

BookMark: Appropriating Existing Infrastructure to Facilitate Scalable Indoor Navigation

Jennifer Pearson, Simon Robinson, Matt Jones

FIT Lab, Computer Science Department, Swansea University, Swansea, SA2 8PP, UK

Abstract

This paper describes an approach to using physical-digital appropriation for navigation, piggybacking off the humble, ubiquitous barcode to facilitate cheap, scalable indoor wayfinding. We illustrate the technique by describing a prototype interface for navigating to specific books within a library. Our design—BookMark—provides visitors to the library with a detailed map to any desired book by simply scanning the barcode on the back of any other book in the library. After describing in detail how our technique is achieved, we move on to show its effectiveness via an in-situ experiment in which we compare our design to standard methods of library navigation. We then present the results of a longitudinal evaluation where we deploy the BookMark application to library visitors for a period of 24 months. We conclude with a discussion of our overall results, what our design means for other pervasive infrastructures, and how best to design for the future of physical-digital appropriation.

Keywords: Appropriation, indoor navigation, barcodes, books, libraries

1. Introduction

With the rise of smartphones and cheap sensors, GPS-based navigation systems have expanded to become the standard routefinding option for drivers and pedestrians alike. In well-mapped environments, then, outdoor navigation is often considered a solved problem. Indoor navigation, however, is more complex – issues such as blocked or reflected signals, and error ranges that are often larger than the indoor spaces being navigated make GPS-based approaches unsuitable for these environments. Instead, indoor navigation most commonly makes use of dead reckoning, beacon- or sensor-based approaches. Each of these techniques have their relative pitfalls – dead-reckoning approaches, for example, suffer from accumulated error over time, whereas other techniques such as beacon- or sensor- based approaches require additional infrastructure.

In this work we aimed to develop a method for indoor navigation that does not require additional infrastructure, and instead uses existing markers and waypoints in the environment. Our focus was on libraries, where visitors require a fine level of granularity of positioning in order to help them find a specific book within collections of many thousands of items. There have been many attempts to streamline the process of locating books within physical libraries, but, as reviewed by Walsh (2011), these usually require some sort of specialist hardware or additional infrastructure within the library itself. Instead, we propose

BookMark: a simple, cheap, scalable solution that makes use of existing infrastructure to facilitate precise navigation to specific books within a library.

Our technique relies on the ubiquity of the digital barcode – an easily recognisable marker that is visible on almost all consumer-level products, including the majority of printed books. Current uses of barcodes are rather unimaginative, and there are very few situations when these markers are used for anything other than their intended purpose. Researchers have previously tried to address this by making digital markers more appealing (e.g., Meese et al. 2013), using attractiveness to hide mundanity. In our view, though, the beauty of barcodes is not in their aesthetics, but in what their visibility and ubiquity affords. We take advantage of these pervasive codes by using them to determine where a user is standing in the physical space. By scanning the barcode on the back of any nearby book, BookMark can determine which book the user is holding, which shelf it was on, and, consequently, where they are positioned in the library. Using this information, and the details of the book they wish to locate, BookMark is then able to suggest a route from one location to the other, guiding the user to the correct shelf.

Other than a cameraphone, our approach requires no specialist hardware on the user side. Despite this, it is able to offer a very fine level of granularity, leading library visitors to the exact shelf containing the book they require. It also requires no additional infrastructure within the library itself, as the only information used for positioning is already printed on the back of most recent books. As a result, the approach is flexible, allowing librarians to

Email addresses: j.pearson@swansea.ac.uk (Jennifer Pearson),
s.n.w.robinson@swansea.ac.uk (Simon Robinson),
matt.jones@swansea.ac.uk (Matt Jones)

automatically update groups of book positions when shelf relocations occur or when new books are added.

The general technique is particularly beneficial in locations such as libraries, where large collections of sorted or categorised objects need to be found by users. In addition, however, the barcode appropriation technique has potential in other situations with varying levels of granularity. For example, supermarket items could easily be appropriated in this manner, directing customers perhaps not to a specific product, but to a category of products (e.g., the cereal section). Similarly, department stores could make use of the technique but with a far coarser level of detail, directing users to entire departments (e.g., the shoe department).

We begin this article with a literature review on the topic of indoor navigation as well as related approaches that have been used to locate books within physical libraries. We then present a study of current library use, including interviews with both borrowers and librarians at Swansea University’s Library and Information Centre. Following this, we discuss the background of our partner library, including how books are currently located and organised, before moving on to describe the BookMark prototype in detail. We then discuss two user evaluations of our design, including a direct comparison with existing item location methods (previously described in [Robinson et al. 2014](#)), and a longitudinal study of the prototype in situ for a period of 24 months. We end with discussion and conclusions about the benefits of the work.

2. Background

While outdoor navigation is largely seen as a solved problem (at least in well-mapped and GPS-covered areas), navigation within buildings is usually far less advanced. In indoor environments, walls and other obstacles block GPS signals, and layouts can be particularly complex, with multiple levels and many alternative route options. There has been a large amount of previous research into ways to support navigation in these contexts, however, and in this section we broadly review these, with a focus on those that are most closely related to our work.

2.1. Indoor navigation

Previous indoor navigation designs can be broadly grouped into three categories: dead-reckoning (e.g., [Foxlin 2005](#); [Li et al. 2012](#)); sensor-based (e.g., [Hub et al. 2004](#)); and, beacon-based (e.g., [Chawathe 2008](#); [Buchanan 2010](#)). Each of these solutions have their own trade-offs as summarised by [Fallah et al. \(2012, 2013\)](#). Most notably, dead-reckoning techniques degrade in accuracy over time as errors accumulate; beacon-based versions often involve large-scale augmentation of the physical environment; and, sensor-based approaches require considerable computational power and custom hardware.

2.1.1. Dead-reckoning and sensor-based approaches

Dead-reckoning approaches estimate a user’s position based on a previously known location. Using data from accelerometers, magnetometers, gyroscopes or other inertial sensors, this technique infers the user’s movement and predicts their current location based on an offset from an initial reading typically retrieved via GPS or, say, a known building entry point. Unfortunately, however, positioning accuracy using this technique degrades over time as inevitable minor errors from sensor data are compounded as the user moves.

Sensor-based approaches, meanwhile, employ detection of environmental features via cameras, laser scanners or other sensors in order to estimate the user’s current location. Key to these methods is a comprehensive pre-generated record of the indoor environment (whether 3D-modelled or photographically recorded) that current positions can be matched against.

An in-depth survey of dead-reckoning or sensor-based techniques is beyond the scope of this work. [Harle \(2013\)](#), however, provides a comprehensive exploration of research in this area, including classification into a broad taxonomy and a discussion of the benefits and tradeoffs of various approaches. [Hu \(2015\)](#) provides a similar review of indoor localisation methods that use visual sensors, with a particular focus on techniques for visually-impaired users.

In our view, the key similarity between dead-reckoning approaches and our work is that the initial set up cost is low, as no additional infrastructure is required. Our technique differs completely, however, in its approach to navigation – our method does not attempt to track the user while they are on their path, and indeed the system has no knowledge of where the user is during the majority of the navigation task. Instead, our approach relies on knowledge of the user’s starting point and their destination, assuming that they are able to follow the map, and will remain on the suggested route. If at any point the user diverges from the map or gets lost, they need only to scan any other nearby book to update their position with the system, automatically retrieving an updated map.

We do not see sensor-based indoor navigation as being closely related to our work due to its reliance on highly accurate models of the user’s environment. While many approaches in this category use a combination of image or distance sensors and inertial measurements in order to lessen the need for model accuracy (or indeed the considerable computational power and custom hardware that can be required to achieve accurate positioning), visual or topographic matching of the user’s position to a known point on a pre-generated model is at the core of this class of methods. Our approach does use a camera to scan ISBN barcodes; however, this can be seen as purely for convenience – our approach relies on the number encoded in the barcode and its known ordering within the collection of items, so is hence far more similar to beacon-based approaches.

2.2. Beacon-based approaches

Beacon-based navigation systems rely on the presence of physical identifiers (i.e., beacons) within the navigable space that can be recognised and used as positional aids. In some cases these may be added specifically for the navigation task at hand (e.g., RFID tags or fiducial markers), while other options more aligned with our approach take advantage of existing infrastructure such as WiFi signals in order to triangulate or estimate positions. While many indoor navigation systems use a combination of several methods (often including inertial or sensor-based techniques in addition to beacons), broadly speaking this class of navigation systems is the most closely related to our work.

