which allows tags to be powered consistently at any position (Figure 3, right).

By sequentially switching between rows and columns the system gathers information on the presence and signal strength of a tag on each row and column. This information is used to estimate a two-dimensional tag location. Depending on the size of a tag, the overlap and response of antennas, a tag can be detected by multiple rows or columns. The design ensures a low capacitance open circuit when an antenna is not active. This prevents overlapping coils from coupling with each other, which would waste energy and degrade localisation.

An alternative implementation could use tiled antennas, rather than rows and columns. Given small enough antennas, this would minimize the number of collisions and maximize tag coupling, but the number of individual antennas would increase proportionally with surface area and become impractical given to the need to route and switch each coil individually.

Capacitive Object Footprints, Touch and Hover

The capacitive subsystem uses an Atmel mXT2954T2 integrated circuit (IC), configured to output raw mutual capacitance deltas (as part of the chip’s debug reports present on




Figure 4. Capacitive sensor image, showing fingers touching, tags on playing cards, a custom NFC tag made of printed circuit board, and an untagged object.

commodity touchscreen chips [14]). We disabled all internal touch processing, signal smoothing, hysteresis, and background subtraction on the IC for maximum control. 41×58 drive and sense lines are used to cover the active sensing area in a typical touch-sensor diamond pattern (Figure 2). The capacitive sensor is made from copper electrodes, which results in a considerably stronger signal to noise ratio than ITO-based transparent sensors on touchscreens.

The IC outputs signed 8-bit capacitance values at 25 Hz. We separate positive and negative values to discriminate input from touch and metal elements, as touches will decrease the mutual capacitance between drive and sense lines while a dielectric or small ungrounded metal objects will increase it, such as an NFC tag’s coil. We refer to the resulting imprints as *capacitive footprints*. Figure 4 shows a capacitive image with a number fingers and objects on the Mat.

When switched into self-capacitance mode, the capacitive IC increases the sensing distance and reports one-dimensional row or column capacitance values with a 100 Hz frame rate. The system interlaces sequential row/column scans and estimates the centre of mass of a hand up to 10 cm above the surface of the Mat. This feature is fed into a recognizer to detect gestures such as swiping across the Mat or drawing shapes in the air. Setting the Mat into hover detection mode is mutually exclusive from multi-touch and object footprint detection. The Mat supports explicit switching through our API as well as operation in self-capacitance mode that switches to mutual-capacitance upon sensing approaching objects.

NFC and Capacitive Coexistence

For the touch and object tracking system to work, NFC communication and capacitive sensing must happen concurrently. Our localised NFC involves switching antennas, and the capacitive sensor operates by switching and charging the lines with a high edge rate. Due to these transient events, and the close proximity of the two sensing layers, there is signal interference caused by the NFC on the capacitive sensing, and vice-versa.

We implement a time-division multiplexing scheme to avoid interference by synchronizing capacitive scanning and NFC activity. In order to maintain a capacitive scanning rate close to its target of 25 frames per second, NFC operations are broken

down into smaller transactions that occur in between capacitive scans. The smallest possible NFC transaction involves energizing the field on a single antenna, determining any tag presence on an antenna, and performing a partial UID collision resolution operation. Operations that take an arbitrary amount of time to complete, such as resolving the UIDs of multiple tags on a single antenna, can operate across multiple capacitive frame periods without slowing down the capacitive frame rate.

The NFC layer sits below the capacitive layer, and acts as a partial ground reference for the capacitive electrodes. The capacitive electrodes do not significantly attenuate the NFC magnetic field, as each electrode appears as an open circuit that does not magnetically couple as a shorted turn. However, they do add significant capacitive loading across the NFC antennas to the extent that the antenna self-resonance drops below the target carrier frequency. This makes impedance matching difficult and forms unwanted resonators when antennas are inactive, coupling energy to other antennas. To mitigate this effect, the layers are separated by 0.8 mm, greatly reducing the capacitive loading on the NFC antennas.

Another challenge relates to the length and number of turns of NFC antennas. Multiple turns are desirable to improve magnetic coupling and therefore power transfer to the tags, which is particularly important for applications where tags harvest energy for purposes beyond NFC communication. However, increasing the turn count also increases the capacitive loading effect discussed above, as the longer wire length acts as a larger capacitive plate. We found a suitable compromise by using a single turn for the longer row antennas, and a double turn for the shorter column antennas. For power harvesting applications, tags are energized by the columns to achieve better performance (while still localized using both columns and rows).

Flexible and Rollable Layer Design

The flexible area of the Mat is designed to be *rollable*, so that it can be wrapped around the rigid baton for transport and storage. It is necessary for some layers to be *stretchable* as well as *flexible*, to accommodate elongation of one layer with respect to another as they are rolled. This would otherwise result in mechanical tension between layers, causing buckling or de-lamination.

The capacitive electrode matrix, shown in Figure 2, is fabricated as a flexible printed circuit (FPC). We have also evaluated a silver screen-printing process that has a lower manufacturing cost but slightly poorer performance as a capacitive sensor due to greater electrical trace resistance. It connects to the rigid PCB in the baton through FPC connectors that are designed as part of the layer. Both processes result in a material which is flexible, but not stretchable—this is not a problem as this layer does not need to stretch or compress, as long as other layers around it are able to do so.

