The StoryPlaces Platform: Building a Web-Based Locative Hypertext System

Charlie Hargood
Creative Technology
Bournemouth University
chargood@bournemouth.ac.uk

Mark J. Weal
Web and Internet Science
University of Southampton
mjw@ecs.soton.ac.uk

David E. Millard
Web and Internet Science
University of Southampton
dem@ecs.soton.ac.uk

ABSTRACT

Locative narrative systems have been a popular area of research for nearly two decades, but they are often bespoke systems, developed for particular deployments, or to demonstrate novel technologies. This has meant that they are short-lived, the narratives have been constructed by the creators of the system, and that the barrier to creating locative experiences has remained high due to a lack of common tools.

We set out to create a platform based on the commonalities of these historic systems, with a focus on hypertext structure, and designed to enable locative based narratives to be created, deployed, and experienced in-the-wild.

The result is StoryPlaces, an open source locative hypertext platform and authoring tool designed around a sculptural hypertext engine and built with existing Web technologies. As well as providing an open platform for future development, StoryPlaces also offers novelty in its management of location, including the separation of location and nodes, of descriptions from locations, and of content from pages, as well as being designed to have run-time caching and disconnection resilience. It also advances the state of the art in sculptural hypertext systems delivery through conditional functions, and nested, geographic and temporal conditions.

The StoryPlaces platform has been used for the public deployment of over twenty locative narratives, and demonstrates the effectiveness of a general platform for delivering complex locative narrative experiences. In this paper we describe the process of creating the platform and our insights on the design of locative hypertext platforms.

CCS CONCEPTS

• **Human-centered computing** → **Hypertext/hypermedia**;

KEYWORDS

Location-Based Narrative, Sculptural Hypertext

ACM Reference format:
Charlie Hargood, Mark J. Weal, and David E. Millard. 2018. The StoryPlaces Platform: Building a Web-Based Locative Hypertext System. In *Proceedings of The 28th ACM Conference on Hypertext and Social Media, Baltimore, USA, July 2017 (HYPERTEXT'18),* 9 pages.
DOI: 10.475/123_4

HYPERTEXT'18, Baltimore, USA

DOI: 10.475/123_4

1 INTRODUCTION

In 2017 Claus Atzenbeck told the hypertext research community that "research on traditional hypertext systems has been fading out over the past decade" [2] - something that very much appears to be true. The early concepts such as Memex [12] and projects like Xanadu [27], led to later systems such as Hypercard [32] and StorySpace [5, 8], which in turn inspired the Hypertext ideas of the 90s such as MICROCOSM [18], and adaptive hypertext in Interbook [10] and AHA [14], as well as "strange Hypertext" experiments such as CardShark [3]. But since then there has been little in the way of new hypertext infrastructure presented in the literature save for a new version of StorySpace [6] and Atzenbeck's own system ODIN [2]. In contrast there have been a number of commercial tools that have become popular in the interactive fiction community, such as Twine[1] and Inform[2].

The rise of the Web undeniably had an impact but these examples show us there is still much to learn from Hypertext implementations in a post Web world. In 1999 Mark Bernstein asked us "Where are the hypertexts?", nearly twenty years later we might ask "Where are the hypertext systems?".

StoryPlaces[3], as a research project, sought to better understand locative narrative (interactive stories that are contextually aware and respond to the readers location). It followed work seeking a generic model for all location aware narrative [25] and has worked both to understanding the poetics of location based narrative [29], its hypertextual patterns and structures [20], and the impact on the writing process [23]. During the project three story deployments were made each

[1] http://twinery.org/ as of 17/4/18
[2] http://inform7.com/ as of 17/4/18
[3] http://storyplaces.soton.ac.uk/ as of 31/1/2018

delivering 6 stories written by a combination of volunteer and paid authors in three different UK cities (Southampton, Bournemouth, and London) - alongside another experimental story created as part of a reflection on the writing process [20]. These stories were delivered using the StoryPlaces platform - a new generic sculptural hypertext system.

In this paper we present the StoryPlaces platform, its features and innovations, the rationale behind its design, and the lessons learned in the engineering of a new generic sculptural hypertext platform.

2 BACKGROUND ON LOCATIVE NARRATIVE

Early locative systems often took the form of tour guides, linking a database of information to specific locations in the environment to create personalised pages designed to be read in-situ [9]. Early systems such as GUIDE had to deal with relatively inaccurate location data [13], adopting strategies where the user was involved in choosing the correct location from a shortlist generated by the application, but later systems could rely on more stable location technologies such as GPS, and the research focus moved to the experience itself. This often took the form of narrative experiences such as a mosaic of personal stories [28] or historical plays that unfold in a given location [7]. The potential of these sorts of experience for heritage applications has been noted for both navigation and education [19], as well as considered as a way of augmenting more conventional exhibition spaces [21].

Educational applications of locative hypertext have also been popular. The Chawton House project combined digital information with physical activities for schoolchildren in the grounds of a period house [34], for example giving children information about the use of a particular area, and then asking the children to write a poem in their notebooks about the household interacting in that space. Other systems use digital activities supported directly by the system, for example 'Gaius' Day in Egnathia' [1] asks users to collect location-based clues that draws them through the exploration of Egnathia, an ancient Roman city in Southern Italy.

Researchers have also explored how the real world can be incorporated into the digital experience. In 'University of Death' [11] players collect and use real world props alongside digital clues to progress. Augmented Reality (AR) is a popular approach - 'Viking Ghost Hunt' [26] allows players to hunt ghosts on the streets of Dublin, using AR technology to visualize their targets. AR can also be used to present information overlaid over real world targets, or even to recreate historical buildings and places using the device [17].

