Zócalo: A Service-Oriented Architecture Facilitating Sharing of Computational Resources in Interactive Narrative Research

R. Michael Young young@csc.ncsu.edu
James Thomas jmthoma5@ncsu.edu
Charles Bevan cfbevan@ncsu.edu
Bradley A. Cassell bacassel@ncsu.edu

Digital Games Research Center and Liquid Narrative Group
Campus Box 8206
North Carolina State University
Raleigh, NC, 27695

Why is it good to share code? There are at least three characteristics of software that arguably make it "an economically obvious thing to share." First, the Internet has made negligible the costs of copying, distribution, and coordination among geographically dispersed parties. Second, software "is a 'non-rival' good: i.e., your use of it does not interfere with my use. Third are the network effects, where additional users help the software become more usable, either indirectly as they advocate for useful changes to be made to the software, or in the case of open-source software, through actually extending the code in useful directions. As a small research community confronting a large and difficult problem, sharing can help us build more useful systems from our limited pool of creative resources by avoiding redundant coding. Over the past decade, the Liquid Narrative group at N.C. State has worked steadily to make our Interactive Narrative software more easily shareable. We developed and refined an architecture, called Zócalo, consisting of three optional, pluggable components for the 1) authoring, 2) planning, and 3) execution of interactive narratives on a variety of game engines, first described by Vernieri and Young (2006).

The Zócalo architecture integrates a suite of intelligent control tools with various commercially available 3D graphical game engines. While these engines are well-suited for building conventional 3D interactive game titles, the representation of the environments that they model don't match well with those typically used by AI researchers. Like most virtual world engines, their internal representations are *procedural* -- they don't utilize any formal or declarative model of the characters, setting or the actions of the stories that take place within them. Consequently, direct integration of intelligent software components is often not straightforward. To facilitate this integration, Zócalo overrides a game engine's default mechanisms for controlling its virtual environment, using instead a service-oriented architecture in which low-level control of the game environment is performed by a customized version of the game engine (via a lightweight client) and high-level reasoning about narrative structure and user interaction is performed remotely by an intelligent service. The service creates directives for action execution in a declarative, STRIPS-style (Fikes and Nilsson, 1971) language that is transmitted to the game engine's client via XML across a generic socket connection. The client then acts as a scheduler and execution monitor, translating the declarative representation of action directives into function calls responsible for initiating the corresponding engine-specific procedures.

The Zócalo system contains three sharable elements:

1. Longbow planning system. The Longbow planning system, first developed as a natural language discourse planner, has been used in a number of systems focusing on story and discourse generation in interactive narrative. Longbow's knowledge representation is readily extensible and has been adapted for use in modeling character intentionality in stories (Riedl and Young 2011), the generation of conflict between character plans (Ware and Young, 2011), and the planning of cinematic camera control coordinated with underlying story (Jhala and Young, 2010). Longbow is written in Allegro Common Lisp and is also available as a Windows DLL for inclusion in custom Windows software. Within Zócalo, Longbow inputs and outputs are specified in XML.

2. Bola, the client-side story manager. Bola provides network communication and XML parsing of the narrative plan produced by the planning system, coupled with an execution manager that schedules the

partially ordered actions of the plan, and a structured interface to guide the writing of action classes that correspond to the functionality in the virtual environment (e.g., the execution of animations, physics simulations, lighting changes) required by the corresponding declarative action representations used by the server's planning systems. This interface ensures that if the story plan encounters an error where an anticipated precondition or effect of an action is unfulfilled, an exception can be propagated to the core game code to allow it to decide how to handle the problem. Current implementations of Bola exist for the Unreal 3 and Unity3D 3.4 game engines.

3. Bowman, a tool for the authoring of story action libraries. Bowman assists authors in defining for the planner the library of actions and objects relevant to the story, as well as the initial state of the game world and the desired state of the world at the end of the story. Bowman reads and writes XML representations for all of these components of the story world to allow for reuse and extensions of these declarative representations for different stories. In addition, Bowman provides socket-based real-time communication with the planner to refine the author's depictions of the world based on the set of plans these depictions cause the planner to produce. Bowman provides richly annotated navigable graphs to depict both the details of individual plans as well as the relationships between various plan alternatives provided by the planner (Thomas and Young, 2006).

Zócalo promotes sharing and reuse of the intellectual capital required for interactive narrative in several ways. First, the planning, authoring, and game client components of Zócalo are decoupled to allow interchangeable alternative implementations of each. Thus, a given interactive narrative produced by Zócalo can be played in games based on the Unreal, Source, or Unity3D game engines; a given story can be authored through Bowman, an alternative tool, or by hand; and the action sequences in the plan could be produced by Longbow, an alternative planner or a case-based reasoning system. A second aspect of the Zócalo architecture that promotes sharing is generic, socket-based communication between the three components allows for each to run on its own computing platform. Finally, encoding all inter-component communications within XML allows the products of the authoring and planning components to be delivered contemporaneously with the plan execution, or stored and delivered later.

Acknowledgements

The authors thank Tommy Vernieri for his significant initial contribution to the development of Zócalo . Julio Bahamón, Stephen P. Cash, David Christian and Justin Harris have also contributed software to the project. The work was supported by NSF CAREER award 0092586 and Graduate Research Fellowship award DGE-0237090.

References

Fikes, R., and Nilsson, N. "STRIPS: A new approach to the application of theorem proving to problem solving." Artificial intelligence, 2, 3-4 (1971), 189–208.

Jhala, Arnav and Young, R. Michael, Cinematic Visual Discourse: Representation, Generation, and Evaluation, in *IEEE Transactions on Computational Intelligence and AI in Games*,volume 2, issue 2, pages 69 - 81, June 2010.

Riedl, Mark O. and R. Michael Young. Narrative Planning: Balancing Plot and Character. *Journal of Artificial Intelligence Research*, vol. 39, 2010.

Thomas, James M. and Young, R. Michael. Author in the Loop: Using Mixed-Initiative Planning to Improve Interactive Narrative, in *Working Notes of the International Conference on Automated Planning and Scheduling (ICAPS) 2006 Workshop on AI Planning for Computer Games and Synthetic Characters*, Cumbria, UK, 2006.

Vernieri, Thomas and Young, R. Michael, "Web services for interactive narrative", 2nd International North-American Conference on Intelligent Games and Simulation GAMEON-NA, Monterey, CA, USA, 2006, 13-17.

Ware, Stephen and Young, R. Michael, CPOCL: A narrative planner supporting conflict, in the proceedings of the Seventh International Conference on Artificial Intelligence and Interactive Digital Entertainment 2011, October, 2011, 91-102.