The NFC antenna array layer sits 0.8 mm below the capacitive layer, separated by a stretchable cotton fabric layer and covered to give a total thickness of 2.5 mm. When rolled around the baton to a diameter of 50 mm, this layer needs to stretch by

13 mm or 3%. To achieve this, we manufacture the layer using Litz wire stitched onto a stretchable cotton fabric substrate. Litz wire is made up from multiple thin copper strands that are individually insulated by a thin enamel with a silk wrap to add to its tensile strength, which benefits the stitching process. The RF properties of Litz wire are also useful in reducing the high-frequency resistance at the NFC carrier frequency of 13.56 MHz. At this frequency, the skin effect is significant, and using thicker single-strand wires would result in current being concentrated on the outside of the conductor, effectively increasing its resistance and degrading power transfer to tags. To fabricate the layer, Litz wire of 30 strands of 40 μm diameter is wound into a bobbin of a consumer sewing machine, and a zigzag stretch-stitch is used in the direction of stretch (Figure 2, top-left). We have also successfully evaluated fabrication using industrial computer-controlled embroidery machines. The terminations of the row antennas are soldered directly to a PCB in the baton while the columns are soldered to a strip of FPC that runs along the top edge of the Mat (Figure 2) that includes local tuning and switching components.

The top external layer, which is the one a user comes into contact with during interaction, is made of a microfibre cloth that is a durable and comfortable touch surface. The bottom external layer, which experiences the most stretch when rolling, is made of an elastic polyurethane fabric. The layer stackup—including both active layers, internal separation and external layers—is laminated by coating adjacent layers with an elastic spray adhesive (3M DisplayMount), and pressure is applied through a heated press.

Tracking system

Our tracking system fuses data from the NFC and capacitive subsystems to generate high-level interaction events (e.g. *TouchDown*, or *TaggedObjectMoved*). While the capacitive subsystem reports footprints at higher resolution, precision, and faster update rates, the NFC subsystem reports the unique identities of tags even if they are not in direct contact with the Mat along with the coarse locations of tags. Zanzibar’s tracking system thus implements a hybrid tracking approach, building on previous approaches that use multiple modalities for simultaneous location and identity tracking (e.g., [12]).

Background NFC frames arrive around once a second, and include information about *all* of the tags detected on the Mat. Their timing varies depending on the number of tags. Requested sub-area scans take higher priority and return *localised NFC frames* that only includes tags in a specified location. In contrast, the capacitive subsystem reports *capacitive frames* to the tracker at a regular rate of 25 Hz.

Our tracking pipeline starts by separating the tracking of touches and object footprints. Touch events (*TouchDown*, *TouchMoved*, and *TouchUp*) are determined using capacitive data alone. To detect and track touches, we implement a tracking algorithm that segments footprints and links them to previously detected blobs. Due to our high capacitive signal quality, we segment touches using an aggressive threshold that discriminates even closely co-located touches. Depending on the presence, offset, or absence of touches, the tracking system

generates API events, which include touch locations, contact area size, and orientations.

For tagged object tracking, we threshold the capacitive frame to extract blobs that are caused by tags. The result serves as a mask to perform a convolution with predefined tag footprints for supported standard tag sizes to detect the likely presence of NFC tags. We then attempt to establish links between each local cluster of high correlations to previously detected clusters below a certain margin. For all resulting links, the Mat emits a *TaggedObjectMoved* event through the API.

For each newly detected cluster, we trigger a localised NFC scan in the region of the footprint to quickly identify the tag. If a returning *localised NFC frame* includes a previously untracked tag ID, we link the tag data with the capacitive footprint and emit a *TaggedObjectDown* event. After each *background NFC frame*, we check whether all currently tracked tags are included. If a tag is not reported after a number of frames, we emit a *TaggedObjectUp* event. If the tracker receives data about a tag without a corresponding capacitive footprint, we infer a hovering tag, and generate a *TaggedObjectDown* event using the coarse location of the tag based on the NFC data alone. We also use background NFC frames to emit *TaggedObjectMoved* events for tags that have no capacitive footprint linked. If the tracker later links such a tag with a capacitive footprint, locations will be reported using the capacitive footprint instead. *TaggedObject* events thus include the tag ID, location, and location confidence—‘coarse’ if tracked only by NFC or ‘fine’ if tracked by a capacitive footprint.

These implementation decisions address a number of challenges that we experienced in real-world use. One issue was merging and separating capacitive footprints that occur if a user moves two tagged playing cards on top of one another. Since this should not affect the application state, our tracker determines *TaggedObjectUp* and *TaggedObjectDown* events only after a full *background NFC frame* has arrived, resolving the ambiguity of whether one of those tags is no longer present. During this brief period of ambiguity, both tags’ locations are derived from a single capacitive footprint. If the cards split, the ensuing *localised NFC frame* resolves the previous ambiguity and correctly restores each association. A second issue we observed may result from a tentative player, holding a tagged object above the Mat, for example while pondering the next game move. Our tracker’s behaviour is consistent in that it reports the tag along with a *coarse* location, leaving it to app developers how the game should respond. Finally, given the update rate of the capacitive sensor, which is lower than current touchscreens, very fast motions may result in erroneous footprint-to-NFC links. In that case, a *background NFC frame* will correct erroneous links within a second.

Tracking Characterization

Accuracy

The accuracy of the tracking system is largely determined by the physical layout of the sensing elements. The NFC antenna array has a pitch of 20 mm, and the signal strength of a single tag measured on multiple antennas is used to estimate its location. Tag location accuracy depends on the number and placement of tags above the Mat, with a worst-case of 20 mm

and a maximum of 16 tags. The capacitive sensor electrodes have a pitch of 7 mm. This determines the sensing resolution of both touch points and tagged objects in direct contact with the surface of the Mat. Like current touchscreens, we interpolate between adjacent electrodes to achieve an accuracy of 2 mm. As a result, a tagged object is sensed with a 20 mm accuracy while held above the Mat, but once it is placed on the surface, its location is derived from its capacitive footprint with the same 2 mm accuracy as touch contacts. This is less accurate than on commodity touchscreens, which often use a 4 mm capacitive electrode pitch. Tag rotation accuracy is a function of the size and position of the additional metal feature with respect to the centre of the NFC tag. For the configuration shown in Figure 1d, the accuracy is approximately 5 degrees.