StoryPlaces is concerned with applications that are closer to interactive fiction, with a focus on story and interactive structure. These locative narrative experiences have been termed locative literature' [22] or 'ambient literature'[16]. Some examples, such as the 'iLand of Madeira', have simple structure, in the case of iLand relying on a mosaic of story nodes that are gradually revealed to the reader [15]. Other examples use more complex rules and draw on not only reader choices but contextual information, for example 'San Servolo, travel into the memory of an island', which considers not only location but also weather [30].

Locative literature can be conceptualised as a hypertext, and in particular as a *Sculptural Hypertext* system, where all nodes are potentially linked, but those links are 'sculpted away' based on rules that are compared to the readers' state [4, 35]. Sculptural hypertext is a good fit for locative systems, as location can be modeled as part of the users state, allowing location to be incorporated seamlessly alongside logical rules (for example, 'Node C is visible once you have read Node B, and you are standing in the Courtyard'. Sculptural Hypertext has been shown to support linear, branching, and open locative structures (described as Canyons, Deltas, and Plains) which when combined are sufficient to describe a wide range of location-based narratives [25]. It is this model that drives the StoryPlaces Platform, and enables the software to support a wide range of locative experiences.

3 THE STORYPLACES PROJECT AND DESIGN PROCESS

StoryPlaces was a two and a half year Leverhulme Trust funded project investigating the structures, poetics, impact, and application of locative narratives. Its name is a deliberate homage to Storyspace [5, 6, 8] and one of its key contributors (Mark Bernstein) who did much to establish the idea of sculptural hypertext [3]. The project was an interdisciplinary collaboration between the authors as computer scientists, literary academics, and writers. Locative fiction has largely existed within bespoke systems created for individual stories. This is expensive, forces authors to acquire technical skills or require collaborators, and is often not sustainable in the long term as maintenance for each application only lasts as long as its initial deployment. A general framework for locative narrative supported by an authoring tool could unlock the creation of locative narrative to a wider base of writers, create a technological standard for longer term community maintenance as an open source project, and provide a platform for the StoryPlaces project to create a range of stories to support its own investigation into locative fiction.

The framework was designed based on a combination of participatory design, co-operative inquiry, and systematic review. The sculptural model was based on the previous work *Canyons, Deltas, and Plains* - as common structures of locative narratives that could potentially support the majority of systems described in the literature [25]. Initial low-fi experiments with creative writers writing locative fiction revealed

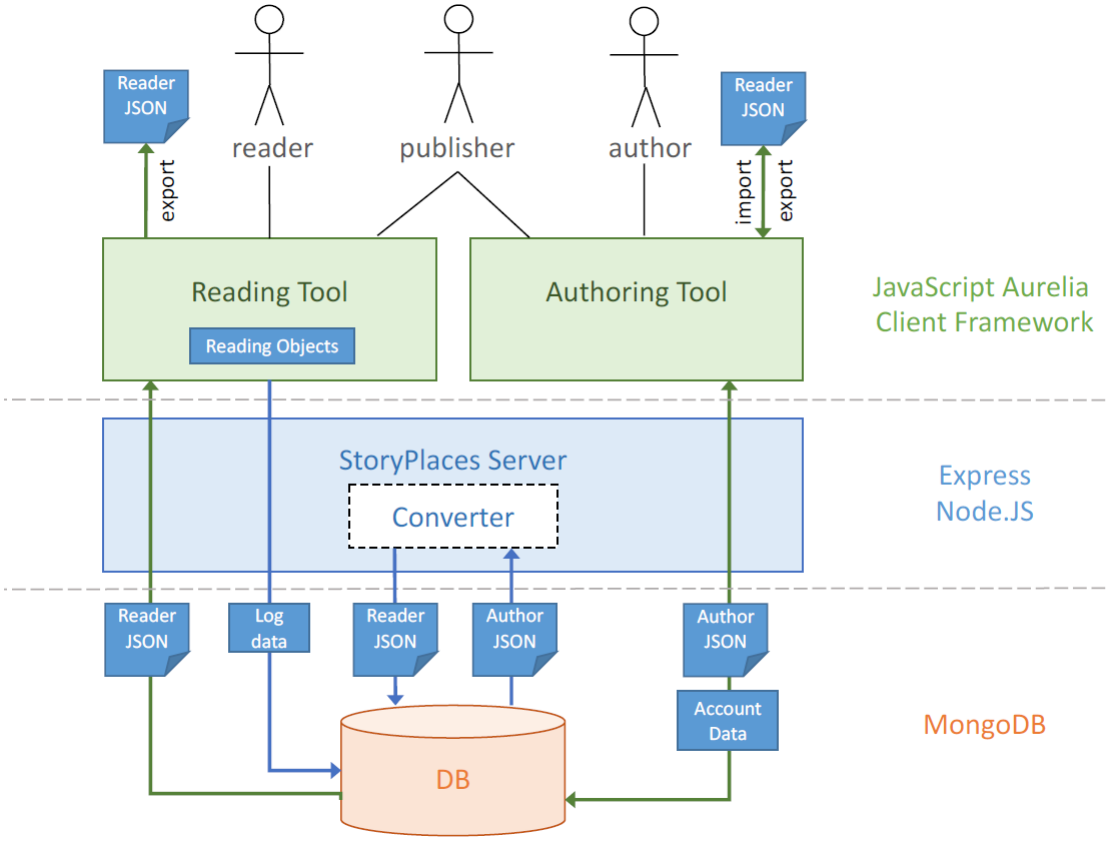

Figure 1: The Storyplaces Architecture, User Interaction, and Data Flow

a number of common patterns and also confirmed the suitability of sculptural hypertext as an underlying model [20]. Following this the team worked with both engineers and writers in a participatory design process to create interface components and explore the requirements from both a technical and artistic point of view. As part of this process, and in the spirit of co-operative inquiry [31], we also personally explored the process of authoring a locative narrative itself [23]. The direct experience of working with the form provided insights into the challenges authors face, which aided later workshops with authors, leading us to create a toolkit of advice for authors to help them balance the demands of story, structure, and landscape [29].

