

The Storytelling Systems of Los Atlantis



Figure 1: From the prototype performance of *Los Atlantis* at the UCLA School of Theater, Film and Television.

Jeff Burke

Peter Gusev

Zoe Sandoval

Jared J. Stein

Zhehao Wang

UCLA Center for Research in Engineering, Media and Performance
(UCLA REMAP)

Los Angeles, CA 90095, USA

{jburke, peter, zoe, jared, zhehao}@remap.ucla.edu

Abstract

UPDATED—February 14, 2017. *Los Atlantis* was a multisite live performance created within a research effort exploring the simultaneous authorship of story and code. This case study describes the systems and interfaces for the authors and production team. The piece’s story followed its characters as they explored the futuristic Archive of a historical city while accompanied by both physical and online audiences. Its systems aimed to make the contemporary city of the

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

CHI'17 Extended Abstracts, May 06 - 11, 2017, Denver, CO, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4656-6/17/05...\$15.00

DOI: <http://dx.doi.org/10.1145/3027063.3053342>

production—in this case, Los Angeles—the location of the historical city and Archive by dynamically integrating local media with the performances’ actions and settings. Subsystems included an “active script” combining control code and traditional text, a distributed media repository on YouTube that fed real-time video manipulation and projection, a media gathering task assignment system, a web interface and separate mobile guide for the audience, and the necessary control and storage mechanisms. The case study concludes by briefly discussing successes and limitations of the process, and research challenges in authoring systems for related future work on *enacted stories*.

Author Keywords

Live performance; theater; film; algorithmic performance; authoring interfaces; dramatic authorship.

ACM Classification Keywords

J.5 ARTS AND HUMANITIES; H.5.2 User Interfaces; H.5.4 Hypertext/Hypermedia; D.2.6 Programming Environments.

Introduction

Los Atlantis was an experimental performance at the UCLA School of Theater, Film and Television in spring 2015. Created over the course of an academic year, the



Figure 2: Media was gathered by the cast and crew members for UCLA's prototype performance of *Los Atlantis*.

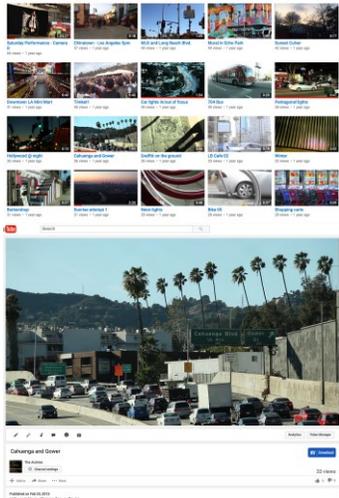


Figure 3: Sample views of the Archive YouTube Channel, containing media from Los Angeles used in the prototype performance.

piece integrated a number of contemporary technology concepts into a multisite theatrical event shared with the public as a prototype production. The underlying concepts included 1) an *active script*, a dynamic document serving a traditional script's purpose in rehearsal, but also as an organizational tool for real-time decision making with code, and as a control interface for technicians during the performance; 2) the use of a *networked media repository* (YouTube) for authored and audience-generated media assets, integrated with a real-time digital signal processing engine; 3) an evolving *show memory* database of information about the audiences and activity specific to each performance; and 4) *data-centric control* based on finite state machines and the use of the Named Data Networking (NDN) future Internet architecture to communicate. [5] Additionally, the performance included an on-site mobile interface (with a bot-like character) and a website for uploading user-generated content and remote viewing of multiple live streams.

Los Atlantis was developed by UCLA REMAP as part of a larger, multiyear research project and corresponding course, to explore the simultaneous authorship of *story* and *code*, in particular to support new relationships among audiences, performers, and authors enabled by contemporary digital technology. First, this case study describes the piece's concept, then briefly introduces subsystems of the performance, each as motivated by a key design goal. It concludes with a short discussion of code-related aspects of the performance itself and of lessons learned and research opportunities for HCI to support authoring at the intersection of story and code.

Concept

Inspired by a quote from Cicero—"The life of the dead is placed on the memories of the living."—the story of *Los Atlantis* followed a small band of characters as they explored the futuristic Archive of a historical city while accompanied by audiences. Audience members were participants in the story and interacted with interfaces of the Archive, both in-person and online. The Archive manifested itself within a physical performance, a YouTube channel (Figure 3), a mobile website for the in-person audience ("the Guide"), and a website for the online audience called the "Observatory."