Update Rate and Latency

The capacitive sensor provides a raw capacitive image at 25 Hz, which determines the update rate of *TouchMoved* and *TaggedObjectMoved* events. The latency of touch events or movement of objects is a function of the signal processing of the capacitive image plus communication overhead, which is currently approximately 100 ms. The latency of detecting tag presence depends on the number of tags that are detected on nearby rows and columns and the scanning rate of the NFC subsystem, and is approximately 1 second in the average case.

False Positives and Negatives

An object presence is only reported once it has been verified through NFC, so there are no false positives on tagged object detection. Due to RF interference, a NFC transaction might fail and result in a false negative. This is mitigated with retries on the NFC scans, triggered by the presence of a candidate object observed in the capacitive image. False positives and negatives on touch are like those experienced on commodity touchscreens; in addition, large conductive objects placed on the Mat might introduce capacitive artefacts that can result in false touch-point positives.

Simultaneous Hover and Touch Limitations

When the hover gesture modality is enabled, self-capacitance readings approximate the position and distance of a hand above the Mat. This is time-multiplexed with mutual-capacitance scans to check whether the state of objects has changed. While an object is detected to be moving, or while a hand is close enough to the Mat to likely touch it, hover sensing is disabled. From an interaction perspective, it is not possible to move an object or touch while *simultaneously* performing hover gestures, but both actions can happen in succession, for example by moving a card on the Mat and performing a gesture above it to activate it. We have found that developers are able to understand and design successful hover gesture interactions in this way.

Design Patterns for Zanzibar Tangibles

We would like to facilitate experiences with Zanzibar that go beyond interaction with simple tokens, which afford only their presence, location and identity as input. To this end, we have developed design patterns and hardware reference designs that simplify the creation of tangible objects and range in cost, complexity, and interactive capabilities.

Simple NFC Sticker Tags

The simplest and lowest-cost way of allowing an object to be tracked by a Zanzibar Mat is to equip it with an off-the-shelf NFC sticker tag at its base. The tracking system supports the ISO14443A tags in 25 mm and 35 mm diameters (Figure 1c), which are available for as low as \$0.05. We have tested tags that use the NXP NTAG21X IC extensively; in addition to a globally unique ID, they provide user-rewritable memory that is accessible through the Zanzibar SDK to store metadata. The patterns described below build on these capabilities.

Rotation

The capacitive footprint of an NFC tag is rotationally symmetric, making it impossible to determine its rotation just from its footprint alone. Figure 1d shows the addition of a metal feature placed near a sticker tag, in this case a circle of copper tape. If present, our tracking system detects the feature to report the rotation of a tag on the Mat. This increases the footprint of the object, but less so than the alternative of using two separate sticker tags.

Power Harvesting and Input/Output

NFC tag ICs with power harvesting capabilities can rectify some of the RF signal to provide up to 10 mA of DC current. Figure 1f shows two different implementations we have developed to use harvested power for on-object input/output, which we refer to as *I/O Tags*. The first implementation is based on the Silicon Craft SIC4311 IC, on a custom PCB that equips the IC with a coil antenna and exposes seven digital input/output (I/O) and power pins. These pins can be used to drive and sense simple interface elements such as LEDs and switches, which can be easily read or set from the Zanzibar SDK.

This design is extended by a second implementation that includes a programmable low-power ARM Cortex-M0 microcontroller and improved power conditioning circuit to store charge in a super-capacitor for later use (for example, to drive an actuator that needs high instantaneous current). The microcontroller can be programmed to run arbitrary code and provides more advanced capabilities such as analogue inputs, pulse-width modulated outputs, and serial interfaces. This design can be used to equip a tangible object with sophisticated sensors or actuators that require more than binary I/O. It can be programmed to, for example, locally process sensor data and trigger an actuator in response, without communication with the Mat over the relatively slow NFC communication channel. If the power requirements exceed the harvesting abilities, both designs support the use of an additional battery.

Extending and Stacking

To enable a Zanzibar Mat to detect tagged objects at remote, elevated or offset locations from the surface, we developed techniques for *extending* and *stacking* tagged objects.

Extenders allow touch and NFC sensing to be extended away from the surface of the Mat (Figure 1g), up to a tested distance of 50 cm. Our implementation connects two tuned PCB coils through a length of Litz wire and includes an NFC IC with in-memory metadata that identifies it as an Extender. When one of the coils is placed on the Mat, it is identified and tracked as a regular tagged object. When an additional tag is placed

in proximity to the second coil, it couples magnetically and is detected as a second object on the Mat, sharing the same location as the Extender. Touching the remote coil capacitively couples down to the Mat, appearing as if the on-Mat coil has been touched directly. Multiple extenders may be placed next to each other to extend to different locations (for example, to detect when a toy figure has been placed in one of several rooms of a multi-storey toy house). Alternatively, it is possible to branch a single coil on the Mat to multiple off-Mat locations (in which case the presence of the figure in the house would be detected, but not which room it was in).