Turning first to manually-placed beacons, previous research has explored the use of barcodes, QR codes and other scannable or detectable markers to provide navigation support. Mulloni et al. (2009), for example, describe a navigation system using custom scannable codes placed at specific points around an indoor space. Their system was used as a conference guide, requiring 37 of its markers to be installed in a 100 m × 200 m area to provide directions to key locations. Butz et al. (2001) chose instead to place infrared transmitters at key points throughout a building, broadcasting navigation directions to users who passed underneath. A similar system was deployed by Ciavarella and Paternò (2004) in a museum context, providing exhibit-specific information, while Chawathe (2008) and Wang et al. (2013) have demonstrated the use of Bluetooth-based beacons for indoor positioning. Other schemes have been proposed to combine both beacon-based and dead-reckoning techniques to improve accuracy. For example, Serra et al. (2010) describe a QR code-based system, using marker scans to precisely locate waypoints, and commercial apps such as PinPoint¹ have used systems of QR codes and NFC tags to locate and guide users (e.g., to find cars in multi-level car parks). Related approaches include that of Coughlan et al. (2006), whose technique used coloured tags for long-range phone camera-based localisation, aimed at people with vision problems. RFID tags have also been used to locate users within indoor spaces, such as in the work of Mantyjarvi et al. (2006) whose system was used for identifying art works in museums. Similarly, Retscher and Fu (2010) used RFID tags, but also incorporated signal strength measurements and inertial sensing to achieve more accurate positioning.

In general, beacon-driven designs such as these require a precisely-registered network of *additional* codes that must be added to objects and walls around the building, or suffer from existing problems such as dead-reckoning error accumulation. One type of beacon-based approach—WLAN triangulation, as demonstrated by Bahl and Padmanabhan (2000)—piggybacks on WiFi signals to estimate location, as we piggyback on the much lower-level technology of barcodes. Some triangulation techniques can be

very precise—Woodman and Harle (2008) achieve 1 m accuracy, for example—and there have also been commercial products using this method.² Bahl and Padmanabhan’s approach used a database of locations and signals to enhance standard triangulation approaches and improve resistance to signal noise via ‘fingerprinting’ user locations. More recently, the same technique has been used with Bluetooth beacons by Ahmetovic et al. (2016) to provide precise navigation assistance for those with visual impairments. However, such triangulation schemes depend on dense and universal coverage, external power, and highly-accurate location maps. In practice, location accuracy with real systems in this sort of context (e.g., libraries or supermarkets) is likely to be around only 10 m to 20 m.³

More similar to our approach is that of Nicholson et al. (2009), who used shelf-edge barcode labels in a small supermarket to help blind users navigate to product areas via verbal directions and waypoints. This particular method is largely similar to the way that large warehouses often store items. For example, Amazon use a ‘chaotic storage’ system (Eostein, 2012), which places items not in order of their category, but rather wherever there is free space available. These approaches rely solely on barcodes to function, as in our design, but, crucially, require a precise and curated mapping of individual items to positions. In our design we take advantage of the natural sorting and grouping of existing collections to minimise this effort. As a result, our approach provides accurate directions to endpoints, but does not update instructions as users move. Previous work exploring navigation user experiences (Brush et al., 2010) has demonstrated that route maps such as those we provide can be useful even when not actively being followed or refreshed to locate the user, while Taher and Cheverst (2011)’s studies of user preferences for indoor navigation found that mobile navigation options should offer a high degree of customisation, allowing users to update their route. We provide this flexibility by allowing users to scan any nearby book’s barcode to update the route map and receive new directions from their current position.

2.3. Library navigation

Libraries know where their books are located as a result of their cataloguing system. To organise the collection, the standard approach is to use the Dewey Decimal or Library of Congress Classification, and store items in a sorted manner throughout the building. Modern libraries usually also have electronic catalogues with records of where each book or resource is located, with precision often at least to the level of a range of shelves on a particular floor. Searching for a title returns the catalogue identifier, and a reader can then use this to find the resource within the ordered collection.

Indoor mobile location-based systems have previously been trialled in libraries to simplify the process of finding

¹<http://pinpoint-navigation.com>

²See, e.g., Aisle411: <http://aisle411.com>

³See, e.g.: <http://skyhookwireless.com> estimated accuracy.

an item from its catalogue number. In this work we demonstrate the use of ISBN barcodes already on library items as navigation signposts, but previous library systems have typically used WiFi appropriation or other beacon-based designs. SmartLibrary (Aittola et al., 2003) and the Newman Project, (Sciacchitano et al., 2006), for example, were early navigation designs that used WiFi triangulation in order to locate users within the library building. Others have proposed navigation systems that require significant additional infrastructure to be added to libraries in order to support navigation. Satpathy and Mathew (2006), for instance, created an RFID-based location system which involved proximity sensors on shelves, books and ID cards, and LEDs on shelf edges to indicate the position of the item in question. Buchanan (2010) employed Bluetooth beacons in conjunction with RFID tags in a similar manner, while Watanabe et al. (2010)’s LiNS system used various types of location sensors in order to detect the user’s proximity to nearby shelves. Finally, Dekel et al. (2011) took a slightly different approach, in which book barcodes could be scanned in order to look up further detail or sign out the item, while QR codes were positioned around the library to be scanned for navigation support.

Walsh (2011) reviews these and other common location technologies that have been used to support library navigation. It is clear that library navigation is a common and much-researched problem, but to our knowledge there are no previous examples of using existing markers and collection ordering information together to provide navigation support. Instead, previous designs have manually generated maps of book locations, or added further markers in order to support navigation. Our barcode appropriation design, then, offers a simpler approach, piggybacking on existing infrastructure and providing a precise enough level of detail to locate individual items in a large library.

3. Current library use study

In order to better understand existing library use, we conducted a series of interviews with 32 library visitors (17 male; 15 female) between the ages of 18 and 65. In addition, we interviewed two librarians to explore their thoughts about our proposed system, and their current techniques for helping library users find items. We also collected examples of users’ current strategies for locating items in order to inform our design.

3.1. Procedure

Participants were met in the library and took part in a short semi-structured interview. We aimed to gather a range of perspectives, so recruited users with various visiting habits. After an IRB-approved ethical consent process, the interview began by asking participants for their perspectives on the ease of use of current library book finding methods on a Likert-like scale of 1 to 7 (1 difficult; 7 easy). We then asked participants to estimate how long they spent

looking for items, and how many items they borrowed on a typical visit. Participants then discussed how they personally approached the task of finding library items, including how they found the correct section and shelf, and how they remembered the details of the item being found. Interviews took between 5 and 10 minutes each and participants were given a gift voucher as incentive for participation.

3.2. Results

Of the 32 participants, 25% visited the library at least once a day, 22% visited around once a week, 34% visited once a month and 19% visited around once a year. When asked how easy or difficult it was to find a book’s location within the library, the average score given was 4.0 (s.d. = 1.23, median = 4). The average of participants’ estimates of the number of books taken out on a typical visit was three (s.d. = 3, median = 2.8), with a minimum of one and maximum of twelve. When asked to estimate how long it takes them to find a book on a typical visit to the library, the average time given was 9 min (s.d. = 4 min 54 s, median = 7 min 30 s). While these results are clearly only broad estimations on participants’ behalf,⁴ more importantly, 84% of participants stated that they had been in a situation where they were unable to find a book they were looking for, and a further 25% of these said that the time taken to find a book had deterred them from visiting the library at a later date.

Turning now to book finding strategies, the majority (91%) of participants reported that they make use of the online search tool provided by the library in order to look up the call number of specific books they are searching for, and often use the PCs within the library itself for this task. Around a third of participants (34%) said that they often browse around the appropriate section for their subject area to find related items. Approaches for remembering the call number of an item between finding it online and locating the shelf varied. A majority of participants (69%) stated that they wrote the number down on a scrap piece of paper (see Fig. 1), while others (19%) said they usually just tried to remember the number. Several participants said they made use of their phone to store the numbers – either scanning the QR codes available on the website (13%), writing down the call number (19%), or taking a photograph of the call number from a computer screen (16%).

When asked about techniques for finding the correct bookshelf, around half of participants either knew where their subject area was located from previous experience (19%) or asked a librarian for directions to the general section (28%) and then used physical signage—such as that shown in Fig. 4—to locate the correct shelf. A further 31% of participants stated that they had on at least one occasion used only the physical signage on the edge of

⁴Note that actual timings for users finding books in the library are investigated in our feasibility study in the next section.

HEK (1995b)
 woodfield (1999)
 RCN (1998)
 Remember
 application
 form!

(a) Example 1: This library user has written down only a partial call number, and includes no titles or other details beyond the last part of each item's locator.

