Event-Centric Control for Background Agents

Alexander Shoulson and Norman I. Badler

Department of Computer and Information Science,
University of Pennsylvania,
Philadelphia, PA 19104-6389, USA
{shoulson,badler}@seas.upenn.edu

Abstract. This paper describes a framework for controlling the varied activities of groups of background characters (representing extras or supernumeraries). Our platform is built upon an event-centric agent control model, which shifts behavior authoring from writing complex reactive agents to authoring particular activities. This approach allows us to achieve diverse, complex, and collaborative activities while the agents themselves stay simple and generic. An event is defined generically on agent roles, and can be dispatched to any set of agents that can fill those roles. This allows us to control macro-level group behavior with a centralized entity called the Group Coordinator that dispatches events to agents based on their situational and locational context (which can be controlled by an author). What results a structure for controlling macroscopic behavior for groups of background agents.

Keywords: Virtual Characters and Agents, Narrativity in Digital Games, Story Generation and Drama Management, Real-time Techniques for Interactive Storytelling.

1 Introduction

With few exceptions, the settings of interactive fiction and simulations are populated by autonomous individuals. Some of these individuals consist of bystanders, pedestrians, or other non-focal characters in the background. In traditional cinematography, these individuals are represented by "extras" – actors performing minor tasks that, on a macroscopic level, affect the mood, atmosphere, or illusion of realism. While a great deal of work has been put into simulating the positioning and locomotion of crowds, less attention has been paid to the performance of a group's coordinated activities within a spatiotemporal context [1, 4–6].

To make the background agents of a scene more useful, they must perform actions consistent with their surroundings. For instance, in an urban simulation, some agents could sit and read a book, but this action is more appropriate in a park than in a movie theater. Furthermore, agents must interact with one another in a coherent manner given their spatiotemporal context, such as having a conversation in the park, or sharing popcorn in the movie theater.

In order to simplify the interactions between agents, we centralize the control of multiple agents in a structure called an event. An event specifies the role each

M. Si et al. (Eds.): ICIDS 2011, LNCS 7069, pp. 193–198, 2011.

participating agent plays and the sequence of actions each role should perform on one timeline. For instance, a transaction event would involve a buyer and a seller, and direct each participant to play the correct animations and sounds in the proper sequence to simulate the exchange of goods. We see a similar system with Smart Events [10], but with some key differences. Smart Events notify subscriber agents, at which point the subscriber agent can decide the manner in which it responds, if at all, on an individual basis. Our events take full control of their involved agents, and suspend those agents' autonomy [8]. When an agent is under the control of a dispatched event, it follows the 'script' provided by the event's description. Events of varying priority can preempt one another, but an agent can only be a member of one event at a time.

This provides two advantages. First, once an event is dispatched, it requires no external oversight, and will properly direct each agent involved to behave consistently until complete. Second, since multiple agents are controlled by a single event, it is easy to describe what groups of agents are doing in the environment. By contrast, suppose we authored all conversations reactively in each agent, so that if an agent greets another, that other agent knows to reply to the first, and so on. It may be difficult to inspect those two agents and know that they are in a conversation, or what kind of conversation they are in, since they are each only responding to the stimuli the other provides. If the two agents were in a 'conversation' event, however, we could immediately inspect those agents, recognize that they were being controlled by a conversation event, and retrieve the parameters of that event for more information on its purpose or mood.

In this paper we present our event-based framework enabling centralized simultaneous absolute control of groups of agents with one logical structure. We then discuss a system that dispatches these events in order to direct groups of simple agents to perform collaborative actions with respect to the context of their environment. We show how the distribution of actions appropriate for a location can be controlled based on environmental factors and time. We then conclude with final thoughts and ongoing work.

2 Event-Centric Behavior Authoring

We cannot feasibly build a single Director that controls every aspect of every agent's behavior in a scalable fashion. Even by abstracting complex behaviors into simpler commands and suggestions on the part of the Director (as in [9]), it is impractical to expect a designer to author the logic of an entity for centralized, active control over every agent at all times for larger-scale groups. On the other hand, we do not want to increase the complexity of each agent, as large groups of significant, high-dimensional characters are difficult to model and coordinate [11]. Furthermore, relying on emergent reactive behavior in each agent to properly perform interactions offers limited control over the distribution of events in the environment. If a type of cooperative behavior only occurs when the situation is favorable and the agents decide to perform it (say, a conversation begins when two agents are facing one another), we cannot guarantee that that event will occur with a spatial distribution or temporal frequency, which may be desirable for expressing the overall mood of a crowd.

We balance this conflict with an event-driven architecture. Conceptually, events sit above groups of agents and below a top-level centralized Director [9]. The Director can concern itself with the problem of deciding which events, broadly speaking, should occur in the environment, and let the 'script' of each event handle the detailed execution of those particular collaborative actions. Since events take full control over the agents involved, the only individual behavior logic agents need is a set of actions to perform so as to look busy when not involved in an event. This keeps agents simple, since any reactions to other agents are handled by the event logic. With background characters we are mostly interested in the broad-scale behaviors that the agents can perform to make the scene feel alive, so we can afford to reduce the fidelity of the individual agents.