In order to detect stacking order of objects, we developed a system of *Stackers* (Figure 1e). This is done in a similar way to Extenders, by having pairs of coils couple between objects. In addition, Stackers use a controllable digital output connected to an RF switch to selectively connect or disconnect the second coil. The default state of the switch is open, so that only the bottom coil is energized when the Stacker is first placed on the Mat. The lowest Stacker in the stack is then detected, but anything above it remains uncoupled. Once the Mat has detected the first Stacker, the Mat instructs the Stacker to close its switch, energizing the top coil and allowing a second Stacker to be detected. This process is repeated until the whole stack is detected. One limitation of this design is that the coil at the base of a Stacker must always be the one placed on the Mat, or on another Stacker, and cannot be placed upside down. Our design supports up to four Stackers before coupling losses make detection unreliable.

Software Development Kit

The Zanzibar SDK supports the development of Zanzibar games and applications in the Unity and Visual Studio development environments. Supported Unity build targets include Windows Universal, Win32, and iOS apps, with possible support for Android. The SDK includes platform-specific libraries that abstract communication with Mats (either via USB, Bluetooth RFCOMM or Bluetooth Low Energy), and cross-platform C# APIs that provide: access to multiple Mats connected to a single device; collections of software objects that reflect the state, location, and rotation of touch points and tagged objects currently on a Mat; high-level interaction events (e.g. *TouchDown*, *TouchMoved*, *TaggedObjectUp*, *TaggedObjectMoved*, *HoverGesturePerformed*), and access to tag memory and control of I/O capabilities (for I/O Tags).

APPLICATIONS

In this section we illustrate the versatility of the Zanzibar platform through a number of game and educational application prototypes. The examples serve to demonstrate the interactive modalities of the Mat and tangible design patterns in use. Different scenarios demonstrate how the Mat can be used with mobile devices to complement tangible interactions with audio, display, or augmented reality capabilities.

Pirate Toybox

The notion of screen-time is a growing concern amongst parents as children gravitate away from toys and towards mobile devices at an increasingly younger age. Figure 5 shows the Pirates Toybox, a screenless playset where toys come to life



Figure 5. The Pirate Toybox screenless playset uses audio to complement physical play, which extends to the top of the tower. I/O Tags actuate and sense interaction with the chest and cannons.

on a Zanzibar Mat through sound effects and dynamic dialogue. The Toybox includes figurines for different characters (pirates and naval officers) that are implemented by augmenting existing toys with a simple NFC tag at their base. The combination and proximity of different characters on the Mat triggers different lines of dialogue—for example, two pirates will banter if placed close to one another, while a pirate will hail an officer with a friendly challenge. An Extender enables the Mat to detect when a character is placed on top of a toy tower, which itself can be placed anywhere on the Mat.

A toy treasure chest, equipped with an I/O Tag, senses when it is opened using a magnet and magnetic reed switch. Opening the chest in the presence of a pirate triggers sound effects and further dialogue from any characters on the Mat. A second I/O Tag is used to sense when a spring-loaded cannon is fired by manually pulling a latch. Movement of a character on the Mat immediately after the cannon is fired is used to infer that the character has been hit, causing them to shout out.

An alternative implementation uses an I/O Tag on the cannon to drive a small motor that actuates the firing latch. This can be triggered by a separate interaction, such as performing a hover gesture. This illustrates the ability to support remote play scenarios, where two children may be playing across networked Mats, and actions on one may trigger physical effects on the other.

Movie Maker

Movie Maker is a story authoring app aimed at children aged 5 to 12. In this age group, socio-dramatic or pretend play are prevalent play patterns, with children experimenting with role playing and ascribing unique personalities to their toys.

With the Movie Maker app, children can tell stories by manipulating toys and props on the Mat to control corresponding graphical avatars, previewing and capturing the result as a movie on the screen of a connected device (Figure 6). Camera and a light source objects allow the scene to be framed and lit. Additional props can be used to change the setting of a virtual scene. For example, a palm tree can change the virtual environment to a tropical setting.



Figure 6. Movie Maker can be used to create video stories by manipulating physical toys. Camera and light props are used to frame and light the scene, which is displayed on a connected device.

Toys are equipped with rotational markers to allow their orientation to be tracked in addition to their position on the Mat. The top surface of toys is augmented with conductive regions that extend down to their base, so that touching the toy appears as a touch point co-located with its position on the Mat. This is used to select a character while recording custom dialogue for it. The camera prop consists of two tags—one attached to the camera body, and a second to the lens that can be slid back and forth to adjust the zoom of the virtual camera. Light sources can be added to the scene by using tangible light props. The props have a physical ring that can be used to specify the colour and brightness of the light.

AR Tower Defence

Tower Defence is an augmented reality (AR) game that combines tangible input, unique object identification, precise object tracking, and multi-touch—all afforded by Zanzibar together with AR graphics from a Microsoft HoloLens device as shown in Figure 7 (left). Players set up defences and control characters on the Mat with tangible objects, and fight waves of oncoming holographic skeletons. A tangible stamp is used as a tool to place defensive towers, creating a tower hologram every time it is pressed against the Mat. Holograms are also used to overlay special effects and animations to tangible objects on the Mat, such as shields and arrows.

Alternatively, players can use a gesture to bring the game off the Mat and into the room at a life-size scale. In this mode the tangibles act as miniature proxies for holograms, allowing the player to direct the action of the AR characters via tactile and spatially multiplexed control.

The Zanzibar Mat communicates directly with the HoloLens over Bluetooth. In order to reconcile the Mat's coordinate system with its own, the HoloLens only needs to track the bounds of the Mat. In our implementation, we display a virtual bounding box that is manually aligned with the Mat at the beginning of the game, from which point the HoloLens tracker takes over. This could also be done automatically by having the HoloLens recognize visual features on the Mat.