HD 6276.G77
 586 2056
 Putty workfare
 a place



(b) Example 2: This user has noted the entire call number and the title of the book. Incidentally, the call number is copied incorrectly, and would not locate the book in question.

Microeconomics / Wyn Morgan, Michael Katz, Harvey Rosen
 Level 2 WEST HB172.M69
 Cambridge History of Modern Britain / R. Floud, P. Johnston
 Level 2 WEST HC253.C36
 Macroeconomics: theories and policies / Richard T. Froyen
 Level 2 WEST HB172.F76
 The Macroeconomic Debate / Bryan Hillier
 Level 2 WEST HB 199 >HLL 2

(c) Example 3: This library user has written down four separate book titles, authors, call numbers and also the general location of each item in the library ("Level 2 WEST").

Figure 1: Examples of how users currently remember call numbers when using the library. These are real notes discarded by users of the library that were collected during the development phase of this work.

shelves to locate books – a method that by participants' own admission was very time consuming. Interestingly, only 38% of participants said that they had made use of the maps located around the library to find books, and none had ever carried around a personal copy of a map to help locate the correct shelf.

3.3. Librarian interviews

We interviewed two librarians (both female, aged 46–55) to gather feedback regarding the proposed library navigation system. One of the participants was a subject specialist, whose focus is on Arts and Humanities; the other manages a team of workers who replace items after they have been returned from library users. We demonstrated a prototype of the application, showing how the general technique could be used to locate any item in the library. Feedback was positive, though one concern raised was that in order to scan items, users of the BookMark system might remove books from the shelves and not return them (or return them in an incorrect position).

Both librarians commented on how library staff typically spend a large proportion of their time helping people find books. Currently, to help users find items, staff often hand out a paper map of the relevant section, and mark approximately where the item is located. If the borrower needs further help, a librarian will normally lead the user to the item's location in person. The subject librarian remarked that a system developed explicitly for this purpose would have a positive impact, freeing up time for other day-to-day tasks.

3.4. Call number recall strategies

It was clear from our interviews that participants' book search techniques, perhaps not surprisingly, revolved primarily around the call number of the item they were looking




Figure 2: Floor plan of the library in which the BookMark prototype was deployed.

for. To investigate this aspect further, on each library visit during the development stage of this work, we also made efforts to document book finding strategies on a more informal basis. We collected notes that library users had discarded, and observed library users as they attempted to find books.

Figure 1 shows examples of the types of discarded notes collected from library users. The types of notes that book searchers use range in detail from partial call numbers (Fig. 1a) to full titles, author lists, call numbers and locations within the library (Fig. 1c). The key aspect in all notes, however, is the full or partial call number itself. This finding led us to ensure that, in our design, once the library user is within the vicinity of the item they are searching for, the BookMark system shows the full call number (rather than, for example, the item's title or author details), in order to help them focus on the primary identifier in their search.

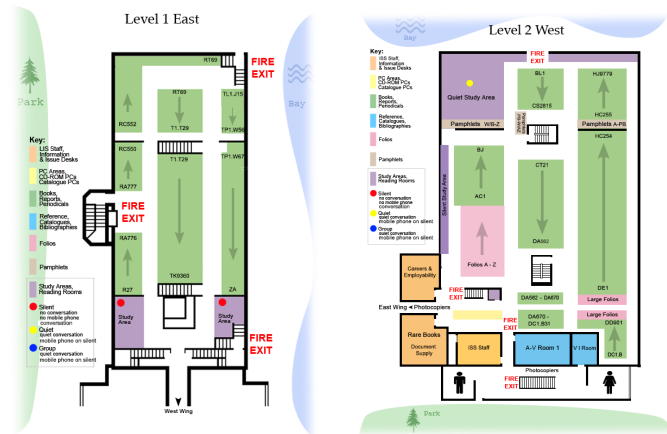


Figure 3: Existing maps of the library. Each level has two separate maps (for West and East wings). Maps are manually updated around once per year as the collection is extended.

4. The BookMark system

BookMark was developed as an Android application, compatible with any phone that has a camera. The app allows users to search the library catalogue to build up a list of items to find. To locate their desired books, users simply scan any nearby item, which produces a personalised map to the position of the book they are looking for, or indeed any item in the library’s collection.

4.1. The library

We worked in partnership with Swansea University library staff for an initial deployment of the BookMark prototype in the Park Campus Library and Information Centre. The library itself is located on a busy campus-based University, and services around 16,000 students and 2,500 staff, with just under one million items loaned per year on average, to 12,000 active borrowers.

The library building spans four floors and two wings (see Fig. 2), covering approximately 11,500 m². It has a total of around 850,000 books located in 3,200 individual bookcases (also known as stacks). Books are ordered using the Library of Congress classification system (Library of Congress, 2014) and can be located within the physical space by searching for a specific call number (e.g., QA76.6 .T44 2007). Currently, visitors to the library can look up the call number of a book they wish to find by using an online search facility from their phone or computer.

The main library building—built in 1920—has several extensions to its core structure, with many corridors and stairwells connecting the different areas. Maps of the basic layout of the building, along with block-level book locations are available from the front information desk (see Fig. 3). Each row of bookshelves in the library also contains a printed sign indicating the range of books contained within it (see Fig. 4). This is the highest level of granularity available within the library – i.e., there are no signs indicating the range of books in specific shelves or stacks within rows.



Figure 4: Current physical shelf numbering within the library. As with floor maps, these end-of-shelf signs are updated when major collection changes happen, but often become out of date due to minor moves and natural reorganisations.

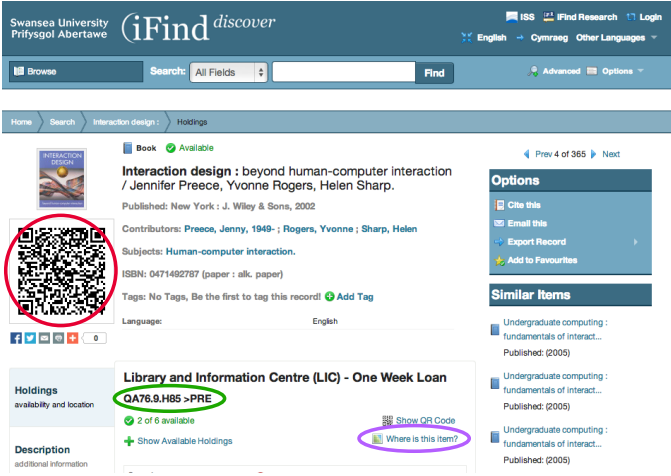


Figure 5: Example of the search feature on the library’s website at the time of initial deployment. The call number and a QR code for transferring the call number automatically to a mobile device are highlighted in red and green. Clicking on the link circled in purple provides the user with a basic map of the item’s floor and wing, as seen in Fig. 3.

The library’s online book search engine can be accessed by visitors to help look up items. Figure 5 shows an example of this resource, which provides users with an item’s call number as well as a QR code that can be used to transfer book details to a mobile device. Once the call number of a book has been identified, library users typically then use the signposts and maps located around the building to find the general area in which a book is located (e.g., by wing and floor only). As with many such large sorted collections, finding individual items can be difficult due to both the sheer number of items and their arrangement within the building.

4.1.1. Mapping the library

In order to provide fine-grained directions to specific bookcases within the library, we first needed to create maps of each floor. These maps included the locations of stairwells, doors, exits, elevators, obstacles—such as pillars or interior walls—and, of course, every bookcase within the physical space. We converted the existing library

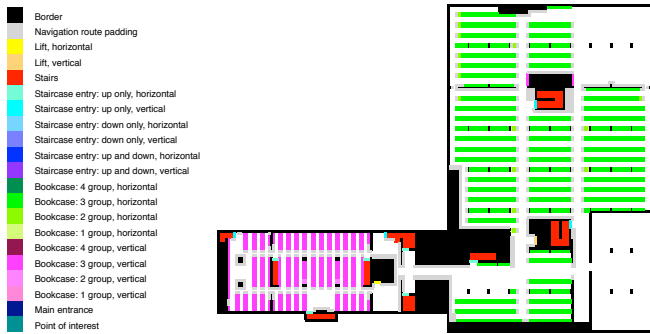


Figure 6: The colour coded SVG map of Level 1 used to generate navigation instructions within the library.

maps to SVG formatted graphics, and colour coded each specific element to assist with navigation. For example, we distinguish between entrance and staircase types, as well as different sized and shaped bookcases and obstacles (see Fig. 6).

The second step in the process of mapping the library was to create a database of call number ranges to allow the system to determine the location of any item within the collection. This provides the core mapping capability of the system, taking advantage of the fact that the ordering of the items is known, and therefore allowing us to calculate the location of any item within the collection in relation to the locations of key items. To create the map, we simply record the call number of the first (key) item in each shelf. From this, we are then able to determine the logical progression of bookshelves in the physical space (see Fig. 7 for an example).