During each performance of *Los Atlantis*, the audience was guided through a series of live performance vignettes by performers and a mobile website. (See Figure 4.) Their participation further expanded the Archive's holdings via photos and videos requested and captured on their mobile devices. *The systems aimed to make the contemporary city of the performance the basis of the historical city encountered in the performance.* This was achieved by dynamically integrating the vignettes with local media. The production discussed here was situated in Los Angeles, and so used media gathered in that city and its surroundings as the Archive's holdings. The audience could view and contribute media online before and during the performance, and was encouraged to do so in some vignettes. Each audience further contributed to the Archive's media, which was then used within subsequent performances and online. The real world became the historical world of the story.

The underlying concept and themes were conceived to explore how code can be authored alongside traditional text to incorporate contextual information and media

assets provided by performers and audience members. The process expanded on previous work that used Google Glass to provide text dynamically to actors and audience through both algorithmic selection and human-in-the-loop control of writers and directors. [2]



Figure 4: Audience led through different sections of the *Los Atlantis* performance at UCLA.

Subsystems

This section briefly introduces the major technical components of *Los Atlantis*, from the perspective of key research and design goals. The article's authors recommend reading it after viewing the supplemental video.

Active Script

Motivating goal: Develop a new form of script itself, one that is reasonable for writers and coders to use to collaborate.

To support the writing of dynamic, context-dependent story elements in a format reasonably familiar to the creators, cast and crew, we developed a new markdown format, Fountain+. It extended the existing Fountain markdown [4] for writing, editing and sharing screenplays. Fountain+ enabled explicit identification of characters, character types, and reference to dynamic

Javascript components. From the markdown text, the parser could generate PDF files using snapshots of dynamic components, or HTML5/Javascript documents containing changing, data-driven elements with the ability to control the performance. A mapping from the markdown to JavaScript Web Components, implementing dynamic functions, enabled authors to write script elements that updated values, sent commands, or read data from other subsystems.

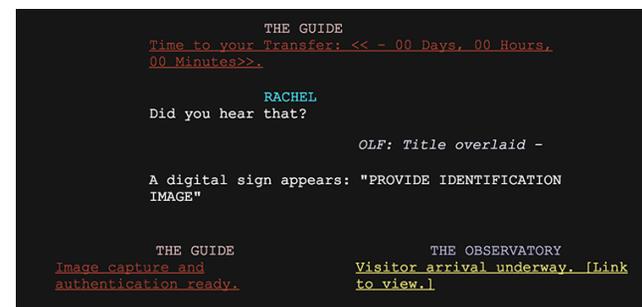


Figure 5: Active script as rendered in a web page, showing dynamic text and dialogue sent to mobile / web elements by an operator's mouse click.

The resulting "active script" evolved to serve three purposes: 1) a dynamically updating version of a traditional script for rehearsal, 2) an organizational tool for code used in real-time decision making, and 3) an interface for the crew to control the show while it was running. (See Figure 5.)

The piece's authors wrote the original text in the popular Final Draft and Scrivener scriptwriting tools, exported it into Fountain format, and then adjusted the markdown. The Fountain+ parser then generated live HTML pages for the crew to control media in the



Figure 6: Internal media gathering task assignment web interface.

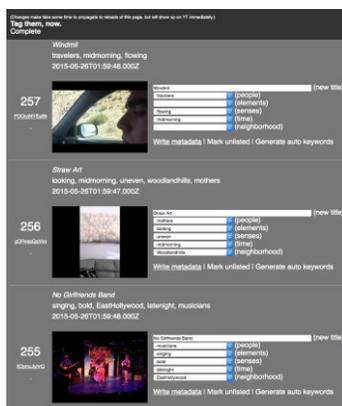


Figure 7: Internal media tagging web interface.

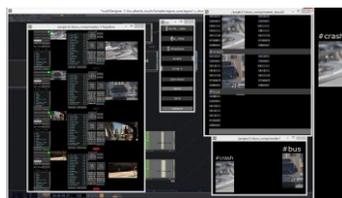


Figure 8: View of TouchDesigner network for loading YouTube media dynamically from the network and manipulating video and metadata in real time.

performance, as well as a PDF of the script, which was given to actors and directors for rehearsal.