This approach is described in more detail in [8], which uses parameterized Behavior Trees that act generically on classes of agents. An event behavior tree, when instantiated, takes full control over the agents in its role fields and executes its internal 'script' on them, which can also take parameters to modulate its behavior (say, the mood of a conversation). When the event terminates, the involved agents return to their idle behaviors and wait to be selected for another event, sitting in a pool of available agents.

3 Ambient Event Selection and Population

To control ambient background agents as the "extras" of a scene, we use a dedicated entity we call a Group Coordinator. We make an important distinction between background characters and principal characters, the latter being pivotal story-driven agents that drive the narrative forward. Principal characters are usually controlled with a planning approach, using actions and preconditions as a dynamic script for the intended plot [3] or with global constraints based on actors' collaborative or competitive goals [2]. We are instead concerned with background agents whose sole purpose is to decorate a scene, so planning and long-term causality are unnecessary. Principal character agents and background agents can both be controlled with the same event (such as an event where a principal agent gives a speech to a cheering crowd of background agents), but a detailed explanation of the interactions between these two categories of agents and the player is outside the scope of this discussion.

The Group Coordinator's sole purpose is to enforce author-specified spatiotemporal distributions of events. Each event is given a category when written and placed in a database. Then, for each location and general time period in the environment, the scenario author specifies a desired distribution over each event category. The Group Coordinator is responsible for enforcing that distribution. If the distance between the desired event distribution and the actual distribution (that is, a normalized count of the currently active events' categories) exceeds a specified threshold, the Group Coordinator identifies the most severely underrepresented event category, selects an event within that category (by some criteria),

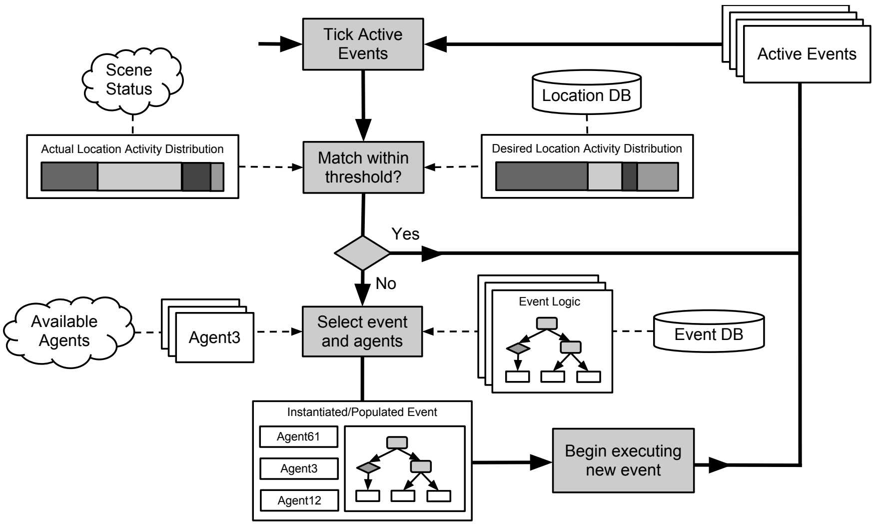

Fig. 1. A conceptualization of the Group Coordinator control process for background groups

and dispatches that event to some agents to perform. The Group Coordinator's decision process is illustrated in more detail in Fig. 1.

The expected distribution of events in a location can change over time or according to some situational factors. An adjustment to the distribution of events for a location can represent effects such as a day/night cycle (a temporal adjustment), or a slow mood shift based on some narrative progression (a situational adjustment). These changes are gradual, since they only affect the selection of new events and have no bearing on events in progress. For instant reaction, such as responding to a fire, it is better to place all of the relevant agents in a high-priority event that instructs each agent to flee (or assist, depending on some decision).

The process of selecting agents to fulfill the roles in an event is more complex. Preliminary work on this issue is discussed in [7], but we will briefly address it here. Agents can be selected based on any number of criteria, including location, age, occupation, and so on. The degree of constraint on the agent selection process influences how readily the Group Coordinator can function. Events that place more restrictions on the selection of agents for its roles will prove more difficult to populate, as it is less likely that agents will be available to meet these criteria. If the event database is populated with a large number of very exclusive events, then the Group Coordinator will have difficulty instantiating events, and will therefore be less able to enforce the desired distribution of events for the environment.

It may be appealing to enforce constraints such as "I cannot purchase an item if I have no money", so that no agent acts inconsistently after prolonged examination. However, the player also has the expectation when observing a shopping mall that there be agents involved in transactions. If no agent can purchase because they have all depleted their monetary allowance, then the player's observation of a scene is inconsistent with her expectations. We see individual agents as less important than the groups they comprise, so we forego some agent-level constraints for the sake of whole-group behaviors.