In this initial deployment, the mapping process took approximately one week, including the creation of maps and scanning of the first item in each bookcase. While this is a relatively time consuming process that needs to be completed manually before navigation is possible, subsequent updates to the physical space or book layouts can be handled incrementally. We have partnered with library staff to incorporate this step into their day-to-day shelf management, and have added administrative functionality into our client application to allow them to update shelf positions as and when items are moved. For small moves, this involves only scanning any item on a shelf (to locate), then scanning the new first item on the shelf. For larger moves, users are able to manually select bookcases in the app in order to update their contents.

4.2. Locating items

In order to locate an item, BookMark relies on the database created in the mapping step discussed in Section 4.1.1. First, each of the call numbers gathered in this mapping process is normalised, transforming the locator code into a sortable form. The key book positions are stored in the location database and associated with bookcase numbers that correspond to those stored in floor maps (cf. Fig. 6). At this point it becomes straightforward to retrieve the




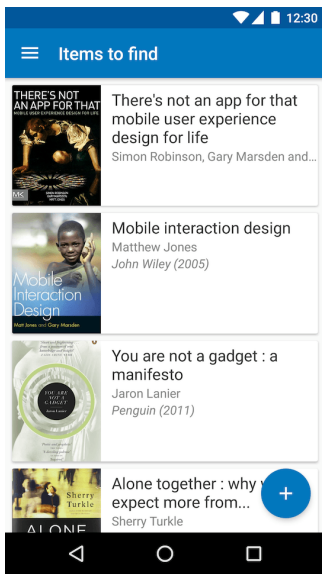
Figure 7: An example of how call numbers flow between shelves in the library. The call numbers of the first books in each of the three bookcases are shown with arrows above the image. From these key locators we are able to determine where any other book can be found. In the example shown, we can deduce that any books with call numbers between GA304.R3 HAR and GB55.F85 2002 will be in the stack highlighted in red, and any books with call numbers between GB55.F85 2002 and GB401.5E26 1999 will be in the bookcase highlighted in blue. Any items with call numbers greater than GB401.5E26 1999 will be located in the yellow stack (or adjacent bookcases), and so on. In this way, by recording the call number of the top-leftmost book of each bookcase in the library, we are able to locate any item, and use this to locate both books that are being searched for, and users who remove a book from a nearby shelf to scan and locate themselves.

bookcase that any item is located within by performing an ordered database query.

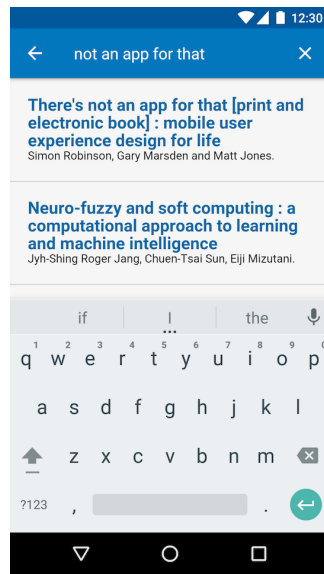
This step of the process is performed by the server-side component of BookMark, which is a simple database application. When the mobile client is requested to retrieve an item location, the server-side process connects to the library's existing book API to retrieve the item's details (including its call number), then looks up the location within the book database and returns a JSON-formatted response.

There are a number of minor complexities in this process, such as the specific ordering of certain items which, despite their call numbers putting them in a particular place in the general library ordering, are either part of a sub-collection or a gift of items that are stored separately. Our system deals with this by grouping them into separate sub-collections for search, and checking the collection type before returning location results.


It may also be worth noting here that not all parts of the building can be accessed from all other parts, as might be expected. For example, Level 4 East is not connected to Level 4 West. In these situations, the application needs to route the user down a level and across a connecting corridor before going back up to the required position.




(a) BookMark's main screen: a list of items to find. Items can be added to this list by searching the library catalogue (see Fig. 8b). Selecting any of the books from the list will take the user to the navigation screen shown in Fig. 8d.




(b) Searching for items. BookMark's search feature uses the existing library search API to find items by keywords, ISBN or call number.




(c) Search results. Search results show the item's details and availability for borrowing. Choosing 'Add to my items' adds the book to the list of items to find as shown in Fig. 8a.




(d) Scanning items to retrieve directions. Once a desired book has been selected from the list in Fig. 8a, users can then scan the barcode of any nearby book to obtain a personalised map to the item they wish to find.




(e) Locating users. In this case, the desired book is on the floor directly below where the user is standing, so the app directs the user to the nearest staircase and asks them to touch the screen when they arrive in order to see the next instruction.



(f) Locating items. When a user arrives on the correct floor, the map is updated to give directions to the precise location of the desired book.



(g) Additional item information. During the final navigation step, the app displays the searched-for item's call number to help identify the correct item. Clicking the 'i' button reveals additional information.



(h) Navigating without precise positioning. BookMark does not restrict users to scanning nearby items – if the user chooses not to scan a book to locate themselves, they receive less-precise instructions starting from the building entrance.

Figure 8: BookMark screenshots, showing the system's various capabilities, and the steps involved in finding items within the library.

4.3. BookMark client application

Figure 8 shows the BookMark client application that we developed. To use BookMark, users first select the items they wish to locate. This can be done within the application by searching the existing online library catalogue (the library offers a reliable internet connection throughout), adding desired books to a list of items to find (see Fig. 8a).

After selecting a book from the list, the user must then inform the app of their current location in order for navigation to begin. This step can be done in one of two ways. The primary method is to scan the barcode on the back of any nearby book (see Fig. 8d). By using the methods described previously to determine where the scanned book is located, we are able to pinpoint the closest bookshelf to where the user is standing, and hence draw a map from here to the desired item’s location see Fig. 8e. If the user is not near to a bookshelf, they can choose to start navigation from the library’s main entrance (see Fig. 8h).

If the desired book is on a different floor to where the user is standing, then the application will first navigate the user to the nearest stairwell or lift, and direct them to tap the screen once they arrive (see Fig. 8e). Upon this action, the app gives directions from the stairwell to the next stage of the route or—if they are now on the correct floor—to the desired book (see Fig. 8f). BookMark uses the A* shortest path algorithm of Hart et al. (1968) for calculating paths between points in the library. The application automatically zooms and pans the map to incorporate the entire route, but can be zoomed in and out manually by using standard pinch gestures to allow users to check details of the route.

It is important to note that since BookMark does not rely on any indoor positioning beyond the user’s barcode scans or arrival confirmations, it cannot follow them along the suggested route in the same way as, say, GPS navigation systems do. That is, the application is not aware of the user’s location along the route at any point after the starting instruction. It assumes, therefore, that the user follows the path set out for them. When directing the user to a new floor, then, it assumes that they will come out at the point directed and will therefore begin on the next floor from that entrance. This lack of tracking could lead to potential issues were the user to get lost; however, we deal with this by allowing users to update the directions given at any point along their path. If the user diverges from the intended route, or gets lost along the way, they can obtain a new map from their current location by simply scanning the barcode of another nearby book.

4.4. Limitations

As a navigation system that is entirely reliant on the position of items within the library’s collection, BookMark has a number of limitations:

Books without barcodes: Older books in the library’s collection often do not have an ISBN barcode, which means that they cannot be used as markers to locate users. A simple solution to this issue is for users to select books

with barcodes when scanning to determine their location. Alternatively, future versions of the app could be extended to take advantage of the various other markers that items often have (such as non-ISBN barcodes, or library RFID tags). It is important to note that all books within the library do have a call number, which means that items without barcodes can still be found for borrowing.

Books in the wrong shelf location: If a user scans a book that is in the wrong physical location (or searches for a book that has been placed in the wrong position), they will be shown an incorrect map. BookMark has no way of detecting that item positions are incorrect; however, we note that this occurrence is no different to the current status quo, where shelf edge signposts are of no help if actual item positions are incorrect. The best remedy in this situation is for users to scan another book to get an updated map.

Reorganisation of books in the library: Libraries are constantly updating their collections, and will sometimes need to reorganise their shelves. If not managed, this process will result in inconsistencies between book locations in the app and physical book locations on library shelves. As noted in Section 4.1.1, we have partnered with library staff to allow shelf stackers to update item positions as and when parts of the collection are relocated, which alleviates this potential issue.