For example, dialogue from *The Guide*, formatted as any character's dialogue would be in Fountain, generated a hyperlink that sent a customized message to each audience member's mobile device, e.g., "Hello Miranda!" because of a "type" associated with that character name. These dynamic elements were implemented as Web Components inserted into the HTML script by the Fountain+ parser when it encountered the appropriate markdown. This also enabled hooks into the Observatory and cues for external media. Markdown text—such as "VQ #40: fade in sunrise video"—would generate hyperlinks to trigger corresponding actions. (Please see supplemental material.)

Streaming Media Repository

Motivating goal: Find a scalable, networked approach to providing a media library for live experimentation.

To support the ongoing upload of media to the Archive and its playback anywhere in the performance's multiple physical sites and configurations, YouTube was employed as the backing store. This unusual choice for live performance traded off dependence on networked services to enable 1) any playback machine to access the full corpus of media without having to download first; 2) web API-based access to the entire corpus; 3) the use of YouTube's upload, storage, and transcoding tools and its other benefits. In practice, the Archive's YouTube channel became an ongoing, public facing media repository for *Los Atlantis*, where all media was uploaded, curated and stored. (See Figure 3.)

Media Gathering & Tagging

Motivating goal: Use folksonomies (tagging) and other approaches familiar to savvy users of social media, to engage performers and audiences with media selection.

To ensure videos covered the story's needs, we created a website for assigning media gathering tasks (Figure 6). Task elements—such as location, time of day, shot (e.g., close-up, low-angle, and out-of-focus), subject, and recommended length—entered into a Google Sheet by the show's authors were then assigned automatically to other participants. (This approach, designed to encourage the cast and crew to see their familiar city anew as they build the media for the Archive, was inspired by the Situationist's explorations of the urban environment. [2]) Assignments were prioritized according to the needs of the script. These tasks guided the performers (and eventually, audience, if interested) to go to different parts of the city at different times, and record people, places and things happening at that time. Media was uploaded to the Archive (in the story) and to the YouTube channel (in the real world).

As a bridge between the story world and the real world, and to limit the organizational burden on participants of adding media, a story-specific folksonomy of tags was created. The audience and performers were able to tag the media when uploaded into YouTube, associating their media with hashtagged concepts from the story. A custom webpage, the Tagger (Figure 7), enabled the authors to preview and (re)tag many videos quickly using the folksonomy.

Real-Time Media Manipulation

Motivating goal: Provide flexibility and high production value for creating the visual environment.

While YouTube provided a robust, distributed store for the show's media, it did not directly provide the real-time media manipulation needed for rehearsal and performances, from speed shifts to color correction, cropping, 3D positioning, etc. So, the team integrated access to YouTube videos into Derivative's TouchDesigner, a platform that provides extensive real-time video processing capabilities for 2D and 3D media. We developed an open-source plugin to enable streaming of any YouTube video into a TouchDesigner processing network simply by providing its URL. (See Figure 8.) Videos were fetched and projected in the space in what was effectively real time. To keep playback buffering to a second or two, short enough to be triggered "just in time" before playback was needed, we used a 10G high-speed experimental network provided by UCLA NOC.

Access to *any* YouTube video, not only those gathered for the Archive, on any machine used for projection, became a powerful and unusual capability for prototyping and rehearsal. Commonly used playback, filtering, image adjustment and compositing elements could be included via markdown written by non-programmers.

Ananke (Control)

Motivating goal: Prototype a distributed control system for live, media-rich, networked performances.

To control the various instances of the subsystems—typically five independent video projectors on four machines controlled by up to six active script interfaces—we developed a simple distributed control system, Ananke, using NDN. [5] Since 2005, REMAP has explored the development of middleware for

interactive systems, e.g., in Kolo and Nebesko, a Java-based middleware and corresponding state machine, respectively, which used hierarchical application-defined names for data exchange. Using hierarchical names to describe all entities in a system provides a consistent and understandable way to implement control and data access. NDN forwards such named data packets directly at the network layer, thus providing several potential advantages over traditional middleware for scaling and robustness.¹

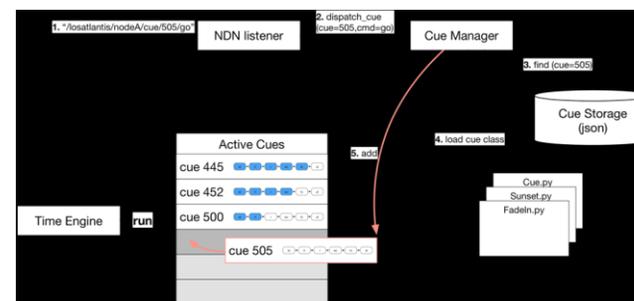


Figure 9: Media event (cue) execution within each instance of the stage machine.