Figure 7. The Tower Defence game combines tangible input with augmented reality holograms (left). Segmented E-Ink displays embedded in playing cards are driven by a power-harvesting I/O Tag (right).



Figure 8. The Letter Plates Montessori exercise provides feedback as words are formed and tactile letter-shapes are traced. Paper overlays on the Mat can provide additional context to the exercise.

E-Ink Battle Cards

Battle Cards is a technical demo where playing cards are embedded with LEDs and segmented E-Ink displays, driven by a power-harvesting I/O Tag (Figure 7, right). Placing pairs of cards on the Mat triggers sound effects, LEDs flash and the numeric value of the card is updated based on the result of the battle. The on-card displays provide an alternative way to communicate dynamic and digital information, instead of relying on the display of a device connected to the Mat. Because these displays are bistable, information displayed on them remains even after the card is lifted from the Mat. The prototype cards measure 3 mm in thickness, but could be reduced closer to 1 mm by using a custom PCB instead of the I/O Tag reference design shown in Figure 1f.

Montessori Exercises

The importance of multi-sensory learning is especially apparent in the Montessori method of teaching, where children use physical objects and self directed activities for all subjects of learning. In a manner similar to [51], we have taken some traditional Montessori exercises for young children and extended them with digital content and feedback using Zanzibar.

The exercises use *overlays*—Mat-sized sheets of paper with guides, graphics, or information printed on them—that can be used to decorate the Mat surface or provide additional context for a specific application. Each overlay is tagged, so that placing it on the Mat displays an associated exercise on a connected device. Touch and object detection still operate through the overlays.

The *Number Rods* exercise uses colour-coded rods to teach single digit addition and subtraction. Red or blue rods of different lengths are placed on a grid overlay. The app sets out problems to be solved, such as: “Given a blue number rod of length 4, which red rod do you need to get to a length of 9?” Experimentation is encouraged (for example, using two shorter rods in place of a longer one), and completion gradually increases the complexity of problems. *Fraction Plates* are circular plate segments, representing fractional numbers. The exercise allows a child to explore different ways to make whole numbers by combining segments of varying fractional

sizes. Finally, *Letter Plates* is an exercise that teaches spelling and familiarity with the alphabet (Figure 8). The materials for this exercise used in a typical Montessori classroom use textured letter shapes, so that a child can trace a letter with their finger while seeing and feeling its shape. In our version of the exercise, we create templates where each letter is laser-cut out of a cardboard tile. A child is asked to spell a word, and once this is completed they are asked to trace the shapes of the letters. The Mat follows the tracing of the letters through touch input, and provides feedback as the exercise is completed.

DEPLOYMENT

In addition to the self-generated applications described above, we report on the results of two engagements where professional game developer studios were given access to Zanzibar development kits. The kits included Mats, tags, and the SDK. We provided documentation, technical support, and feedback, while game design and implementation was carried out independently by the two teams.

Duel of the Captains

Duel Of the Captains (DotC) is a two player board game where players compete to destroy and defend their own spaceship. It is an original game created specifically for Zanzibar by Asobo Studio (Figure 9). DotC is played on a Mat with an overlay depicting two spaceships segmented into various functional areas (helm, shields, weapons, and teleportation bays). The overlay also includes virtual buttons to fire weapons and complete a turn, and areas where cards should be played. Each player has a set of figurines of characters with different specializations (pilot, weapons officer, shield operator), and a deck of cards with additional offensive and defensive bonuses. Every turn players can play a card from their hand, move a figurine to a bay to boost ship function, and choose whether to charge or fire their weapons. Once a turn has been completed, the effect of the player’s decisions and their outcomes are animated on-screen. Figurines use I/O Tags to drive LEDs embedded in their base that change colour to indicate if a figurine has been placed in a valid or invalid location. Cards are embedded with commodity 35 mm NFC sticker tags.



Figure 9. *Duel of the Captains* is a hybrid physical/digital board game, developed specifically for Zanzibar by an independent professional video game development studio. Two players battle against each other on a Mat using an overlay, figurines, cards and touch controls. The outcome of every turn is animated onscreen.

The game team consisted of six people, including a producer, designer, developers and graphic artists. Development of a fully working game was completed in three weeks, including Zanzibar integration. The issues faced by the team were less about making the technology work, and more about designing successful physical/digital interactions. For example, a point of discussion was around managing player attention between the physical objects and on-screen graphics: how to maintain feeling of playing with someone across a game board, while leveraging a proximate display to drive gameplay without dividing attention. This led to the turn-based design, where the focus remains on the board while each player makes their tactical decisions (helped by on-figurine LED placement feedback), and where only at the end of a turn are players rewarded with the animated outcome of the battle.

Rival Books of Aster

Rival Books of Aster is a popular strategic collectible card game (CCG) available for PC and iOS. Players build a collection from 200 unique digital cards and then craft custom decks to battle with in online matches that typically last from 5-10 minutes. Differently from other games of the same genre, the position of cards relative to each other has a significant impact on their effect. Stitch Media, the developer of Rival Books, adapted the game into a hybrid physical/digital version using Zanzibar. In the Zanzibar version the Mat acts as the player's control surface, with physical playing cards and spell-casting hand gestures.