5. Feasibility study

Before embarking on a long-term deployment of the BookMark design, we performed a smaller-scale feasibility study in order to test the design’s viability and effectiveness, and gain subjective feedback from participants. In addition, while our focus in this work is not about optimising book finding per se, we used part of this trial as an opportunity to compare BookMark against the standard static map-based method for locating books within the library. We recruited 32 participants (16 male; 16 female) between the ages of 18 and 65 (average: 32) to take part in the study.

5.1. Procedure

The goal in this study was for participants to locate 10 books from a selection of 20 stored over two levels in one wing of the library (Levels 1 and 2 East). Half of the participants used BookMark to find items; the remainder were allocated as a control group, and used a standard library map. Each participant in the control group was required to mirror the specific combination of books found by a participant in the BookMark group. That is, participant 1 in the control group was given a list of items that were chosen by participant 1 in the BookMark group, and so on. Ages and genders were split equally between both groups. The study followed an IRB-approved ethical consent process, and participants were given a gift voucher in return for participation. Each study session took around 20 to 30 minutes to complete. The procedure for each of the groups was as follows:

5.1.1. BookMark group

For the 16 participants using the BookMark system, the application displayed a list of 20 books (10 on each of the two levels), randomised between participants and showing no indication of the location of the book other than the title, author and call number – that is, we customised the BookMark application to remove the book cover and other details in order to focus purely on book finding. This randomisation of books in the list resulted in all 16 participants having a different ordered combination of books to find during the study. Timings were automatically logged by the application in order to allow direct comparisons with the control group described in the next section. The recorded time taken to locate a book started from when the participant selected a book from the list, and ended when they found the correct book and marked in the application that it had been found.

At the start of the study, each participant was met individually in the library for a study briefing and to obtain informed consent. At this point, each participant was also asked to give Likert-like ratings for how easy or difficult they felt it was to find book locations within the library on a typical visit. Participants were then taken to the correct wing of the library and given a demonstration of the system. This demonstration involved finding a book on a different level to the one they were currently standing on, introducing them to the change level scenario shown in Fig. 8e.

Participants were then asked to use the application to find 10 books within the list of 20 shown. This involved selecting a book to search for, then scanning any other book to retrieve a custom map. At any point during the study, participants could choose to scan further books to update the map. Once they had found the correct book, participants were asked to locate it on the shelf (to prove that it had been found), then mark that item as found in the application and move on to the next until they had located 10 items. Upon completion of these 10 tasks, each participant took part in a semi-structured interview, and gave responses to Likert-like ratings of the ease and speed of finding books using the BookMark technique. Participants were also asked to give general comments on the system’s usability and overall features.

A member of our research team accompanied the participant at all times to ensure their safety and to observe behaviour. No assistance was given to participants during the study tasks, nor were they prompted as to which books to scan to find their location – participants were free to scan any item. This random selection of books by participants aimed to ensure a fair and unbiased evaluation of the system by removing any influence over participants’ choice of items to find, or the order and approach by which they were found.

5.1.2. Control group

Instead of using our system to navigate to shelves, the 16 participants in the control group were given printed

maps of the two floors used for the BookMark group (see Fig. 3 (left) for an example). Again, these users were asked to locate 10 books, but in this case from a printed list with the title, author and call number clearly visible. Unlike the BookMark group, where users were able to select 10 random books from a list of 20 items, each participant in this study was required to mirror the specific combination of books found by a participant in the previous study. That is, participant 1 in the control group was given a list of items that were chosen by participant 1 in the BookMark group, and so on. This mirroring of book lists allowed us to directly compare the time taken to find each book via the two different methods. Apart from the book finding method, all other aspects of the study were the same for the control group as for the BookMark group, with questions focused on the map method, rather than the BookMark method. Timings for the control group were recorded by a researcher observing the participant, beginning when the participant selected an item on the paper list, and ending when they located the book on the shelf.

5.2. Results

All participants successfully found 10 books in both the BookMark and control groups. The average times taken to find items were 73 s for the BookMark group and 117 s for the control group.

5.2.1. BookMark group

The average of participants’ responses for ease of finding items was 4.0 (s.d. = 1.2) for typical library use, compared to 5.9 (s.d. = 0.6) using BookMark (where 1 represented “*extremely difficult*” and 7 was “*extremely easy*”). A Wilcoxon signed-rank test shows a significant difference between these scores ($p < 0.005$, $W+ = 7$, $W- = 129$), indicating that BookMark makes locating books easier than users’ previous experience in the library.

An early anticipated limitation with the BookMark design (see Section 4.4) was the prevalence of books that do not include a barcode, and so cannot be used to locate users. To test whether this was an issue, one of the post-study questions asked participants to rate (again on a Likert-like scale) how easy or difficult it was to find a starting book with a scannable barcode on the back. The average score given for this question was 6.6 (s.d. = 0.6) suggesting that a lack of barcodes to scan was not a significant problem for participants.

Qualitative comments made by users during the post-study interview sessions reinforced these findings. For example, one participant stated: “*it was certainly better than all my other experiences navigating in the library*”, while another said: “*it would save a lot of time*”, and a third commented: “*the first time I came in here I was completely lost, so something like this would be very helpful*”.

When asked if BookMark might change the way they use the library, all participants responded positively. One participant, for example, stated: “*I’d definitely be encouraged to visit more often – I wouldn’t have to get three or*

four books in one trip but could come in more frequently just to find the book I need on that specific day. Until now I considered it a challenge and wasn't prepared to do challenging things all the time, but with this I can just get a book; visit again two or three days later and get another one", while another commented: "It might mean I use the library more often, or be more inclined to use the library. It just makes it a little less intimidating when you've got a useful workable map".

Observations during the book finding tasks supported ratings for the system's effectiveness and speed, but also illustrated methods of recovering from potential errors. Of the 160 books found by the BookMark group, there were 15 instances (2.5%) where participants became "lost" – that is, they misread the map and needed to regain their bearings mid-search. Nine of these occasions (six participants) resulted in the rescanning of a different book to generate a new map. One participant even commented when rescanning a book: "right – I am lost; but it's OK because I can help myself." On the remaining six occasions (three people), participants got back on track by using other methods, such as counting physical shelves, searching for landmarks (such as doors, pillars or stairways), or matching the call number of the desired book to physical signage. In the vast majority of the 160 book finding tasks, participants moved, uninterrupted, directly from their position to the location of their desired book.

5.2.2. Control group

The mean time to find a book for the control group was 117s (min = 23s, max = 611s, s.d. = 81s) compared with 73s (min = 15s, max = 239s, s.d. = 40s) using BookMark. A Welch's t-test conducted on this result shows it to be highly significant: $p < 0.0001$, $t = 6.07$, $df = 233$.

As noted earlier, the focus of this work is not on speed of book retrieval; rather, we are interested in the potential for appropriating existing markers as navigation aids. Despite this, the control group's performance demonstrated that BookMark was faster than using a printed map 76% of the time, with an average time difference of 44s in BookMark's favour.

5.3. Discussion

It is clear from the results of this feasibility study that the use of existing markers on library items as navigation markers leads to both speed and user experience improvements. Currently, with the exception of a librarian leading them to exactly where a book is located, the use of printed maps is really the 'best case scenario' for users of this library. While these maps do show the general location of where book sections lie, they do not provide shelf-level precision, and users will generally be required to use the map to locate the general section, and then look for physical signage on the side of shelves to locate their desired book. In comparison, the BookMark approach allows users to get precise directions to any item within the entire library collection. This benefit is highlighted in participants'

comments and the ratings given for various aspects of the BookMark group's tasks.

Turning to the time taken to locate books, it is worth noting that this measure also includes the time taken to search within a shelf – a task that neither BookMark nor a printed map can aid with. This particular part of the search varied between participants as the understanding of call numbers was not as clear to some as it was to others. As is often the case, participants' book finding result times do not closely match the estimated durations recorded in our interviews with current visitors to the library. This is perhaps to be expected, however, as the study here was performed over only two floors on one wing of the library, whereas participants in the library use study were responding about the whole library building. In this feasibility study, participants were told at the start that the books existed in these locations, which greatly reduced the area within which the search needed to be conducted.

6. Longitudinal deployment

As part of our investigation into scalable indoor navigation via appropriation, we wanted to observe use of the application in a more naturalistic setting – that is, actual users of the library as opposed to participants recruited and incentivised to take part. To achieve this goal, we conducted a longitudinal deployment of the BookMark system over a period of two years (December 2013–2015).

6.1. Hypotheses

From our earlier trials of the BookMark technique it was clear that users were able to use scanning of nearby books to help them find items quicker than existing methods. In this trial, then, we focused on broader questions around particular use-cases and the effectiveness of the system. In particular, we were interested in the scanning behaviours of users, including the situations in which they scan items, and the factors that are related to this. We had two hypotheses:

- H1** Users will be more likely to scan a nearby book's barcode to find items after they have already found the first item in a single visit (e.g., they will start from the library entrance for the first item);
- H2** Users will be more likely to scan a nearby book's barcode to find items that are further away (for example, if it is necessary for them to change to a different level).