A complete discussion is outside the scope of this paper, but a simple illustration of NDN's utility is that both control messages (used for triggering pre-built sequences) and real-time adjustment of TouchDesigner parameters (used for experimentation in rehearsal) could be represented consistently using named data and basic NDN primitives. For example, as shown in the supplemental figures, each TouchDesigner node listened for incoming data that corresponds to its name, e.g., /losatlantis/nodeA. Data names under this prefix included: 1) a media event, or "cue" namespace,

¹ Please see <http://named-data.net> for more information.



Figure 10: Audience website, story background page.



Figure 11: Mobile guide for in-person audience members.



Figure 12: Live streaming interface for the online audience.

through which the node receives state transition commands, and 2) a control namespace directly referring to elements in the TouchDesigner media processing network, which internally represents elements as a hierarchically structured set of named nodes for filtering, playback, and so on.

To handle the cue logic, a state machine extending a traditional theatrical control paradigm ran on each machine that needed it. The system supported any number of overlapping and simultaneously running cues on any number of projection machines, triggered independently or together. (See Figure 9.) Each cue was represented in the system as a Python class with six distinct states: wait in, fade in, run, wait out, fade out and auto follow (see Figure 1). Typical cue types could be customized by changing Python class properties that corresponded to transition durations, image adjustments, etc. via the network. In a few cases, complex behaviors were needed (e.g., generating “floating” images from media supplied by the audience). These used custom Python classes.

Show memory

Motivating goal: Like the Archive itself, the piece’s systems should have a memory of past performances.

In addition to the media of the Archive, the *Show Memory*, implemented as an Apache Cassandra database, recorded state that could be used by the active script or other subsystems. The Show Memory tracked 1) *participant demographics* for the audience, prospective audience, cast, and crew; 2) *show run* information about each scheduled performance: time, type, status, that cross-referenced to participants; 3) a *media* table with links to the Archive’s YouTube videos,

and metadata such as tags and descriptions associated with each video; 4) a *tag* table that enumerates all tags used by the Archive’s media; 5) a list of *cues*, or media events; and finally 6) an *events log* that stores a running, high granularity list of events, including participant feedback.

Interfaces for the Audience: Website & Mobile Guide

Motivating goals: Create personalized engagement with the audience and bridge the story and real world.

The web served as a bridge between the real world and the story world. A public facing website, the Observatory, handled in-person and online audience registration and profile creation, integrated a media gallery of videos from the YouTube channel, and embedded the YouTube Live broadcast, all within the context of the story. (See Figures 10 and 12.) It also provided a chat room where audience members could interact (via a custom web-based XMPP client) and receive announcements from the Archive.

In addition to the Observatory, in-person audience members gained access to the Guide, a mobile webpage also using XMPP for real-time chat (Figure 11). Audience members checked-in with their mobile phones at the beginning of each show, which matched them to a pre-existing user profile, and began a message exchange with the Guide. From this interface, they could read about the story and its characters, select to follow a character and receive custom messages, and engage in a chat with the Guide itself (operated by crew members).

Standard chat dialogue within the Observatory and the Guide were written into the active script with simple templates for addressing the audience; they were triggered by operators who could also improvise messages as the characters. These messages could include HTML multimedia, such as images and GIFs. Messages and feedback from both the Guide and the Observatory were logged in the Show Memory, and could be referenced later to customize video projections and other show elements.

Los Angeles Performance

The prototype staging involved more than fifty people. The core concept and framework inspired a series of improvisational exercises for the performers, which generated vignettes authored by students within the course and REMAP researchers. Incorporating and creating a through line for these vignettes, a full-length script was authored by a small group, including project director Jeff Burke, actor Israel Lopez, and theater professor J. Ed Araiza, who directed the production. Initial media about L.A. was gathered by these authors, crew and actors in an iterative process influencing the system's rules themselves. (See Figures 2 and 13.)