The Stitch Media team was enthusiastic to combine the experience of collecting and manipulating well-made physical cards with a visually impressive digital game that crunches numbers and tracks complicated play states. As the computer handles the tracking of all game states, it eliminates the possibility of cheating, either intentionally or by mistake. The team felt that their audience would respond positively to the notion of the experience. As with DotC, difficult design choices mainly related to the distribution of game information between the physical objects and screen. For instance, it is considered a given that in a CCG (or any board game) the cards have game-

play statistics on them, even if those statistics change during the game. In a physical game, it is expected that players will track those changes via tokens and other mechanisms, while in a digital game, statistics are simply updated on-screen. In the Zanzibar version, the team found that leaving statistics off the physical cards altogether presented significant advantages: Displaying statistics only on the screen reduces confusion during gameplay by minimizing ambiguity. It also provides a mechanism to balance the game over time by changing the values or rules of any card digitally without making the physical artefact obsolete, future-proofing the physical product while allowing the game to evolve.

CONCLUSION

We presented Zanzibar, a platform for tangible interaction. The Zanzibar Mat combines capacitive sensing and NFC communication in a novel way, which allows multitouch and hover gesture input to coexist with physical object manipulation and control. The hardware design of the Mat addresses the challenge of providing NFC coverage over a large area, in a scalable and efficient manner. Our tracking system fuses capacitive sensing data and NFC to provide robust tracking of objects. We documented our fabrication techniques to make the Mat rollable: a form-factor that makes the device portable, while providing a surface area large enough to support interaction with multiple tangible objects. Our implementation builds on the existing NFC ecosystem and extends it with custom NFC tag designs that can be used to build tangible objects of varying complexity and interactive capabilities.

One characteristic of a useful platform is its ability to support varied and compelling applications. To this end, we have developed a set of scenarios that showcase different capabilities of the platform, illustrating the various interaction modalities it can support and highlighting how Zanzibar can be used together with existing computing devices (including augmented reality displays) to create a range of distinct physical/digital experiences. Our application examples have focused on play and education, two areas that we are passionate about and where the benefits of tangible interaction are well accepted. We have further validated our platform by deploying it with professional game developers, and learnt from their experience creating physical/digital game hybrids. We hope that our work motivates and enables others to continue to explore new ways to play, learn, and interact.

ACKNOWLEDGMENTS

We thank the following people for their help: John Friend, Phil Spencer, Ian Ormesher, Martin Grayson, Anja Thieme, Jim Holbery, Yuki Machida, Ruody Liu, Ken Sadahiro, Rob Connon, Rob Disano, Matt Hoesterey, Matt Turnbull, Ken Lobb, Richard Smith, Jonathan Lester, Matt Mickelson, Chris Bishop, Abi Sellen, Andy Wilson and the teams at Asobo Studio and Stitch Media.

REFERENCES

1. Alissa N. Antle. 2007. The CTI Framework: Informing the Design of Tangible Systems for Children. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction (TEI '07)*. ACM, New York, NY, USA, 195–202. DOI : <http://dx.doi.org/10.1145/1226969.1227010>
2. Ayah Bdeir. 2009. Electronics As Material: LittleBits. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction (TEI '09)*. ACM, New York, NY, USA, 397–400. DOI : <http://dx.doi.org/10.1145/1517664.1517743>
3. Andrea Bianchi and Ian Oakley. 2013. Designing Tangible Magnetic Accessories. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction (TEI '13)*. ACM, New York, NY, USA, 255–258. DOI : <http://dx.doi.org/10.1145/2460625.2460667>
4. Scott Brave, Hiroshi Ishii, and Andrew Dahley. 1998. Tangible Interfaces for Remote Collaboration and Communication. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work (CSCW '98)*. ACM, New York, NY, USA, 169–178. DOI : <http://dx.doi.org/10.1145/289444.289491>
5. Michael Buettner, Richa Prasad, Alanson Sample, Daniel Yeager, Ben Greenstein, Joshua R. Smith, and David Wetherall. 2008. RFID Sensor Networks with the Intel WISP. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems (SenSys '08)*. ACM, New York, NY, USA, 393–394. DOI : <http://dx.doi.org/10.1145/1460412.1460468>
6. Paul H. Dietz and Benjamin D. Eidelson. 2009. SurfaceWare: Dynamic Tagging for Microsoft Surface. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction (TEI '09)*. ACM, New York, NY, USA, 249–254. DOI : <http://dx.doi.org/10.1145/1517664.1517717>
7. Paul Dourish. 2004. *Where the action is: the foundations of embodied interaction*. MIT press.
8. George W. Fitzmaurice, Hiroshi Ishii, and William A. S. Buxton. 1995. Bricks: Laying the Foundations for Graspable User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '95)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 442–449. DOI : <http://dx.doi.org/10.1145/223904.223964>
9. Audrey Girouard, Erin Treacy Solovey, Leanne M. Hirshfield, Stacey Ecott, Orit Shaer, and Robert J. K. Jacob. 2007. Smart Blocks: A Tangible Mathematical Manipulative. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction (TEI '07)*. ACM, New York, NY, USA, 183–186. DOI : <http://dx.doi.org/10.1145/1226969.1227007>
10. Tobias Grosse-Puppenthal, Christian Holz, Gabe Cohn, Raphael Wimmer, Oskar Bechtold, Steve Hodges, Matthew S. Reynolds, and Joshua R. Smith. 2017. Finding Common Ground: A Survey of Capacitive Sensing in Human-Computer Interaction. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 3293–3315. DOI : <http://dx.doi.org/10.1145/3025453.3025808>
11. Ken Hinckley, Seongkook Heo, Michel Pahud, Christian Holz, Hrvoje Benko, Abigail Sellen, Richard Banks, Kenton O'Hara, Gavin Smyth, and William Buxton. 2016. Pre-Touch Sensing for Mobile Interaction. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2869–2881. DOI : <http://dx.doi.org/10.1145/2858036.2858095>
12. Christian Holz and Patrick Baudisch. 2013. Fibernio: A Touchscreen That Senses Fingerprints. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, New York, NY, USA, 41–50. DOI : <http://dx.doi.org/10.1145/2501988.2502021>
13. Christian Holz, Senaka Buthpitiya, and Marius Knaust. 2015. Bodyprint: Biometric User Identification on Mobile Devices Using the Capacitive Touchscreen to Scan Body Parts. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3011–3014. DOI : <http://dx.doi.org/10.1145/2702123.2702518>
14. Christian Holz and Marius Knaust. 2015. Biometric Touch Sensing: Seamlessly Augmenting Each Touch with Continuous Authentication. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*. ACM, New York, NY, USA, 303–312. DOI : <http://dx.doi.org/10.1145/2807442.2807458>
15. Michael S. Horn, Erin Treacy Solovey, and Robert J. K. Jacob. 2008. Tangible Programming and Informal Science Learning: Making TUIs Work for Museums. In *Proceedings of the 7th International Conference on Interaction Design and Children (IDC '08)*. ACM, New York, NY, USA, 194–201. DOI : <http://dx.doi.org/10.1145/1463689.1463756>
16. Hiroshi Ishii. 2008. The Tangible User Interface and Its Evolution. *Commun. ACM* 51, 6 (June 2008), 32–36. DOI : <http://dx.doi.org/10.1145/1349026.1349034>
17. Hiroshi Ishii and Brygg Ullmer. 1997. Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '97)*. ACM, New York, NY, USA, 234–241. DOI : <http://dx.doi.org/10.1145/258549.258715>
18. Sergi Jordà, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner. 2007. The reacTable: Exploring the Synergy Between Live Music Performance and Tabletop Tangible Interfaces. In *Proceedings of the 1st International Conference on Tangible and Embedded*