6.2. Procedure

We published the BookMark application on Google Play,⁵ and built awareness of its capabilities via a short internal promotion campaign. Visitors to the library were made aware of the app's existence via posters, flyers and TV screens around campus for a 1-month period. An

⁵Available at: <https://play.google.com/store/apps/details?id=ac.robinson.bookmark>

IRB-approved consent screen when first installing the application made it clear that the system was part of a study, and that anonymised usage statistics would be remotely recorded. Users who did not consent were not able to use the application. No monetary or other incentives were given.

6.3. Measures

For this experiment we gathered data primarily from participants' anonymised usage logs, recorded automatically while using the BookMark application. Each logged event consisted of:

- An event timestamp and unique user identifier;
- The type of action performed (e.g., a search, book found, item added to list, etc.).

Certain event types also included other fields – for example, a book identifier or rating response.

In order to gather more qualitative data, we included two short questionnaires in the application. The first is displayed after an item is found using the app, and asks users to rate the accuracy of the navigation directions given on a four-point Likert-like scale from “*I couldn't find the item*” to “*Exactly the right location*”. This questionnaire is optional – users can skip the screen to go directly to their next item.

After a user has located three books using BookMark, the app asks them to fill out a short, optional, feedback questionnaire. These questions ask users to compare BookMark to their previous experience of finding books in this library. Specifically, users are asked to rate, on a scale of 1 to 7 where 7 is high, the extent to which BookMark is slower or faster than previous methods of finding items; the extent to which BookMark is easier or harder to use than previous methods; and, the extent to which BookMark has improved their overall experience in this library.

From logged event information we can analyse a range of activities, including the number of books found, the number of ISBN scans needed to find a book, and the number of floor changes taken per book, for example. We are also able to determine how long it takes to find a book, calculating this as the difference between the time when a map to a book is generated by the application and the time a user either exits the app or provides an accuracy rating. When combined with the qualitative responses from questionnaires, these logs provide a broad picture of overall usage.

Finally, the app also allows users to leave feedback in text form at any point, and we used this in combination with the Google Play analytics, rating and comment system in order to track broader usage and user feedback over the two-year period.

6.4. Results

As of 1st December 2015 the BookMark application had been downloaded by 863 unique users on Google Play over the preceding two-year period. Of these, the number

of users who opened the app and completed a search for a book was 710 (approximately 5% of the student body of the University). A total of 8317 books were found across all users, with an average of 11.7 books found per person (min = 1, max = 222, s.d. = 18). The average number of books found per visit was 4.8 (min = 1, max = 38, s.d. = 4.7), with an average of 161 s (2 min 41 s, s.d. = 385 s) to find each book.

6.4.1. Filtering logged data

It is important to ensure that analysis of longitudinal study data accurately reflects actual user behaviour. As such, we chose to process the raw logged data to ensure its accuracy. We filtered data based on two conditions:

Incomplete searches: As might be expected, not all book searches result in users actually finding the physical copy of an item. In our design, we asked users to report when they found a book via their rating of the accuracy of the app's navigation instructions. However, many users did not get to this point, and simply switched away from or exited the app. While it is probable that many of these users did indeed find the book they were looking for, we cannot assume that this is the case for all occurrences. In future versions of the system, we could partially address this issue by combining BookMark with the library's checkout system to streamline the book borrowing process and ensure that any completed book searches are tracked. However, it is important to note that many users do not check out the books they find, instead making use of them within the library's reading spaces and returning them to the shelves without ever taking them off the premises.

In order to ensure complete accuracy of the analysed data, then, we chose to only include results from navigation instances in which the user reached the post-navigation survey question stage (whether they answered the question or not). While this approach is likely to exclude a number of valid and complete book searches, it also gives a relatively concrete bound to the book location activity, and as such, we argue, provides a more verifiable result.

Durations: It became clear during initial analysis that some users apparently chose to test the application before actually using it to find books, resulting in very small navigation times (i.e., sub-15 s, all beginning from the library main entrance). Conversely, some navigation timings were recorded in the region of several hours, suggesting perhaps that users started a search while off-campus, then travelled to the library before completing navigation, or simply paused navigation to resume at a later time.

Our approach to this issue was to assume that searches that are within a reasonable timeframe after the user is first shown the map to their desired item are more likely to be genuine navigation sessions. Instead of choosing arbitrary bounds for this timeframe, we consulted the results from our feasibility study, selecting both the minimum and maximum times taken to find a book across the 160 books found in the BookMark group (min: 15 s; max: 239 s). However, while

this initial experiment was conducted over two floors and one wing of the library, the publicly released version of BookMark navigates over all four floors and two wings of the building. We therefore assumed that the best-case scenario for the minimum time could still feasibly be as low as 15 s, but the worst-case for the maximum time could be as high as 239 s over two further floors and a second wing (i.e., $239 \times 2 \times 2$), giving an estimated maximum duration of 956 s (15 min 56 s).

6.4.2. Filtered results

Applying these filtering conditions to the raw logged data—that is, including only book navigation instances that both were within the 15 s to 956 s window and reached the stage of rating location accuracy—resulted in a total of 4405 books being found by 604 unique users over 1241 library visits.

Users visited the library and used the BookMark app to find items 2.1 times on average over the two-year period (min = 1, max = 17, s.d. = 2.0). The average number of books found per person was 7.3 (min = 1, max = 139, s.d. = 11.1), with 2.8 books found per visit on average (min = 1, max = 8.2, s.d. = 1.3).

The average time taken to find a book was 146 s (2 min 26 s, s.d. = 158 s). The distribution of times taken to find a single item is shown in Fig. 9: the vast majority of items were found in less than eight minutes.

As discussed previously, BookMark allows users to locate themselves in one of two ways: either by scanning barcodes on other items within the library, or by starting navigation at the library entrance. During the two-year deployment, 1618 of the 4405 books (37%) were found via scanning of other books' ISBN barcodes. Users needed to scan, on average, 1.4 books (min = 1, max = 13, s.d. = 1.0) per navigation instance to locate the book they were looking for.

We hypothesised (H1) that users would be less likely to scan a nearby book's barcode to find the first item in a single visit to the library. Of the 1618 navigation instances which involved scanning barcodes, 1367 (84%) were the second or later book found within a single session. This indicates that our hypothesis is correct, and users most likely began from the library entrance to find their first item, but once within the library stacks tended to scan a nearby item to locate themselves.

Our second hypothesis was that the number of barcode scans necessary would be greater if a floor change was required (H2). Figure 10 shows these behaviours in detail, and indicates that the majority of cases where an ISBN barcode was scanned (21% of all navigation instances) took place when it was *not* necessary to change levels. We are therefore not able to confirm our second hypothesis.

There are several possible explanations for this behaviour. For instance, the Library of Congress call number system groups books on a common topic within the same physical area – it may be that users needed only to find

the first item in their list to locate their topic area, and then chose to manually navigate and browse around the collection. Interestingly, 45% of book navigation instances did not include an ISBN scan, but *did* require changing levels. However, this result is most likely due to the physical organisation of the library at which our deployment took place. Users enter the library on Level 3, and the majority of items are stored on Levels 4, 2 or 1, requiring a change of floor in most cases.

Turning now to uses of the app that did not involve barcode scanning (i.e., for which navigation started at the entrance to the library). Of the 2787 (63%) books for which this was the case, 499 were the first books found during a library visit, suggesting that 11% of items were found after navigating straight from the library entrance. The remaining 52% of items were found when users had already found one or more books, but did not use a barcode to locate themselves to find further items. We interpret this as indicating that these users did not need to locate themselves precisely within the indoor space, but instead used the app solely to retrieve the position of the item they were searching for, and interpreted this manually while already near to the correct location.

6.4.3. Ratings and comments

A total of 1351 responses to the post-navigation questionnaire were recorded. 88.4% (1194 responses) of items were found in exactly the right place as indicated by the BookMark app. The majority of the remaining responses were that the item's location was slightly inaccurate (*"a few bookcases away"*; 127 responses, 9.4% overall), with only 30 books placed in completely the wrong location (2.2% overall). No user chose the *"I couldn't find the item"* option.

Over the two-year period there were 150 responses to the in-app questionnaire. Overall, users felt that BookMark was faster to use than previous methods (5.9 out of 7 on average). Similarly, users rated the ease of use in comparison to previous methods at 5.8 out of 7 on average, indicating that BookMark is also easier to use to locate books than other options. Finally, users gave an average result of 5.6 out of 7 for the extent to which their overall experience had been improved, suggesting that BookMark has made a positive impact on the majority of its users' library experiences.