The research and design process went in parallel with iterative development on the technology. First in the form of a server and web-based reading client, used for the project's public deployments, and then later including an authoring tool that encapsulated the lessons learned about designing locative narratives. These are detailed in the next Section.

4 THE STORYPLACES PLATFORM

The StoryPlaces platform is an open source generic platform for delivering locative hypertext. Its development followed, and builds upon, the CDP locative narrative model [25]. The framework is made up of three individual applications: *reading-tool*, *authoring-tool*, and *server*.

Architecture

While the *server* functions as host, storage, and validator for content, the *reading-tool* and *authoring-tool* provide the interface and control to content for different modes of usage. The *authoring-tool* is for the management, creation, and editing of stories, while the *reading-tool* is for reading published content. Authors authenticate with the *server* using a Google account. This allows their stories to be stored centrally by the server and potentially published to others. Readers are anonymous, so their reading state (stories started and progress) are held locally within their web client using cookies and only uploaded to the server for analysis purposes. A sequence of usage for StoryPlaces might be summarised as follows:

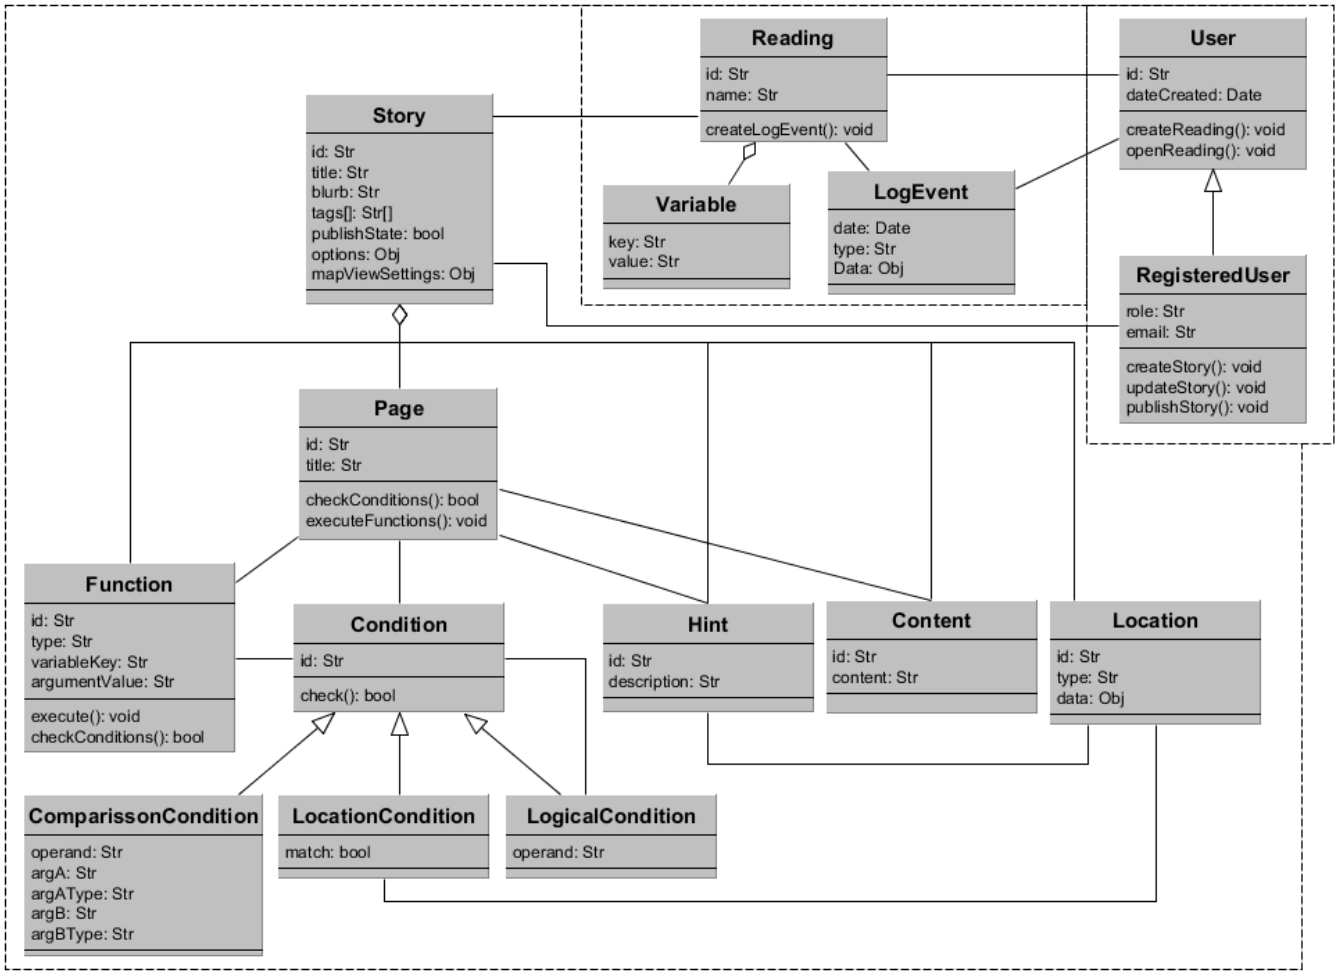

Figure 2: The Storyplaces Schema as a class diagram depicting the story model, runtime elements, and user elements