Audience members first encountered *Los Atlantis* on the web, and were invited within the context of the story to visit the Archive—thus becoming participants in the action and contributors to the story's progression. The audience was asked to contribute media by 1) supplying a photo of themselves upon registration online or on-site, which was later juxtaposed with those of the performers, making the audience characters within the story, and/or 2) uploading video to be added to the Archive's YouTube Channel, to contribute their own history of Los Angeles as described above.



Figure 13: Rehearsal for *Los Atlantis* prototype at UCLA.

Throughout the performance, four rolling projector towers, each with a laptop running TouchDesigner and listening for NDN commands, were used to create the piece's projected media; they were moved by actors and crew from space to space as needed. Video of the performance from four independent cameras, moved from site to site with the projector towers, streamed to YouTube Live, and was then integrated into the audience website (See Figures 14–16.)



Figure 14: Interstitial moment within the *Los Atlantis* performance in L.A., including audience member with The Guide and images of audience members in the background.



Figure 15: Entry and Antechamber to the Archive.



Figure 16: Vignettes inside the Archive.

Discussion

Los Atlantis required engagement by the entire crew, performers, and even the audience with its systems—from both a storytelling and operational perspective. This departed from REMAP’s previous performance works, which were either conceived and created in tight collaboration with a small group of technologically savvy creators or were specific parts of a larger traditional work. [1]

Supporting this new level of engagement across all participants demanded both general approaches that leveraged familiar infrastructure (YouTube) and custom interfaces that were specific to *creating* the story (e.g., Task and Tagging websites). The former were fairly successful given a high level of familiarity amongst the public with common online and media platforms, suggesting an opportunity to incorporate them more extensively in the future if APIs are straightforward and accessible, which is not always the case. But, when coupled with contemporary media tools, such as the integration of access to YouTube into TouchDesigner, there was a useful and provocative bridge between online media at scale and local experimentation and rehearsal.

At the same time, there appear to be few contemporary tools for creating subsystems like the “active script” and its combination with the “show memory”. Together, these formed a living document that could be seen from one perspective as dramatic literature (to some extent) but also incorporated dynamic components that could select media on-the-fly and reference data from past rehearsals and performances. While there are *some* examples in research and commercial use, emerging from the rich tradition of hypertext authoring,

there are very few that are modern, robust, and focused on general support for authoring dynamic texts, especially at the level of the script itself.² To the authors’ knowledge, the ones that do exist are not easily integrated with contemporary web technologies.

Projects like *Los Atlantis* present many other challenges, e.g., how to support iterative development of code and story efficiently, avoiding congested workflows that require “rendering” the script and, instead, streamline experimentation, as done for media by integrating streaming video with a real-time processing engine. Our university audiences were quite open to participating and suspending disbelief, as long as the interfaces and tools met their standards for modern “apps.” However, it was difficult to engage them in advance with content creation that was within the story—e.g., submitting photos or contributing online to the Archive—since they were not yet engaged with the characters and the material. Further, the mobile Guide was often forgotten when it was not somehow critical to the progression of the performance. These challenges suggest more integration within the story is critical, heightening the importance of these elements as fundamental to the audience experience.

Acknowledgements

The authors thank the cast, crew, researchers, and other participants of *Los Atlantis*. The Los Angeles-based firm Creable created the Observatory and Guide website designs. The project was supported in large part by a Google Focused Award on the Future of Storytelling.

² This suggests a rich area for experimentation in HCI and other fields to develop such authoring tools and techniques.

REFERENCES

1. Jeff Burke, et al. 2016. Embedding expression: Pervasive computing architecture for art and entertainment. *Pervasive and Mobile Computing* 2.1: 1-36.
2. Jeff Burke and Jared J. Stein. Location-based and Audience-aware Storytelling. 2017. <http://howlround.com/location-based-and-audience-aware-storytelling-grace-plains-and-bodies-for-a-global-brain>.
3. Guy Debord. 1958. Theory of the Dérive. Internationale Situationniste. <http://www.cddc.vt.edu/sionline/si/theory.html>.
4. Nima Yousefi. An open source implementation of the Fountain screenplay formatting language. 2012. Retrieved September 19, 2016 from <https://github.com/nyousefi/Fountain>.
5. Lixia Zhang, et al. 2014. Named data networking. *ACM SIGCOMM Computer Communication Review* 44.3: 66-73.