Appreciation of the BookMark app has also been seen in the Google Play ratings, with an average score of 4.5 out of 5 from 28 respondents. In the final two months of deployment a library server update led to a six-day period during which users were unable to search for items (though all other functions of the app continued to work) – the app's only negative reviews were left as a result of this downtime.

There were only 19 feedback messages left during the deployment period. Of these, eight were positive comments such as: *"Very good, worked really well"* and *"Amazing app! Absolutely love it! Brilliant idea. Well done and thank you :)"*. A further four were centred around ideas to improve the design – for example: *"make a 'get me out of here' button"*




Figure 9: Histogram of time taken to find books. The average time taken (146 s) and the 95th percentile (489 s) are indicated by the solid and dashed vertical lines, respectively. The histogram shows a clear long tail of results, with the vast majority of navigation instances taking far less than eight minutes. Note that this graph includes only filtered data, as discussed in [Section 6.4.2](#).




Figure 10: Scatter plot showing the number of ISBN scans against the number of floor changes the user made before finding an item. Circles are scaled according to the number of occurrences, and labelled with the percentage of navigation instances that fall in each group (unlabelled circles represent 0.02% to 2% of results). Note that this graph includes only filtered data, as discussed in [Section 6.4.2](#).

to help me find my way out once I have my book!”, and two people reported the search functionality failure. The remaining five comments were filler text or unintelligible.

7. Discussion

Our studies of the BookMark system show the value of appropriating existing infrastructure—in this case barcodes—to provide indoor navigation support. Rather than modifying or creating entirely new infrastructure, we have shown how piggybacking off a widely used current framework can provide a fast, low-cost, scalable method of indoor routefinding. In contrast to previous research approaches, our technique also requires no building modifications (e.g., signage or beacons) or additional hardware.

Turning first to the core technique used in the BookMark system. During the longitudinal deployment, 37% of navigation tasks were completed with the assistance of

one or more ISBN barcode scans. We see this as a successful uptake rate, especially considering that there is no requirement to scan any items, and that the app displays a normal map direct to any item by default. In addition, other external factors—such as the physical library layout, and the list order of items being searched for—are likely to have lessened the benefit of scanning nearby items for a particular subset of navigation tasks.

It is clear from feedback and questionnaire responses during the longitudinal study that users perceive BookMark to be better than existing methods of locating books. Accuracy ratings were particularly high – 89% of respondents reported that the book they were searching for was in exactly the right place. In terms of ease of use, Likert-like ratings in comparison to other navigation methods yielded results of 5.9 and 5.8 out of 7 from the feasibility and longitudinal studies, respectively. Overall, then, the uptake rate during the longitudinal deployment, and feedback gathered from each of our studies, illustrated both the technique’s

ease of use, and its practical application for day-to-day item finding tasks. Users clearly saw value in the approach, and enjoyed using the app to help them find items in the library.

Turning now to the effectiveness of the general method, it is clear from our feasibility study that BookMark is a faster way of locating books when compared to existing approaches (e.g., using the call number of an item and a printed map of the library). During our two-year deployment, the average time to find a book was greater than in the feasibility study (146 s vs. 73 s), but it is important to see this in the context of use: during the deployment users had a far larger space to navigate within, and usage was in their own time, without external motivation, in a more natural setting. Indeed, compared with participants' own estimations of the normal time taken to find items (as captured in our current library use study), the BookMark method is over three times faster on average. Finally, when asked to rate the speed of the BookMark approach in comparison to previous methods, users during the longitudinal deployment gave an average score of 5.9 out of 7. Overall, then, it is clear that library users who chose to use the BookMark approach saw the technique as clearly beneficial, and a genuinely useful way of getting a precise map from their location to the item they were looking for.

8. Conclusions

While there has been a great deal of research interest in efficient and accurate methods of offering navigation support for indoor spaces, most existing methods have drawbacks around the level of additional infrastructure or maintenance required, or provide directions that degrade in accuracy as the user reaches their destination.

Our approach offers a simple and reliable way of supporting navigation in indoor spaces by taking advantage of existing features of the items that users are attempting to find. The technique's usefulness ranges from precise positioning to "where am I?" queries to recovering from errors – see, for instance, the 60% of participants in the feasibility study who rescanned a new book to regain their bearings. It also reduces the time and cost of implementation by reusing an already widely implemented infrastructure.

8.1. Design implications

While barcodes have been largely overlooked in HCI until now, many researchers have been considering the utility and design of future tags and object identifiers as part of a wider Internet of Things agenda. For example, presently there is much interest in the use of embedded wireless chip identifiers (the most common commercial implementation being NFC), but less attention to existing frameworks that might be adopted instead.

We have been able to piggyback on barcodes for two reasons. Firstly, barcodes are visible and highly recognisable. The first machine-scannable barcode design used

ultraviolet ink, invisible to the user so as not to detract from product packaging. This approach failed, partly due to the code being hidden from the person scanning. The type of information piggybacking we have demonstrated would not have been possible if it were not for the visual properties of the barcodes themselves.

In contrast, RFID and NFC tags, and their ilk, are generally hidden from view, which greatly reduces their appropriability. Furthermore, such tags can be encrypted or use proprietary formats – rightfully so in some security-critical cases, but in others this can seriously restrict the ways in which they can be reused for other purposes. Barcodes are based on an open specification, which greatly broadens the extent to which they can be appropriated.

We have demonstrated how it is possible to piggyback on existing ubiquitous infrastructure without any additional alterations. The question we would like to pose, then, is: how can we make future marker designs appropriable enough for others do to the same? We postulate that developers and designers of future schemes should consider how to make their designs both visible and open to other uses. We therefore suggest that designers ensure that any new digital markers offer visual, haptic or other affordances, and openness wherever possible, to encourage piggybacking as a method of creating new and exciting uses for these infrastructures.

8.2. Limitations and trade-offs

The BookMark system discussed and evaluated here was deployed to explore the real-world feasibility of an entirely infrastructure-free navigation tool. While there are limitations, as previously discussed in [Section 4.4](#), it is clear that the approach is viable in this type of navigation space.

In adaptations of the system for different spaces, or in wider library deployments, it could be valuable to combine the approach with other navigation techniques discussed in previous work. One obvious modification could be to partially mark-up the navigable space with a small number of fixed barcodes or other markers (e.g., at entrances, or at the start of major book range groups). Taking this approach would mean that users would not always need to remove books from the shelf to scan their barcodes, and that the scanned position would remain fixed, hence trading the need for infrastructure for higher certainty about user positioning. However, it is important to note that for fixed markers other than those at entrances, over time this approach would likely suffer from one of the original problems identified in the library, where fixed-position signage tends to become outdated and inaccurate (e.g., see [Fig. 4](#)).

8.3. Future work

There are clearly various improvements that could be made to the BookMark app itself. For example, letting users create a multi-book list and automatically generating the shortest path between all of the items would improve



Figure 11: Barcode piggybacking granularity: (a) shelves in libraries; (b) product zones in supermarkets; (c) departments in department stores.

usability in several of the scenarios that our longitudinal deployment participants seem to have employed.

However, more interesting to us is the potential usage of the technique in other contexts. As we have discussed previously (cf. Robinson et al. 2014), there are many other situations in which the general technique could be used. In places such as libraries, the approach can provide very detailed maps, illustrating exactly where a book is located (Fig. 11 (a)); in others it might suggest only a broader area (Fig. 11 (b) or (c)). Depending on map granularity, then, our piggybacking approach can offer macro, micro and also larger-scale navigation that is entirely supported by the items that the user is searching for.

Acknowledgements

This work has been supported by EPSRC grant numbers EP/J000604/2 and EP/M00421X/1. We would like to thank the librarians and IT support staff at Swansea University Library for giving us permission to deploy BookMark, and for the help and guidance they have provided throughout the process.

References

Ahmetovic, D., Gleason, C., Kitani, K.M., Takagi, H., Asakawa, C., 2016. NavCog: Turn-by-turn smartphone navigation assistant for people with visual impairments or blindness, in: Proc. W4A '16, ACM. pp. 9:1–9:2. DOI: [10.1145/2899475.2899509](https://doi.org/10.1145/2899475.2899509).

Aittola, M., Ryhänen, T., Ojala, T., 2003. SmartLibrary – location-aware mobile library service, in: Proc. MobileHCI '03. Springer. volume 2795 of *LNCS*, pp. 411–416. DOI: [10.1007/978-3-540-45233-1_38](https://doi.org/10.1007/978-3-540-45233-1_38).