(1) **Author** creates a story using the *authoring-tool* which is stored on the *server*
(2) **Author** uses the *authoring-tool* to create a preview of the story and tests it using the *reading-tool*
(3) **Author** refines their story using the *authoring-tool* and marks it as ready for publication; updates *server*
(4) **Publisher** reviews the story and may publish through the *authoring-tool*; updates *server*
(5) **Reader** encounters the story on the *reading-tool* which loads from the *server* and selects it for reading
(6) **Reader** creates a Reading for the story which is held locally in their instance of the *reading-tool*
(7) **Reader** opens the reading and reads the story, their reading state is stored locally in the *reading-tool*, and usage data is logged on the *server*

The architecture of the StoryPlaces platform, and the flow of data between its parts, is summarised in Figure 1. The authoring tool uses an extended version of the Story JSON format, which includes additional information used to scaffold the creation of the stories (for example, information used to support pattern based authoring - an approach discussed in more depth in previous work [24]). When the story is published (either in a limited way by the author for testing, or in a more permanent way by the publisher for others to discover) a converter within the server transforms the Author JSON into Reader JSON, which can then be served to the Reading Tool. Openness was an important design goal, and so the author allows Reader JSON to be imported/exported using the authoring tool, and the reading tool also allows Reader JSON to be exported.

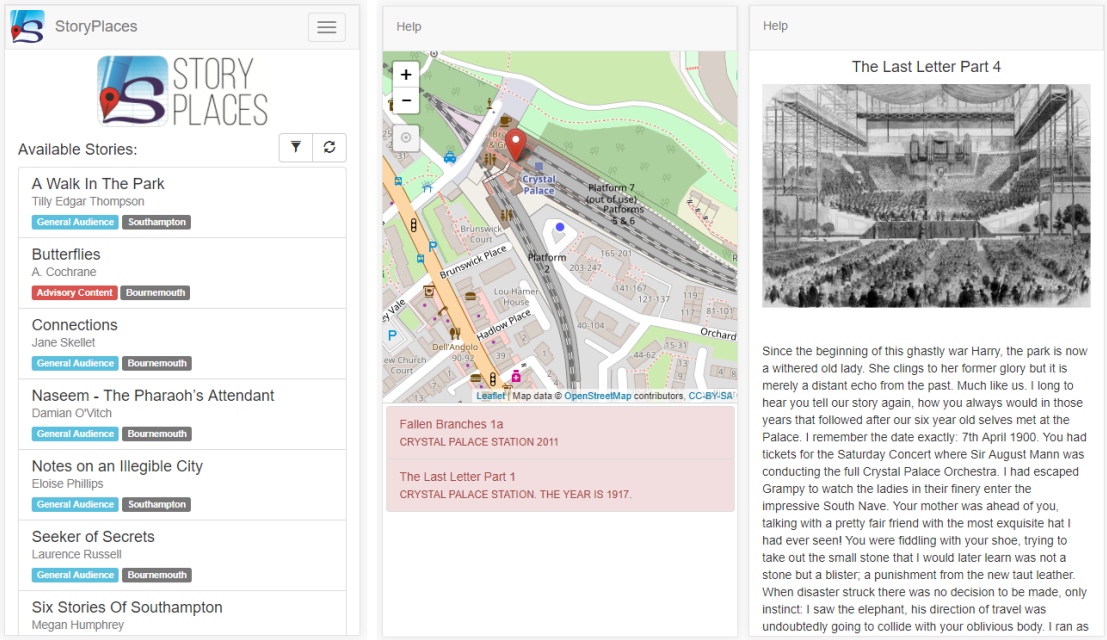

Figure 3: Three screenshots of the StoryPlaces reading tool displaying (from left to right) selecting a story, navigating to a location during a reading, and reading a page

StoryPlaces is powered by a flexible Sculptural Hypertext engine that resides in the *reading-tool*. Sculptural Hypertext can be compared to a state machine defined by pages, rules, and constraints. Consequently the StoryPlaces **Story** object is modeled in terms of *Page*s with *Condition*s that determine their visibility, and *Function*s which may be executed by pages to change the state, thus changing page visibility.

The StoryPlaces **Reading** object stores the current state information, which serves as the instance of a users specific reading of a particular story. Readings are made up of a reference both to **User** and **Story**, and contain a list of *Variable*s which are key-value pairs of strings read by *Condition*s and written to by *Function*s. Finally **User** and **LogEvent** complete the first order objects, handling user role and identity, and logging of actions within the system.

The *Condition* object was separated into three types: *ComparisonCondition*, *LocationCondition*, and *LogicalCondition*. *ComparisonCondition*s handle standard Sculptural Hypertext constraints - comparing variables. *LocationCondition*s serve the inherently locative aspect of the medium the system works with, this is a contextually aware constraint that utilises a *Location* object, and considers whether the user currently satisfies this based on data and sensors available. At present StoryPlaces defines location within or outside of a defined GPS centered circle. However, the *Location* object is extensible. Finally *LogicalCondition* enabled the construction of more complex logical statements by taking reference to other *Condition*s and a defined logical function such as AND or OR. Through the use of all three types the engine can support sculptural hypertext with largely limitless complexity in structure through logical hierarchies of variable comparisons and location checks. While *Page*s make reference to *Condition*s in a way typical of sculptural hypertext, StoryPlaces also allows the *Function*s on *Page*s to have *Condition*s controlling their execution and the changes of state. The combination of conditional functions and logical hierarchies of conditions serves to provide the sculptural hypertext engine with a fully functioning state machine.