- Interaction (TEI '07)*. ACM, New York, NY, USA, 139–146. DOI: <http://dx.doi.org/10.1145/1226969.1226998>
19. Lego. 2017. Lego Dimensions. Website. (18 Sep 2017). Retrieved September 17, 2017 from <https://www.lego.com/en-us/dimensions/products>.
 20. Hanchuan Li, Can Ye, and Alanson P. Sample. 2015. IDSense: A Human Object Interaction Detection System Based on Passive UHF RFID. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2555–2564. DOI: <http://dx.doi.org/10.1145/2702123.2702178>
 21. Mattel. 2017. Mattel Brings Powerhouse Brands to the iPad With Apptivity Play. Website. (11 Sep 2017). Retrieved September 11, 2017 from <http://news.mattel.com/news/mattel-brings-powerhouse-brands-to-the-ipad-with-apptivityTM-play>.
 22. David Merrill, Jeevan Kalanithi, and Pattie Maes. 2007. Siftables: towards sensor network user interfaces. In *Proceedings of the 1st international conference on Tangible and embedded interaction*. ACM, 75–78.
 23. Hendrik Müller, Jennifer Gove, and John Webb. 2012. Understanding Tablet Use: A Multi-method Exploration. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services (MobileHCI '12)*. ACM, New York, NY, USA, 1–10. DOI: <http://dx.doi.org/10.1145/2371574.2371576>
 24. Nintendo. 2017. Nintendo Amiibo. Website. (18 Sep 2017). Retrieved September 17, 2017 from <http://www.nintendo.com/amiibo>.
 25. Uwe Oestermeier, Philipp Mock, Jörg Edelmann, and Peter Gerjets. 2015. LEGO Music: Learning Composition with Bricks. In *Proceedings of the 14th International Conference on Interaction Design and Children (IDC '15)*. ACM, New York, NY, USA, 283–286. DOI: <http://dx.doi.org/10.1145/2771839.2771897>
 26. Alex Olwal and Andrew D. Wilson. 2008. SurfaceFusion: Unobtrusive Tracking of Everyday Objects in Tangible User Interfaces. In *Proceedings of Graphics Interface 2008 (GI '08)*. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, 235–242. <http://dl.acm.org/citation.cfm?id=1375714.1375754>
 27. Osmo. 2017. Osmo. Website. (11 Sep 2017). Retrieved September 11, 2017 from <https://www.playosmo.com>.
 28. Seymour Papert. 1980. *Mindstorms: Children, computers, and powerful ideas*. Basic Books, Inc.
 29. James Patten, Hiroshi Ishii, Jim Hines, and Gian Pangaro. 2001. Sensetable: A Wireless Object Tracking Platform for Tangible User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*. ACM, New York, NY, USA, 253–260. DOI: <http://dx.doi.org/10.1145/365024.365112>
 30. Clément Pillias, Raphaël Robert-Bouchard, and Guillaume Levieux. 2014. Designing Tangible Video Games: Lessons Learned from the Sifteo Cubes. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3163–3166. DOI: <http://dx.doi.org/10.1145/2556288.2556991>
 31. Ben Piper, Carlo Ratti, and Hiroshi Ishii. 2002. Illuminating Clay: A 3-D Tangible Interface for Landscape Analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. ACM, New York, NY, USA, 355–362. DOI: <http://dx.doi.org/10.1145/503376.503439>
 32. S Price, Y Rogers, M Scaife, D Stanton, and H Neale. 2003. Using ‘tangibles’ to promote novel forms of playful learning. *Interacting with Computers* 15, 2 (2003), 169–185. DOI: [http://dx.doi.org/10.1016/S0953-5438\(03\)00006-7](http://dx.doi.org/10.1016/S0953-5438(03)00006-7)
 33. Jun Rekimoto. 2002. SmartSkin: An Infrastructure for Freehand Manipulation on Interactive Surfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. ACM, New York, NY, USA, 113–120. DOI: <http://dx.doi.org/10.1145/503376.503397>
 34. Jun Rekimoto and Masanori Saitoh. 1999. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, 378–385. DOI: <http://dx.doi.org/10.1145/302979.303113>
 35. Jun Rekimoto, Brygg Ullmer, and Haruo Oba. 2001. DataTiles: A Modular Platform for Mixed Physical and Graphical Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*. ACM, New York, NY, USA, 269–276. DOI: <http://dx.doi.org/10.1145/365024.365115>
 36. Eric Schweikardt and Mark D. Gross. 2006. roBlocks: A Robotic Construction Kit for Mathematics and Science Education. In *Proceedings of the 8th International Conference on Multimodal Interfaces (ICMI '06)*. ACM, New York, NY, USA, 72–75. DOI: <http://dx.doi.org/10.1145/1180995.1181010>
 37. Arnan Sipitakiat and Nusarin Nusen. 2012. Robo-Blocks: Designing Debugging Abilities in a Tangible Programming System for Early Primary School Children. In *Proceedings of the 11th International Conference on Interaction Design and Children (IDC '12)*. ACM, New York, NY, USA, 98–105. DOI: <http://dx.doi.org/10.1145/2307096.2307108>
 38. Skylanders. 2017. Skylanders Video Game. Website. (18 Sep 2017). Retrieved September 17, 2017 from <https://www.skylanders.com/>.