4 Conclusions

Adding background agents to a scene is important for narrative; well-populated areas represented in a story are less convincing if only a few simulated actors inhabit them. Individually, however, these agents are not meaningful to the player, and so we have no reason to personify them beyond their specific role in the scene. Since ambient characters may only play one or two minor roles in a scene before they are discarded, we do not want to waste a complex behavioral repertoire on each of them. At the same time, we do not want to author a diverse pool of simple agents that are specialized to certain ambient tasks.

Event centrism provides a means to retain agent action diversity while mitigating individual agent complexity. When complex behaviors, especially interactions and coordination between agents, are abstracted in a centralized event structure, behavioral repertoires for agents expand when needed. Any convenient agent can be selected to perform a complex sequence of actions while being neither unduly sophisticated nor unnecessarily specialized. Furthermore, we can describe and control the world based on what is happening in it, rather than just who or what is in it. Events capture and denote activity in the world.

As alluded to in Sec. 3, we would like to author narrative progression as sequences of events as well. Background agents are not given narrative qualities (like a backstory) and are not integral to the plot, but they can be used to react to the actions of principal plot characters. An event could easily coordinate both a principal character giving a speech on a podium, and a crowd of background agents gathering and responding to the oration. Events provide an opportunity to develop a story in the domain of plot occurrences.

Currently this framework is being implemented in what we call the Agent Development and Prototyping Testbed (ADAPT), which features toolkits for agent modeling, behavior authoring, and full simulation on top of the Unity game engine. We expect to evaluate the degree to which an event-centric agent control methodology simplifies the task of adding multi-agent behaviors to the environment. Since event centrism allows direct control over the spatiotemporal distribution of behaviors across agents, we expect groups controlled in this manner to act more believably on a macroscopic level than groups of agents authored solely with agent-centric reactive control logic. Understanding the consequences of these design decisions will allow us to better balance the degree to which we depend on either agent sophistication or event-centrism, and how specifically we design our events with respect to agent selection and roles.

Acknowledgements. The research reported in this document/presentation was performed in connection with Contract Number W911NF-10-2-0016 with the U.S. Army Research Laboratory. The views and conclusions contained in this document/presentation are those of the authors and should not be interpreted as presenting the official policies or position, either expressed or implied, of the U.S. Army Research Laboratory, or the U.S. Government unless so designated by other authorized documents. Citation of manufacturers or trade names does not constitute an official endorsement or approval of the use thereof. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation heron.

References

1. Farenc, N., Boulic, R., Thalmann, D.: An Informed Environment Dedicated to the Simulation of Virtual Humans in Urban Context. Computer Graphics Forum 18(3), 309–318 (1999)
2. Kapadia, M., Singh, S., Reinman, G., Faloutsos, P.: A Behavior-Authoring Framework for Multiactor Simulations. IEEE Computer Graphics & Applications (2011)
3. Li, B., Riedl, M.O.: Creating Customized Game Experiences by Leveraging Human Creative Effort: A Planning Approach. In: Dignum, F. (ed.) Agents for Games and Simulations II. LNCS, vol. 6525, pp. 99–116. Springer, Heidelberg (2011)
4. Pelechano, N., Allbeck, J.M., Badler, N.I.: Controlling Individual Agents in High-Density Crowd Simulation. In: Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA 2007, pp. 99–108 (2007)
5. Pelechano, N., Allbeck, J.M., Badler, N.I.: Virtual Crowds: Methods, Simulation, and Control. In: Synthesis Lectures on Computer Graphics and Animation. Morgan & Claypool Publishers (2008)
6. Peters, C., Ennis, C.: Modeling Groups of Plausible Virtual Pedestrians. IEEE Comput. Graph. Appl. 29, 54–63 (2009)
7. Shoulson, A., Garcia, D., Badler, N.I.: Selecting Agents for Narrative Roles. In: Proceedings of the 4th Workshop on Intelligent Narrative Technologies (2011)
8. Shoulson, A., Garcia, F., Jones, M., Mead, R., Badler, N.I.: Parameterizing Behavior Trees. In: Proceedings of the Fourth International Conference on Motion in Games, MIG 2011 (2011)
9. Si, M., Marsella, S.C., Pynadath, D.V.: Directorial Control in a Decision-Theoretic Framework for Interactive Narrative. In: Iurgel, I.A., Zagalo, N., Petta, P. (eds.) ICIDS 2009. LNCS, vol. 5915, pp. 221–233. Springer, Heidelberg (2009)
10. Stocker, C., Sun, L., Huang, P., Qin, W., Allbeck, J.M., Badler, N.I.: Smart Events and Primed Agents. In: Safonova, A. (ed.) IVA 2010. LNCS, vol. 6356, pp. 15–27. Springer, Heidelberg (2010)
11. Yu, Q., Terzopoulos, D.: A Decision Network Framework for the Behavioral Animation of Virtual Humans. In: Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA 2007, pp. 119–128 (2007)