Implementation

*Page*s also reference *Hint* and *Content* objects. Locative Narrative is unlike most hypertexts in that it demands its reader satisfies location conditions during the reading. The Reading Tool therefore can give readers guidance on where potential pages are; these are pages where the comparison and logical conditions are met, but the location conditions are not. The reader can then move to that location to open and read the page. The *Hint* object describes how these potential pages should be described to the user. StoryPlaces currently supports textual directions and points on a map (the points are defined using their own Location object, and therefore need not match the Location conditions themselves). *Content* objects deals with the actual content of a page, specifying its text and any multimedia (currently images, audio, and video

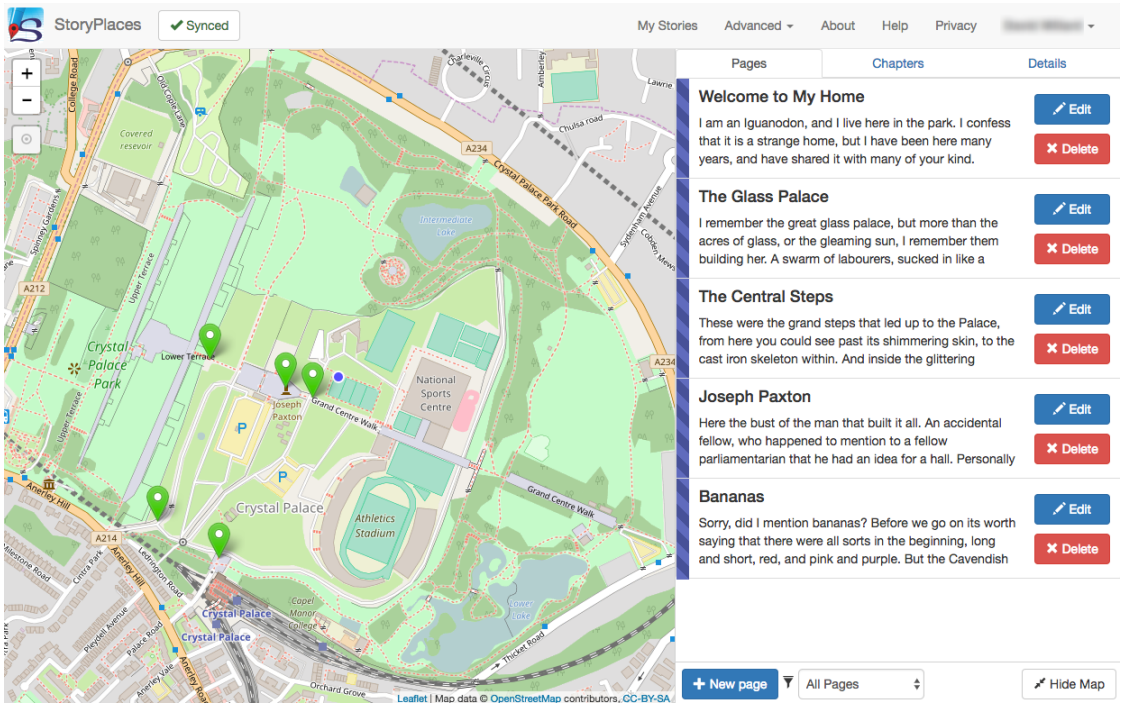

Figure 4: The StoryPlaces authoring tool (username blurred)

are supported). StoryPlaces is therefore unusual in that it does not view a node (a Page) as a wrapper for content, but rather as a relationship object that binds together content, conditions, functions, and hints. This will be explored further in Section 5. Beyond this structural schema many objects also include more mundane meta data such as titles for *Page* and **Story**, blurb and tags for **Story**, and name for **Reading**. The entirety of our Sculptural Hypertext engine schema is summarized in a slightly simplified class diagram in figure 2.

The StoryPlaces Platform has been implemented as a node.js web application with the reading tool and authoring tool implemented in the Aurelia JavaScript client framework. Web applications are able to access hardware such as the camera and GPS on mobile devices while ensuring a degree of platform neutrality between Android, iOS, and any other mobile operating system outside of browser compatibility. Web applications also avoid the barrier of working with OS application distribution services such as the app store, and allow StoryPlaces to be accessed on both mobile and more conventional devices through the use of responsive CSS. This is especially useful as both tools allow location to be spoofed, so stories can be authored and tested on a desktop machine as well as in-situ.

The StoryPlaces server is implemented using node.js, express.js, a mongoose schema, and operates a mongo database for its objects. It exposes an API which allows its client applications CRUD access to its objects. The reader-tool and authoring-tool client side applications are implemented principally in aurelia.js and make use of open street map for their map based applications. Figure 3 displays screenshots of the implemented reader-tool on a mobile device, and Figure 4 shows the authoring tool on a desktop device. StoryPlaces is fully open source and available under the MIT license.

5 HYPERTEXT SYSTEM INNOVATIONS

StoryPlaces is unusual as a generic web-based hypertext system that is focused on location. The platform therefore demonstrates several innovations in both its treatment of location within its hypertext model, and in the power of its sculptural model and engine.

Building upon the Sculptural Premise

Nearly two decades ago the idea of sculptural hypertext sought to demonstrate that hypertext could arise out of a system that "sculpted away" links at runtime [4, 35], similarly to the way that 'guard fields' controlled the availability of links in the StorySpace system - in essence a sculptural hypertext is one where every node is linked to every other, but each of those links has guard fields. The latest version of StorySpace (version 3) contains a sculptural engine, and extends the power of guard fields to include nested logical structures,

such as '(visited(A) & unvisited(B)) | clicked(Anne)' to mean true if the reader has visited node A but not B, or if they have previously clicked the word Anne [6].