39. Andrew Spielberg, Alanson Sample, Scott E. Hudson, Jennifer Mankoff, and James McCann. 2016. RapID: A Framework for Fabricating Low-Latency Interactive Objects with RFID Tags. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5897–5908. DOI : <http://dx.doi.org/10.1145/2858036.2858243>
40. Hideyuki Suzuki and Hiroshi Kato. 1995. Interaction-level Support for Collaborative Learning: AlgoBlock—An Open Programming Language. In *The First International Conference on Computer Support for Collaborative Learning (CSCL '95)*. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 349–355. DOI : <http://dx.doi.org/10.3115/222020.222828>
41. Paula Te. 2015. TADCAD: A Tangible and Gestural 3D Modeling & Printing Platform for Building Creativity. In *Proceedings of the 14th International Conference on Interaction Design and Children (IDC '15)*. ACM, New York, NY, USA, 406–409. DOI : <http://dx.doi.org/10.1145/2771839.2771865>
42. John Underkoffler and Hiroshi Ishii. 1999. Urp: A Luminous-tangible Workbench for Urban Planning and Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, 386–393. DOI : <http://dx.doi.org/10.1145/302979.303114>
43. Simon Voelker, Christian Cherek, Jan Thar, Thorsten Karrer, Christian Thoresen, Kjell Ivar Øvergård, and Jan Borchers. 2015. PERCs: Persistently Trackable Tangibles on Capacitive Multi-Touch Displays. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*. ACM, New York, NY, USA, 351–356. DOI : <http://dx.doi.org/10.1145/2807442.2807466>
44. Danli Wang, Cheng Zhang, and Hongan Wang. 2011. T-Maze: A Tangible Programming Tool for Children. In *Proceedings of the 10th International Conference on Interaction Design and Children (IDC '11)*. ACM, New York, NY, USA, 127–135. DOI : <http://dx.doi.org/10.1145/1999030.1999045>
45. Roy Want, Kenneth P. Fishkin, Anuj Gujar, and Beverly L. Harrison. 1999. Bridging Physical and Virtual Worlds with Electronic Tags. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, 370–377. DOI : <http://dx.doi.org/10.1145/302979.303111>
46. Malte Weiss, Julie Wagner, Yvonne Jansen, Roger Jennings, Ramsin Khoshabeh, James D. Hollan, and Jan Borchers. 2009. SLAP Widgets: Bridging the Gap Between Virtual and Physical Controls on Tabletops. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 481–490. DOI : <http://dx.doi.org/10.1145/1518701.1518779>
47. Pierre Wellner. 1991. The DigitalDesk Calculator: Tangible Manipulation on a Desk Top Display. In *Proceedings of the 4th Annual ACM Symposium on User Interface Software and Technology (UIST '91)*. ACM, New York, NY, USA, 27–33. DOI : <http://dx.doi.org/10.1145/120782.120785>
48. Andrew D. Wilson. 2005. PlayAnywhere: A Compact Interactive Tabletop Projection-vision System. In *Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology (UIST '05)*. ACM, New York, NY, USA, 83–92. DOI : <http://dx.doi.org/10.1145/1095034.1095047>
49. Neng-Hao Yu, Li-Wei Chan, Seng Yong Lau, Sung-Sheng Tsai, I-Chun Hsiao, Dian-Je Tsai, Fang-I Hsiao, Lung-Pan Cheng, Mike Chen, Polly Huang, and Yi-Ping Hung. 2011. TUIC: Enabling Tangible Interaction on Capacitive Multi-touch Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 2995–3004. DOI : <http://dx.doi.org/10.1145/1978942.1979386>
50. Zowie. 2017. Zowie Playsets. Website. (19 Sep 2017). Retrieved September 19, 2017 from <http://www.piernot.com/proj/zowie/>.
51. Oren Zuckerman, Saeed Arida, and Mitchel Resnick. 2005. Extending Tangible Interfaces for Education: Digital Montessori-inspired Manipulatives. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. ACM, New York, NY, USA, 859–868. DOI : <http://dx.doi.org/10.1145/1054972.1055093>