Bahl, P., Padmanabhan, V., 2000. RADAR: An in-building RF-based user location and tracking system, in: Proc. INFOCOM '00, pp. 775–784. DOI: [10.1109/INFCOM.2000.832252](https://doi.org/10.1109/INFCOM.2000.832252).

Brush, A.J., Karlson, A., Scott, J., Sarin, R., Jacobs, A., Bond, B., Murillo, O., Hunt, G., Sinclair, M., Hammil, K., Levi, S., 2010. User experiences with activity-based navigation on mobile devices, in: Proc. MobileHCI '10, ACM. pp. 73–82. DOI: [10.1145/1851600.1851616](https://doi.org/10.1145/1851600.1851616).

Buchanan, G., 2010. The fused library: integrating digital and physical libraries with location-aware sensors, in: Proc. JCDL '10, ACM. pp. 273–282. DOI: [10.1145/1816123.1816165](https://doi.org/10.1145/1816123.1816165).

Butz, A., Baus, J., Krüger, A., Lohse, M., 2001. A hybrid indoor navigation system, in: Proc. IUI '01, ACM. pp. 25–32. DOI: [10.1145/359784.359832](https://doi.org/10.1145/359784.359832).

Chawathe, S., 2008. Beacon placement for indoor localization using Bluetooth, in: Proc. ITSC '08, IEEE. pp. 980–985. DOI: [10.1109/ITSC.2008.4732690](https://doi.org/10.1109/ITSC.2008.4732690).

Ciavarella, C., Paternò, F., 2004. The design of a handheld, location-aware guide for indoor environments. *Personal Ubiquitous Comput.* 8, 82–91. DOI: [10.1007/s00779-004-0265-z](https://doi.org/10.1007/s00779-004-0265-z).

Coughlan, J., Manduchi, R., Shen, H., 2006. Cell phone-based way-finding for the visually impaired, in: 1st International Workshop on Mobile Vision.

Dekel, A., Kirkpatrick, S., Noach, N., Schiller, B., 2011. Amazon-on-earth library navigator, in: Proc. Mobility '11, IARIA. pp. 14–21.

Eostein, Z., 2012. An inside look at an Amazon warehouse. URL: <http://news.yahoo.com/inside-look-amazon-warehouse-163514829.html>.

Fallah, N., Apostolopoulos, I., Bekris, K., Folmer, E., 2012. The user as a sensor: navigating users with visual impairments in indoor spaces using tactile landmarks, in: Proc. CHI '12, ACM. pp. 425–432. DOI: [10.1145/2207676.2207735](https://doi.org/10.1145/2207676.2207735).

Fallah, N., Apostolopoulos, I., Bekris, K., Folmer, E., 2013. Indoor human navigation systems: A survey. *Interacting with Computers* DOI: [10.1093/iwc/iws010](https://doi.org/10.1093/iwc/iws010).

Foxlin, E., 2005. Pedestrian tracking with shoe-mounted inertial sensors. *IEEE Computer Graphics and Applications* 25, 38–46. DOI: [10.1109/MCG.2005.140](https://doi.org/10.1109/MCG.2005.140).

Harle, R., 2013. A survey of indoor inertial positioning systems for pedestrians. *IEEE Communications Surveys Tutorials* 15, 1281–1293. DOI: [10.1109/SURV.2012.121912.00075](https://doi.org/10.1109/SURV.2012.121912.00075).

Hart, P.E., Nilsson, N.J., Raphael, B., 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* 4, 100–107. DOI: [10.1109/TSSC.1968.300136](https://doi.org/10.1109/TSSC.1968.300136).

Hu, F., 2015. Emerging Techniques in Vision-based Indoor Localization. Technical Report. The City University of New York.

Hub, A., Diepstraten, J., Ertl, T., 2004. Design and development of an indoor navigation and object identification system for the blind, in: Proc. SIGACCESS '04, ACM. pp. 147–152. DOI: [10.1145/1028630.1028657](https://doi.org/10.1145/1028630.1028657).

Li, F., Zhao, C., Ding, G., Gong, J., Liu, C., Zhao, F., 2012. A reliable and accurate indoor localization method using phone inertial sensors, in: Proc. UbiComp '12, ACM. pp. 421–430. DOI: [10.1145/2370216.2370280](https://doi.org/10.1145/2370216.2370280).

Library of Congress, 2014. Library of congress classification. URL: <https://www.loc.gov/catdir/cpsol/lcc.html>.

Mantyljärvi, J., Paternò, F., Salvador, Z., Santoro, C., 2006. Scan and tilt: Towards natural interaction for mobile museum guides, in: Proc. MobileHCI '06, ACM. pp. 191–194. DOI: [10.1145/1152215.1152256](https://doi.org/10.1145/1152215.1152256).

Meese, R., Ali, S., Thorne, E., Benford, S., Quinn, A., Mortier, R., Koleva, B., Pridmore, T., Baurley, S., 2013. From codes to patterns: Designing interactive decoration for tableware, in: Proc. CHI '13, ACM. pp. 931–940. DOI: [10.1145/2470654.2466119](https://doi.org/10.1145/2470654.2466119).

Mulloni, A., Wagner, D., Barakonyi, I., Schmalstieg, D., 2009. Indoor positioning and navigation with camera phones. *Pervasive Computing* 8, 22–31. DOI: [10.1109/MPRV.2009.30](https://doi.org/10.1109/MPRV.2009.30).

Nicholson, J., Kulyukin, V., Coster, D., 2009. ShopTalk: Independent blind shopping through verbal route directions and barcode scans. *The Open Rehabilitation Journal* 2, 11–23. DOI: [10.2174/1874943700902010011](https://doi.org/10.2174/1874943700902010011).

Retscher, G., Fu, Q., 2010. Continuous indoor navigation with RFID and INS, in: Proc. PLANS '10, IEEE. pp. 102–112. DOI: [10.1109/PLANS.2010.5507242](https://doi.org/10.1109/PLANS.2010.5507242).

Robinson, S., Pearson, J., Jones, M., 2014. A billion signposts: Repurposing barcodes for indoor navigation, in: Proc. CHI '14, ACM. pp. 639–642. DOI: [10.1145/2556288.2556994](https://doi.org/10.1145/2556288.2556994).

Satpathy, L., Mathew, A.P., 2006. RFID assistance system for faster book search in public libraries, in: Proc. CHI EA '06, ACM. pp. 1289–1294. DOI: [10.1145/1125451.1125691](https://doi.org/10.1145/1125451.1125691).

Sciacchitano, B., Cerwinski, C., Brown, I., Sampat, M., Lee, J., McCrickard, D., 2006. Intelligent library navigation using location-aware systems: the Newman project, in: Proc. ACM-SE '06, ACM. pp. 371–376. DOI: [10.1145/1185448.1185531](https://doi.org/10.1145/1185448.1185531).

Serra, A., Carboni, D., Marotto, V., 2010. Indoor pedestrian navigation system using a modern smartphone, in: Proc. MobileHCI '10, ACM. pp. 397–398. DOI: [10.1145/1851600.1851683](https://doi.org/10.1145/1851600.1851683).

- Taher, F., Cheverst, K., 2011. Exploring user preferences for indoor navigation support through a combination of mobile and fixed displays, in: Proc. MobileHCI '11, ACM. pp. 201–210. DOI: [10.1145/2037373.2037405](https://doi.org/10.1145/2037373.2037405).
- Walsh, A., 2011. Blurring the boundaries between our physical and electronic libraries: Location-aware technologies, QR codes and RFID tags. *The Electronic Library* 29, 429–437. DOI: [10.1108/026404711111156713](https://doi.org/10.1108/026404711111156713).
- Wang, Y., Yang, X., Zhao, Y., Liu, Y., Cuthbert, L., 2013. Bluetooth positioning using RSSI and triangulation methods, in: Proc. CCNC '13, IEEE. pp. 837–842. DOI: [10.1109/CCNC.2013.6488558](https://doi.org/10.1109/CCNC.2013.6488558).
- Watanabe, K., Takahashi, T., Ando, T., Takahashi, K., Sasaki, Y., Funakoshi, T., 2010. LiNS: A library navigation system using sensors and smartphones, in: Proc. BWCCA '10, IEEE. pp. 346–350. DOI: [10.1109/BWCCA.2010.94](https://doi.org/10.1109/BWCCA.2010.94).
- Woodman, O., Harle, R., 2008. Pedestrian localisation for indoor environments, in: Proc. UbiComp '08, ACM. pp. 114–123. DOI: [10.1145/1409635.1409651](https://doi.org/10.1145/1409635.1409651).