StoryPlaces hierarchy of constraints has similar expressive power, allowing *arbitrary nested constraints* comprised of logical or comparative statements. It also includes locative constraints, meaning it is possible to mix logical and location requirements. For example, to express that a page is visible if the reader is in the garden, or if they have previously met Anne and are in the driveway. The use of contextual constraints extends to time, and StoryPlaces supports *time-based constraints* in the form of checks against

time frames that are either absolute (e.g. can only be accessed from 11.00-13.00) or relative to user action (e.g. only available 1 hour after reading node 6). Time and location are two examples of contextual constraints that rely on information beyond the system itself, other examples - such as weather or social media context - would be possible, but remain as future work.

Previous sculptural hypertext systems have often made constraints a part of the page itself (for example as an embedded statement in a domain specific language), but StoryPlaces keeps its conditions and functions outside of the page where they are referenced by the page via ID. Thus in StoryPlaces complex constraints are in fact an explicit named hierarchy of JSON objects. This makes reusing complex constraints easy, as all constraints are first class and can be referenced, but does make creating complex constraints by hand cumbersome (although this is easily hidden behind a GUI interface in the authoring tool).

One of the things that we have observed when working with StoryPlaces is that these constraints encourage a much more complex structure than a normal branching narrative. For example, we can begin to see intricate networks of dependent nodes in two of the larger StoryPlaces stories: Isle of Brine [23] and Fallen Branches [29], that would be cumbersome to create using calligraphic models.

StoryPlace's functions allow for the simple manipulation of variables (via set and increment), the system also supports *conditional functions*, which are functions with their own constraints that limit when they are activated. This allows for a greater degree of author control over state changes and avoids the need for duplicated pages with differing functions for different conditions. For example, imagine that we want to move to Act 2 once all five of our main characters have been encountered, regardless of the order in which they were met; using conditional functions we can do this by creating a function that increments the Act, and placing a constraint on it that is a check that character count is 5. We can now attach this function to every page in Act 1, and it will only trigger once the condition is met. Without conditional functions we would have had to duplicate the five pages that introduce the main characters so that there is a second version that replaces the first when the other four have been met, and which changes the Act.

Locative Demands

Sculptural hypertext lends itself to locative narrative as it more easy supports open structures with many possible navigational paths, which is a good match for readers navigating a physical environment [20, 25]. But, location does make certain demands of the hypertext model.

We have already described how *LocationCondition*s serve to expand the scope of hypertext state from the text itself to the users context. LocationConditions reference a Location object which defines an area. In StoryPlaces these are *first class locations* detached from their conditions and referenced by ID. This means that a story includes a set of locations that are defined independently of the content and logical rules of the story, these locations could therefore be easy swapped for alternatives, instantly remapping a story onto a new area. At present this is not supported by the reader, but can be done directly with the JSON story definitions.

StoryPlaces also uses *first class content*, also referenced by ID. This is because it is sometimes conceptually easier to model different states as different nodes that share the same content, rather than as a single node with very complex constraints. This was particularly useful for Isle of Brine [23] where a thread of narrative was replicated in three Acts set in entirely different areas, but was only visible to the reader in their current Act. In this case three nodes were created to account for the three states in the three Acts but only a single content object was necessary.

Unlike a traditional hypertext where a user selects a link on the screen, navigation in StoryPlaces is done through a combination of screen and physical navigation. Thus guidance is necessary to help the user find locations in their environment to trigger new content. Although StoryPlaces uses a map, this is theoretically possible in a number of ways, and therefore StoryPlaces also *separates hints from locations*, using the Hint object to direct readers rather the location information attached to the constraints. StoryPlaces hints can contain both a textual direction (that is shown on the screen) and a GPS location (that is shown on the map), but modeling these explicitly means that they are used at the authors discretion and it is perfectly possible to have a node with no hint, or a hint with the description or the location missing. To our knowledge StoryPlaces is the first locative system to have separated direction hints from actual location and given authors this control. It also creates a conceptual

place in the model where more sophisticated hint information could be placed in future iterations of the software (such as audio cues, or AR models).

Finally, as a system designed to be used in-the-wild and on-the-move StoryPlaces cannot make the same sort of assumptions about users' connectivity as more traditional desktop hypertext systems. Consequently StoryPlaces undertakes *client-side caching for disconnection resilience*, creating a local cache of the story, content and usage data for later synchronization in the event of loss of connectivity. This is applied in the authoring tool as well as the reader, which will hold a queue of changes locally and resync with the server once connectivity has been restored. Browser based local storage limits prevent StoryPlaces from caching all media, however the system is typically able to cache the story structure, text, and images client side. This does not include audio and video due to the storage limits, which represents a future challenge for these types of web-based application.

6 CONCLUSIONS AND FUTURE WORK

In this paper we have presented the StoryPlaces Platform - a generic web-based locative hypertext system with a powerful Sculptural Hypertext engine. StoryPlaces is the first general platform for locative narrative, and a rare example of a new and novel hypertext system [2].

The StoryPlaces Platform has been used to create and deliver 21 different stories in various locations across the UK (6 in Southampton, 7 in Bournemouth, 6 in London, 1 in Exeter, and 1 in Tiree). The majority (18) of these stories were released at launch events in their host cities where they were publicly advertised through local partners and attended by the general public. The stories include linear works, branching narratives, open narratives, short stories of 20 minutes and long stories of 4 hours, text based works, audio works, and mixed media works. Our readers have been a mixture of academics, reading group members, and the general public. Our authors have included 2 members of the research team, 3 commissioned professional authors, and 17 volunteer writers with a mixture of professional, academic and amateur backgrounds.

The StoryPlaces Platform is open source software comprised of two web-based clients (the *reading-tool*, and *authoring-tool*) implemented in the JavaScript Aurelia framework, that communicate with a Node.JS *server*, using the JSON data format, and MongoDB for storage. The clients use responsive CSS to render to both mobile and desktop environments.

The StoryPlaces Platform demonstrates both an evolution of Sculptural Hypertext, and an innovative response to the demands of locative narrative. Specifically its sculptural engine supports arbitrarily complex (nested) conditions, conditional functions, and contextual constraints (in the form of temporal and location conditions). To deal with locations the hypertext model separates locations from pages (potentially allowing remapping of stories to new spaces), separates hints from locations (meaning that directions to locations are independent of the definition of those locations), and separates content from pages (treating pages, or nodes, as a mapping between content, constraints, hints, and functions, rather than as a wrapper for content). The system also supports runtime caching for both reading and authoring tools in order to provide disconnection resilience.

Our Future work includes both software developments and new locative applications. Work is currently underway to use the framework to recreate a number of classical locative experiences from the research literature whose own bespoke systems have fallen out of maintenance and functionality - this includes 'The Chawton House Project' [34], and 'iLand of Madeira' [15] with permission from the original authors.

Further work is also required to explore alternative interfaces for reading, and understand their impact on the reader experience. For example, StoryPlaces' current reliance on screen-based navigation takes the reader's attention, distracting them from their surroundings. It's possible that an audio based interface may provide a less demanding alternative that would allow users to spend more time looking at the locations themselves rather than their device. We are also considering options that would allow for the platform to support location in other ways, including classifications of locations that are not tied to one specific GPS point, allowing stories to be mapped dynamically to new areas as required. Finally while locative narrative (and a majority of hypertext) is designed as an individual experience we are also exploring how the StoryPlaces' sculptural engine might be extended to support multi-participant narratives, where sculptural state is shared between several readers, meaning that a readers' experience is affected by the choices of others [33].

Our hope is that the StoryPlaces Platform will lower the barriers for authors who want to experiment with locative literature, and act as an extensible software platform that can be used by researchers and developers for further experimentation into what is possible with locative storytelling.

7 ACKNOWLEDGMENTS

This work was undertaken as part of the StoryPlaces project funded by The Leverhulme Trust (RPG-2014-388). We would like to acknowledge the assistance of our co-designers and developers who have contributed to the StoryPlaces project: Verity Hunt, Heather Packer, Petros Papadopoulos, Yvonne Howard, Will May, Phillip Hoare, James Jordan, Patrick McSweeney, Andrew Day, Kevin Puplett, Callum Spawforth, Victoria Dawson, James Cole, and Katie Lyons.

REFERENCES

[1] C Ardito, P Buono, M.F Costabile, R Lanzilotti, and T. Pederson. 2007. Mobile games to foster the learning of history at archaeological sites. *Visual Languages and Human-Centric Computing* (2007).

[2] Claus Atzenbeck, Thomas Schedel, Manolis Tzagarakis, Daniel Roßner, and Lucas Mages. 2017. Revisiting Hypertext Infrastructure. In *Proceedings of the 28th ACM Conference on Hypertext and Social Media*. ACM, 35–44.

[3] Mark Bernstein. 1998. Patterns of Hypertext. In *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia*. ACM, New York, NY, USA, 21–29.

[4] Mark Bernstein. 2001. Card shark and thespis: exotic tools for hypertext narrative. In *Proceedings of the twelfth ACM conference on Hypertext and Hypermedia*.

[5] Mark Bernstein. 2002. Storyspace 1. In *Proceedings of the thirteenth ACM conference on Hypertext and hypermedia*. ACM, 172–181.

[6] Mark Bernstein. 2016. Storyspace 3. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media*. ACM, New York, NY, USA.

[7] M. Blythe, J. Reid, P. Wright, and E. Geelhoed. 2006. Interdisciplinary criticism: analysing the experience of riot! a location-sensitive digital narrative. *Behaviour & Information Technology* 25, 2 (2006), 127–139.

[8] Jay David Bolter and Michael Joyce. 1987. Hypertext and creative writing. In *Proceedings of the ACM conference on Hypertext*. ACM, 41–50.

[9] J. Broadbent and P. Marti. 1997. Location aware mobile interactive guides: usability issues. In *Proceedings of the Fourth International Conference on Hypermedia and Interactivity in Museums*. 162–172.

[10] Peter Brusilovsky, John Eklund, and Elmar Schwarz. 1998. Web-based education for all: a tool for development adaptive courseware. *Computer Networks and ISDN Systems* 30, 1-7 (1998), 291–300.

[11] B.S. Bunting, J. Hughes, and T. Hetland. 2012. The Player as Author: Exploring the Effects of Mobile Gaming and the Location-Aware Interface on Storytelling. *Future Internet* 4, 1 (2012), 142–160.

[12] Vannevar Bush and others. 1945. As we may think. *The atlantic monthly* 176, 1 (1945), 101–108.

[13] N Davies, K Cheverst, K Mitchell, and A Efrat. 2001. Using and determining location in a context-sensitive tour guide. *Computer* 34, 8 (2001), 35–41.

[14] Paul De Bra, David Smits, and Natalia Stash. 2006. The design of AHA!. In *Proceedings of the seventeenth conference on Hypertext and hypermedia*. ACM, 133–134.

[15] Mara Dionisio, Valentina Nisi, and Jos P. Van Leeuwen. 2010. The iLand of Madeira Location Aware Multimedia Stories. In *Proceedings of the Third Joint Conference on Interactive Digital Storytelling (ICIDS'10)*. Springer-Verlag, Berlin, Heidelberg, 147–152.

[16] Jonathan Dovey. 2016. Ambient literature: Writing probability. *Ubiquitous Computing, Complexity and Culture* (2016), 141–154.

[17] Mihai Duguleana, Raffaello Brodi, Florin Girbacia, Cristian Postelnicu, Octavian Machidon, and Marcello Carrozzino. 2016. Time-Travelling with Mobile Augmented Reality: A Case Study on the Piazza dei Miracoli. In *Digital Heritage. Progress in Cultural Heritage: Documentation, Preservation, and Protection*, Marinos Ioannides, Eleanor Fink, Antonia Moropoulou, Monika Hagedorn-Saupe, Antonella Fresa, Gunnar Liestøl, Vlatka Rajcic, and Pierre Grussenmeyer (Eds.). Springer International Publishing, Cham, 902–912.

[18] Andrew M Fountain, Wendy Hall, Ian Heath, and Hugh C Davis. 1990. MICROCOSM: An Open Model for Hypermedia with Dynamic Linking.. In *ECHT*. 298–311.

[19] Chiara Garau and Emiliano Ilardi. 2014. The âĂIJNon-PlacesâĂİ Meet the âĂIJPlaces:âĂİ Virtual Tours on Smartphones for the Enhancement of Cultural Heritage. *Journal of Urban Technology* 21, 1 (2014), 79–91. DOI:http://dx.doi.org/10.1080/10630732.2014.884384 arXiv:https://doi.org/10.1080/10630732.2014.884384

[20] Charlie Hargood, Verity Hunt, Mark Weal, and David E. Millard. 2016. Patterns of Sculptural Hypertext in Location Based Narratives. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media*. ACM, New York, NY, USA.

[21] Fotis Liarokapis, Panagiotis Petridis, Daniel Andrews, and Sara de Freitas. 2017. *Multimodal Serious Games Technologies for Cultural Heritage*. Springer International Publishing, Cham, 371–392. DOI: http://dx.doi.org/10.1007/978-3-319-49607-8_15

[22] Anders Sundnes Løvlie. 2009. Poetic Augmented Reality: Place-bound Literature in Locative Media. In *Proceedings of the 13th International MindTrek Conference: Everyday Life in the Ubiquitous Era (MindTrek '09)*. ACM, New York, NY, USA, 19–28. DOI:http://dx.doi.org/10.1145/1621841.1621847

[23] David E. Millard and Charlie Hargood. 2017. Tiree tales: a co-operative inquiry into the poetics of location-based narratives. In *Proceedings of the 28th ACM Conference on Hypertext and Social Media*.

[24] David E Millard, Charlie Hargood, Yvonne Howard, and Heather Packer. 2017. The StoryPlaces Authoring Tool: Pattern Centric Authoring. In *Authoring for Interactive Storytelling 2017 Workshop @ ICIDS 2017*.

[25] David E. Millard, Charlie Hargood, Michael O. Jewell, and Mark J. Weal. 2013. Canyons, Deltas and Plains: Towards a Unified Sculptural Model of Location-based Hypertext. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media*. ACM, New York, NY, USA, 109–118.

[26] K. Naliuka, T. Carrigy, N. Paterson, and M. Haahr. 2010. A narrative architecture for story-driven location-based mobile games. In *New Horizons in Web-Based Learning*. Springer, 11–20.

[27] Theodor Holm Nelson. 1999. The unfinished revolution and Xanadu. *ACM Computing Surveys (CSUR)* 31, 4es (1999), 37.

[28] V. Nisi, I. Oakley, and M. Haahr. 2008. Location-aware multimedia stories: turning spaces into places. *Universidade Católica Portuguesa* (2008), 72–93.

[29] Heather S Packer, Charlie Hargood, Yvonne Howard, Petros Papadopoulos, and David E Millard. 2017. Developing a WriterâĂŹs Toolkit for Interactive Locative Storytelling. In *International Conference on Interactive Digital Storytelling*. Springer, 63–74.

[30] F. Pittarello. 2011. Designing a context-aware architecture for emotionally engaging mobile storytelling. *IFIP Conference on Human-Computer Interaction* (2011), 144–151.

[31] Peter Reason and Hilary Bradbury. 2005. *Handbook of action research: Concise paperback edition*. Sage.

[32] Ted Smith and Steve Bernhardt. 1988. Expectations and experiences with HyperCard: a pilot study. In *Proceedings of the 6th annual international conference on Systems documentation*. ACM, 47–56.

[33] Callum Spawforth and David Millard. 2017. A framework for multi-participant narratives based on multiplayer game interactions. In *International Conference on Interactive Digital Storytelling*. Springer. https://eprints.soton.ac.uk/414445/

[34] M.J. Weal, D. Cruickshank, D.T. Michaelides, D.E. Millard, D.C.D. Roure, K. Howland, and G. Fitzpatrick. 2007. A card based metaphor for organising pervasive educational experiences. In *Pervasive Computing and Communications Workshops, 2007*. IEEE, 165–170.

[35] Mark J. Weal, David E. Millard, Danius T. Michaelides, and David C. De Roure. 2001. Building Narrative Structures Using Context Based Linking. In *In Hypertext '01. Proceddings of the Twelfth ACM conference on Hypertext, Aarhus, Denmark*. 37–38.